



THE SCIENCE,
PSYCHOLOGY & PHILOSOPHY OF GAMBLING

SQUARES
& SHARPS,
SUCKERS
& SHARKS

JOSEPH BUCHDAHL

JOSEPH
BUCHDAHL

SQUARES
& SHARPS,
SUCKERS
& SHARKS

THE SCIENCE, PSYCHOLOGY
& PHILOSOPHY OF
GAMBLING

HIGH STAKES

SQUARES & SHARPS, SUCKERS & SHARKS

People have been gambling, in one form or another, for as long as history itself. Why? Money, entertainment, escape and a desire to win are all traditional explanations. Arguably, however, these are secondary considerations to a higher order purpose: a craving for control. Gambling offers a means of gaining authority over the unknown, granting us a sense of control over uncertainty. Almost always that sense is illusionary – gambling, including betting and investing, is essentially random – yet for many it is nonetheless profoundly rewarding. This book attempts to explore the reasons why. Along the way, it examines:

- The science of probability and uncertainty
- Why gambling is often condemned
- The difference between expectation and utility
- The irrationality of human beings
- Evolutionary perspectives on gambling
- Luck and skill
- Market efficiency and the wisdom of crowds
- Why winners take all
- Cheating
- Why the process matters more than the outcome

Since 2001 Joseph Buchdahl has been providing quantitative football and tennis data for betting analysis, and independent verification for sports betting advisory services. He is also the author of *Fixed Odds Sports Betting* and *How to Find a Black Cat in a Coal Cellar*.

Contents

[God Does Play Dice](#)

[Cleopatra's Nose](#)

[To Gamble or not to Gamble: is there a Question?](#)

[The Three Rs: Risk, Reward and Rationality](#)

[The Harder I Work, the Luckier I Get](#)

[Monkeys Throwing Darts](#)

[Ginsberg's Theorem](#)

[Winner Takes All](#)

[A Market for Lemons](#)

[The Fox and the Hedgehog](#)

[Bibliography](#)

[About the Author](#)

[Copyright](#)

GOD DOES PLAY DICE

Albert Einstein once famously said that God does not play dice, expressing his contempt for the idea that the universe is governed by probability and believing instead that everything is causally deterministic. According to 19th century determinism, if someone could know the precise location and momentum of every atom in the universe, their past and future values for any given time could then be calculated from the laws of classical mechanics. Laplace's Demon, as this thought experiment became known, has provided the beacon of hope to all gamblers that it is fundamentally possible to predict the future. Sadly, quantum mechanics, the science of the 20th century, demonstrated that both Einstein and Laplace were wrong. Not only does God play dice, but he doesn't know what the outcome will be.

The quantum mechanical world of the atom may not, at first sight, have a great deal to do with the spin of a roulette wheel, predicting the outcome of a football match or the value of a share, although more than one might imagine, as we shall see. Yet the significance of the distinction between these two ideas of determinism and uncertainty lie at the very heart of understanding the science of gambling and the psychology of gamblers. Human beings love to find patterns; indeed, they've evolved that way (because pattern recognition is cognitively less energy-intensive). And they love to find causal explanations for those patterns, even when none actually exists. Randomness, by contrast, is not a concept easily understood and embraced, but failure to do so ensures that the majority of gamblers, including even those in the arenas of sports and finance where theoretical advantages exist, find themselves on the wrong side of the profit line. Furthermore, almost all of those who do make money from such gambling markets do so purely by chance.

This is not an idea that most gamblers find palatable, since it has implications for the very reasons why we choose to gamble in the first place. Gambling is connected to an intrinsic desire to control one's destiny, to manipulate luck in order to validate and find meaning in life. Gambling,

it turns out, is as natural as a faith in God, and for more or less the same reasons. No wonder, then, that those of a more religious persuasion, both past and present, have attempted to condemn it as something immoral. If all (or almost all) of gambling, including sports betting and financial investing, is just uncontrollable chance, what, then, is the point of it?

Spoiler alert: this book will not provide you with a winning system. On the contrary, having read it you will understand why, if I had made such a claim, it would probably no longer be valid. My intention, then, is not to help you become a more profitable gambler but rather, hopefully, a wiser one, through a deconstruction of three core areas associated with gambling: its science, psychology and philosophy. In doing so I hope to explore the reasons why some of us gamble, why others condemn it, why still others exploit it for selfish intentions, why most of us lose whilst a few winners take all, and finally why gambling, or at least the way some gamblers think, might actually be good for our decision making.

Whilst I will be examining various domains of gambling, including games of pure chance (at the casino) as well as games that theoretically offer an element of skill (poker, sports and the world of finance), my background as a sports data analyst predicates that much of my material will focus on betting. In particular, I will be using data that I have collected over the past 14 years to investigate why so few sharps¹ actually manage to beat the market, and why the remaining squares are really just randomly chucking darts. Following this, I will also review a few examples of the shady practices that take place in the world of gambling, exploring some of the reasons why sharks might choose to prey on suckers and why the latter allow themselves to fall victim. Finally, I will conclude by examining what makes a good gambler, and why when faced with decision making under uncertainty, it pays to focus more on the process than the outcome.

In writing this book, I have adopted a multidisciplinary approach, taking the reader on an explorative journey into domains as varied as economics, behavioural and evolutionary psychology, neuroscience, quantum mechanics, chaos and complexity theory, game theory, history and ethics, as well as the more familiar territory of probability upon which all of gambling hinges. With that in mind, let's begin this journey by first delving into the world of uncertainty, and an investigation into the length of Queen

Cleopatra's nose.

¹ Whilst the term ‘sharp’ has at certain times been used to describe players who exploit others in games of chance, for example ‘card sharp’, here I define a ‘sharp’ player as a gambler with a positive expectation acquired through something more than chance, whilst the term ‘shark’ is reserved for those who intentionally prey on others, the ‘suckers’ (who fall for the sales pitch), for their own financial gain. Finally, ‘squares’ are considered players who have no positive expectancy and are merely winning and losing as a consequence of luck.

CLEOPATRA'S NOSE

Blaise Pascal, a 17th century French mathematician and one of the founding fathers of probability theory, once famously remarked: “*Cleopatra’s nose, had it been shorter, the whole face of the world would have been changed.*” Had her nose been smaller, he hypothesised, she would have lacked the dominance and strength of character which a large nose in the Egyptian first century BC epitomised. As a consequence, Julius Caesar and Marc Antony would not have fallen under her spell, wars would not have been fought, and today we might all be speaking Latin. The ‘Cleopatra’s Nose’ theory is basically the proposition that chance has a massive role to play in the evolution of history. And so, of course, it does in gambling.

We have probably all had similar ‘Cleopatra’ insights, thinking about how things might have happened differently given tiny changes to insignificant starting points. If Steven Gerrard had woken up a second later than he did on that fateful day in April 2014 when Chelsea beat Liverpool, would he still have slipped over? If Mark Robins hadn’t scored his 56th minute goal against Nottingham Forest in the 3rd round of the FA Cup on 7 January 1990 would Manchester United have won 13 Premiership titles and would Alex Ferguson have been knighted?

Pascal’s thought experiment laid the foundations for what would ultimately come to be known as chaos theory. We’ll consider how this theory, more commonly known as the butterfly effect, has implications for the success of our predictions about the future; but first, a brief history of probability. Ironically, it all began with gambling.

A Brief History of Probability

Probability, the subject matter that defines all of gambling, did not gain any rigorous academic attention until the 16th century when the Italian mathematician Gerolamo Cardano developed the first statistical principles, and in particular the notion of odds as the ratio of favourable to

unfavourable outcomes, thereby expressing probability as a fraction (the ratio of favourable outcomes to the total number of possible outcomes), a concept that is still used by bookmakers and casinos today. Critically, Cardano recognised the significance of possible combinations that contribute to a ‘circuit’ – the total number of possible combinations. For example, when throwing a pair of 6-sided dice, he recognised that there are not 11 but 36 possible outcomes. Yet Cardano may never have realised what he was on the verge of discovering. Indeed it remains unclear whether he developed his elementary rules of probability for the purposes of gambling – he was a consummate gambler – or for the purposes of defining a new theory of mathematics. This task fell to two French mathematicians, the first of whom we have already met at the start of this chapter.

In 1654 Blaise Pascal was asked by his friend Chevalier de Méré to consider the problem of points. The problem of points concerned a game of chance, called balla, where two players had equal chances of winning a round. Each player contributed equally to a prize pot, and agreed in advance that the first player to have won a certain number of rounds would collect the entire prize. Chevalier de Méré asked Pascal to consider how a game’s winnings should be divided between two equally skilled players if, for some reason, the game was ended prematurely. Originally considered in 1494 by another Italian mathematician, Luca Pacioli, the problem remained unsolved, even by Cardano. Pascal decided to correspond with his friend and colleague Pierre de Fermat (famous for Fermat’s last theorem) on the matter. The work that they produced together signalled an epochal moment in history, defining a new field of mathematics: probability theory. In doing so they introduced the concept of mathematical expectation or expected value, understood by every gambler with more than a passing interest in numbers.

Given that human beings have been playing games of chance for many thousands of years, it is perhaps surprising that it took so long for the subjects of probability and randomness to be considered formally at all. Undoubtedly, the equivalence most societies and cultures prior to the Enlightenment had perceived between chance and pre-ordained divination according to God (or gods) accounts for much of the explanation. Yet the ancient Greeks, being more intellectually enlightened than most of the 2,000 years that followed them, also ignored the problem. Despite

understanding that more things *might* happen in the future than actually *will* happen, they never chose to formalise this mathematically. In all probability (pun intended), the reason was that the Greeks had little interest in experimentation and proof by inductive inference; they preferred proof by logic and deduction instead. By contrast, the Enlightenment heralded a birth of a new freedom of thought, a passion for experimentation and a desire to control the future.

Pascal was also a deeply religious man, and he reconciled his new theory of probability, and the propositions it advised for unfinished games of balla, as a matter of moral right. Other exponents of probability theory, such as Jacob Bernoulli, a 17th century Swiss mathematician, would also blur the distinction between mathematics and morality. As such, how wagers in games should be settled, and how value should be assigned to their stakes, came to be understood in terms of religious morality and Divine will. Indeed, even one of Adam Smith's defining works that marked the birth of capitalism was named the *Theory of Moral Sentiments*.

Pascal used his new mathematics to pose a question, which has become known as Pascal's Wager: "*God is, or He is not. But to which side shall we incline? Reason can decide nothing here.*" Which way we should wager will be defined by four propositions: 1) you bet that God exists and he really exists – infinite gain; 2) you bet that God doesn't exist but he does exist – infinite loss; 3) you bet that God exists and he doesn't exist – finite loss; and finally 4) you bet that God doesn't exist and he doesn't – finite gain. Essentially, Pascal was asking us to consider the relative value of the cases where God does and does not exist, even if it happens that the distinction represents a 50-50 proposition. The answer, to Pascal at least, was obvious: why risk eternal damnation betting against God, when betting for God, through means of living a pious life, involves a considerably smaller outlay, regardless of whether God exists or not. As such, Pascal's Wager represented the beginnings of behavioural decision theory, or the theory of decision making under uncertainty, which Daniel Bernoulli, Jacob's nephew, would advance during the following century.

Moral Certainty

Thus far, probability theory had concerned itself merely with games of chance, where the probabilities of possible outcomes could be calculated *a priori* from mathematical principles. Such mathematics is pretty much all that is required for a casino offering games such as roulette, craps and keno to manage its liabilities (particularly an online casino that won't suffer from the vagaries of imperfect roulette wheels and dice), since expected values for all these games can be calculated exactly.

In 1703, two years before his death, Jacob Bernoulli wrote to his friend Gottfried Leibniz, a German mathematician and philosopher (famed for the development, alongside Sir Isaac Newton, of calculus) commenting on the oddity that we can know the odds of rolling a five rather than a three with a pair of dice, and yet are unable to precisely calculate the chances that a man of 20 will outlive a man of 60. In a stroke, in making a crucial distinction between reality and abstraction, Jacob had identified the (moral) conundrum that has plagued speculators of sports and finance ever since. Many outcomes, and more importantly outcome expectancy, cannot be known with perfect precision.

Jacob Bernoulli wondered whether the problem might be solved by examining a large number of pairs of each age. In doing so, he was implicitly recognising that the past must provide some key to predicting the future. Leibniz was not impressed: "*Nature has established patterns originating in the return of events, but only for the most part.*" For Leibniz, a finite number of historical observations would inevitably provide too small a sample from which to formalise a mathematical generalisation about nature's intentions. Jacob's response provided a revolution in statistics. His intellectual leap was to be the first to attempt to measure and define uncertainty, and in doing so calculate a probability empirically via inductive inference that a particular value lies within a defined margin of error around the true value, even when that true value remains unknown. For Jacob, probability was a degree of moral certainty and differed from absolute certainty as the part differs from the whole.

As such, Jacob Bernoulli's method of inductive inference involves estimating probabilities from what happened after the event, that is to say, *a posteriori*. For his solution to work, it requires one key assumption: under similar conditions the occurrence or otherwise of an event in the future will follow the same pattern as was observed in the past. Jacob recognised the

significance of the limitation this assumption implied, and in doing so revealed the uncertain nature of the world we live in.

Jacob Bernoulli's work on *a posteriori* estimation of probabilities led to his formulation of the law of large numbers. Frequently confused by gambling squares with the law of averages, the law of large numbers states that, as a sample size of independent trials (for example coin tosses) grows, its average should move closer and closer to the expected value. A key word here is 'independent'. In roulette, for example, each spin of the wheel is independent of the previous one, and its outcome has no memory of the last. The probability of the ball landing on red occurring after 3, 5, 10 or any number of consecutive blacks remains 50% (discounting the effect of the zero or zeros). Misunderstanding of this law has cost many a gambler dear. On 18 August 1913 at the Monte Carlo Casino, the roulette ball landed on black 26 times in a row with a probability of 1 in 136,823,184². Of course, one should remember that every other sequence of red and blacks (and zeros) was just as likely, but for human beings programmed to see and interpret patterns, far less memorable. Gamblers lost millions incorrectly believing that, according to the erroneous interpretation of the law of averages, a red must surely be more likely to appear after successive increases in the sequence of consecutive blacks to restore the balance of randomness. Unsurprisingly, the gambler's fallacy is also known as the Monte Carlo fallacy. It is probably the most frequently expressed fallacy in all of gambling.

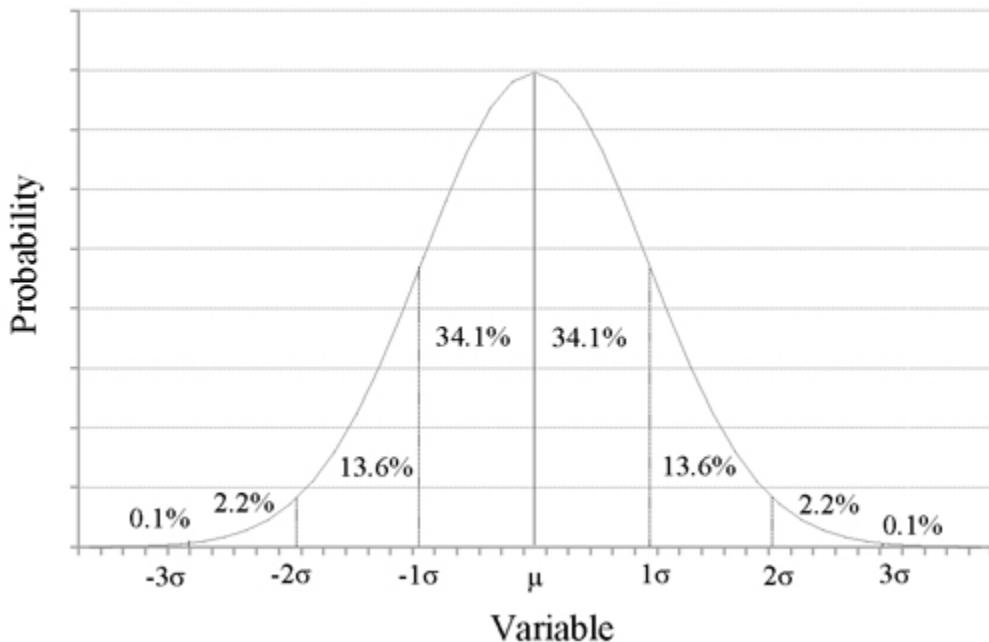
Jacob Bernoulli illustrated his law of large numbers by means of a hypothetical urn filled with 3,000 white pebbles and 2,000 black pebbles. Initially, this ratio is unknown to us. Our task is to estimate it through the process of iteratively withdrawing and replacing the coloured pebbles, each time noting the colour. The larger the number of pebbles we draw, the nearer we should expect the ratio of drawn white and black pebbles to approach 3:2, the true ratio. Jacob calculated that it would take 25,550 drawings to demonstrate a moral certainty with 1 part in 1,000 that the result we should obtain would lie within 2% of the true ratio. Jacob clearly demanded a high price for moral certainty. Others may well have accepted 'truth' long before. Indeed, acceptance of a scientific hypothesis reliant on similar proof by statistical inference requires a moral certainty of just 1 in

20. Doubtless, there will be explanations for this weaker insistence on moral truth, but the consequences will be far reaching; a lot of what is claimed as scientific evidence will be nothing more than meaningless statistical association arising by chance. For that matter, a lot of people who claim to be able to beat the financial market or to be able to predict the outcome of sporting contests actually fail to demonstrate a meaningful standard of moral certainty when subjected to proper scrutiny. We will return to that in later chapters.

Normality

Another of Jacob's nephews, Nicolaus Bernoulli, continued his uncle's work on probability theory and the estimation of uncertainty. Whilst Jacob calculated the number of trials one would require to define the error between an observed value and a true value, Nicolaus chose to start from the other end; given a fixed sample of observations, in this case the ratio of male to female births, what is the probability these would fall within a specified margin of error? Another French statistician, Abraham de Moivre, turned his attention to how well Nicolaus' samples represented the world from which they were drawn. De Moivre was already familiar to the gambling fraternity, with his publication in 1718 of *The Doctrine of Chances*, the first serious textbook on probability theory. Indeed, the first edition had the subtitle: *a method for calculating the probabilities of events in play*. De Moivre observed that establishing moral certainty via Jacob Bernoulli's experimental method of counting would be so laborious as to be of little practical use. Solving the problem by combining both calculus and the binomial theorem³, de Moivre observed how a set of random samples would distribute themselves about an average value. The larger the number of samples he observed, the smoother the shape of that distribution became. In effect, he expanded the binomial distribution to the infinite limit and discovered the normal distribution curve, with its own mathematical expression. Students of high school mathematics will remember its bell-shape, with many observations clustered around the mean and fewer further away.

The Normal Distribution



Whilst de Moivre's normal distribution couldn't calculate the precise chance that a man of 20 will outlive a man of 60, it could answer the question: if the true chance is assumed to be a particular number, what is the probability that our observations of the longevity of men aged 20 should occur. In effect, de Moivre was one of the founding fathers of statistical hypothesis testing.

De Moivre's mathematics allows us easily to determine when a set of data is normally distributed by means of its standard deviation, a measure of the spread or variance of the data within the distribution. When observations are normally distributed, those values less than one standard deviation away from the mean account for just over 68% of the data set; two standard deviations from the mean account for about 95%; and three standard deviations account for over 99%. The normal distribution is immensely powerful as it helps to define instances of real world phenomena consisting of independent observations that occur simply by chance. A beautiful illustration of this can be seen by means of a quincunx machine, originally devised by Sir Francis Galton in 1889. Many are available online⁴. Normal distributions are more improbable when observations are path dependent, that is to say, the probability of the next one occurring is

dependent on, or causally determined by, the previous one. In the absence of path dependency, it's usually a pretty safe bet that the phenomenon we are observing is random. That is to say, it has no cause. Not that de Moivre interpreted it that way; he was so astonished by the orderliness of randomness that he attributed it to Divine Providence, or in his words *Original Design*.

Many worldly phenomena find themselves normally distributed, for example intelligence, height, weight, blood pressure and many other physical and genetic characteristics that show no systematic differences across populations, life expectancies (for humans as well as batteries), annual crop yields and rainfall, batting averages in major league baseball and, much to the disappointment of a perennial stream of deniers, so also most of the daily movement of stock prices. A random process essentially means it has no memory. Without a memory, how can future observations possibly be predicted from preceding ones, and perhaps more importantly, how can we hope to make a profit?

A corollary of the normal distribution is that more extreme variables will tend to move closer to the average on subsequent measurements. The phenomenon was first uncovered by Sir Francis Galton, the Victorian polymath, as he experimented with his quincunx machine and the heredity of sweet peas. In cross breeding trials, Galton noted a tendency for the size of the offspring to show a smaller (but still normal) distribution than that of the parents. Crucially, whilst the offspring of larger parents tended to be smaller, the offspring of smaller parents tended to be larger. Galton described this tendency as reversion or regression to the mean. It is important to realise that there is no requirement for any teleological cause for this regression in a strictly deterministic sense, merely a random process that sees extremes become less extreme. As if to demonstrate this point, paradoxically, regression to the mean is not time dependent; if subsequent measurements are more extreme, the tendency will be for their earlier ones to be closer to the average. Regression to the mean, then, is entirely reversible.

Crucially, this principle informs us not that things **must** return to the average, just that they have a **tendency** to do so. After each successive black on that fateful day in Monte Carlo, there remained the tendency that the overall sequences of reds and blacks would revert towards the average

of 50-50, but this did not imply that it had to. Roulette balls don't have memories; they simply obey the laws of probability. 'What goes up must come down' is as much a fallacy as a belief in the law of averages. What goes up has a tendency to come back down, but it doesn't have to, nothing is making it do so. As Jordan Ellenberg clarifies in *How Not to be Wrong: the Hidden Maths of Everyday Life*, the law of large numbers works not by balancing out what's already happened, but by diluting it with new trials.

It is easy to see how gamblers might make incorrect inferences about patterns they perceive as offering the potential for profitability, if they fail to consider the implications of regression to the mean. An increase in the price of a mutual fund or an upturn in the fortunes of a football team might easily be misconstrued as having causal explanations when in fact they represent nothing more than statistical quirks. Considerable research into the financial markets has demonstrated evidence of regression to the mean. One particular example is noteworthy. On 1 April 1994 Morningstar, the investment research and management firm known for its ratings of mutual funds, published the performance of a basket of mutual fund categories for two five-year periods, comparing the five years to March 1989 with the subsequent five years to March 1994. All funds above the mean in 1989 (13.6% growth) were below the mean in 1994 (13.1% growth) and vice versa. International stocks, for example, had grown by 20.6% in the five years to March 1989, contrasted with just 9.4% in the subsequent five years. Small Company funds on the other hand underperformed the market to 1989 with a growth of 10.3% but managed to outperform it over the following five years, seeing 15.9%.

So what's an investor to do if such movements demonstrate little more than a random walk underpinned by regression to the mean? Well, 'buy low, sell high' may be excellent folklore advice in this context. Indeed, such a contrarian strategy may account for much of the success experienced by legendary investors such as Warren Buffett. The problem, of course, is knowing when low is low and high is high. In the real world of finance, regression won't manifest itself as a simple linear trend to smooth out extremes. Regression will be dynamic, sometimes overshooting, sometimes undershooting, fluctuating around a mean which itself will not necessarily be stable, such that normality itself is an ever-changing benchmark.

Another example from the world of sport exemplifies regression to the

mean beautifully: the new manager effect. The new manager effect concerns the idea that new football managers appear to improve the success of a football club relative to its performance under the old manager just prior to his sacking. The data on that appear pretty conclusive. Analysing managerial turnover across 18 seasons (1986 to 2004) in the Dutch premier division, Bas Ter Weel⁵ revealed noticeable patterns of prior decline and subsequent improvement centred on the sacking of one manager and the appointment of a new one. Crucially, however, almost the same pattern could be observed where managers had not been sacked. How so? Ter Waal was unequivocal in his explanation: *“If managers do not matter for differences in performance across firms and quality does not vary across managers, the only observed performance change following turnover would be mean reversion.”* David Sally, co-author of *The Numbers Game: Why Everything You Know About Football is Wrong*⁶, emphasises the point:

“In the same way that water seeks its own level, numbers and series of numbers will move towards the average, move towards the ordinary. The extraordinary... is followed by the ordinary... the ordinary is what happens. The average is what happens more often than not.”

Ter Waal’s research has been replicated for other football leagues, most particularly in Germany and Italy.

Laplace’s Demon

In one sense, the development of probability theory throughout the Enlightenment was at odds with the pervading culture of the 17th and 18th centuries. The Age of Reason, personified in Sir Isaac Newton and codified in his famous laws of motion and gravitation, had ushered in a new era of scientific determinism. If probability theory was describing a world of chance and randomness, what use was it when it came to ascribing effects to prior causes to explain and predict why it is that things happen? Of course, the forefathers of probability theory were still very much grounded in scientific rationality, and considered their new mathematics as offering valuable tools with which to make predictions about the future. We have already observed how de Moivre submitted to the power of ‘Original Design’, an epistemological position echoing back to earlier ideas of Divine

Predestination that had been subsumed during the Enlightenment. Jacob Bernoulli, too, believed that if “*all events from now through eternity were continually observed (whereby probability would ultimately become certainty), it would be found that everything in the world occurs for definite reasons.*”

19th century polymaths were cast under the spell of scientific determinism too. Henri Poincaré, a French philosopher, physicist and mathematician, insisted that chance is only a measure of our ignorance.

“Every phenomenon, however trifling it be, has a cause, and a mind infinitely powerful, and infinitely well-informed concerning the laws of nature could have foreseen it from the beginning of the ages. If a being with such a mind existed, we could play no game of chance with him; we should always lose.”

Furthermore, in a world of cause-and-effect, Poincaré insisted, we can invoke the laws of probability to make predictions about future stock prices, the value of life insurance policies and even the weather.

Perhaps the most significant and earliest articulation of scientific determinism can be attributed to Pierre-Simon Laplace, a French astronomer and mathematician, who in 1814 published the following postulate which subsequently became known as Laplace’s Demon.

“We may regard the present state of the universe as the effect of its past and the cause of its future. An intellect which at a certain moment would know all forces that set nature in motion, and all positions of all items of which nature is composed, if this intellect were also vast enough to submit these data to analysis, it would embrace in a single formula the movements of the greatest bodies of the universe and those of the tiniest atom; for such an intellect nothing would be uncertain and the future just like the past would be present before its eyes.”

Clearly we can see where Poincaré took his inspiration from. Evidently, Laplace did not believe in luck. Indeed, he was convinced there was no such thing. Put simply, Laplace and Poincaré were arguing that everything happens for a reason, and provided we (or our demon) know enough about the initial conditions, it should simply be a question of mathematics to be able to predict how and why they happen; music to the ears of every financial investor and sports bettor, no doubt.

Nevertheless, both Laplace and Poincaré appear to have held reservations. For his part, Laplace warned against the tendency to assign a particular cause to an outcome when in fact only chance was at work. In

doing so, he was unmistakably aware that all of us are prone to find significance, or as Jacob Bernoulli would put it, moral certainty, in patterns that undeniably have no meaning at all. 26 consecutive blacks on a roulette wheel is clearly a pattern that conjures all sorts of emotional responses, and for some misguided wagering. A random series of some reds and blacks over 26 wheel spins elicits no such response, indeed it would never be consigned to memory at all. And yet both sequences are just as probable (or improbable), and just as random as each other.

Perhaps more significantly, Poincaré understood that sometimes the distinction between randomness and determinism becomes blurred. Some events that appear to be lucky are in fact deterministic, but slight variations in the initial conditions change the evolution of successive cause-effect iterations such that the final outcome may bear no resemblance to another with a similar, but slightly different, starting point. In uncovering this sensitivity to initial conditions, Poincaré indicated that randomness and determinism appear distinct only because of long term unpredictability. A very small cause, which eludes our capacity to analyse, determines a considerable and observable effect; hence we say that it is due to chance. As such, prediction becomes impossible and we have a random phenomenon. This was the birth of chaos theory. Laplace, it would appear, was right: luck is merely evidence of incomplete knowledge.

To most people, chaos theory is more popularly known as the butterfly effect. The name of the effect, coined by Edward Lorenz, a 20th century American mathematician and pioneer of chaos theory, is derived from the metaphorical example of the simple flapping of the wings of a butterfly somewhere in the world influencing the outcome of a major weather system a couple of weeks later somewhere else on the planet. Exemplifying Poincaré's sensitivity to initial conditions, we can reason that, had the butterfly flapped its wings in a slightly different manner, the successive perturbations to the air around it, and subsequently to the wider atmosphere at large, manifested through a process of non-linear feedback, would result in a completely different weather pattern a couple of weeks hence. It is for this reason that Poincaré explained why meteorologists had such limited success making weather forecasts.

Essentially chaos theory reveals that often we have too little information to apply the laws of probability, and even if we try we can never be

absolutely certain about causation. This takes us back nicely to de Moivre's samples and his normal distribution. No matter the quality of our sample data we can never extrapolate with 100% certainty what they inform us about the underlying 'truth' of the population. The best we can do is infer that a hypothesis under scrutiny should either be rejected or not rejected, but never accepted with absolute certainty. Today this is known as the principle of falsifiability.

It requires little effort to transform a simple linear system into a chaotic unpredictable one. Consider for example a simple pendulum and start it swinging. Its motion will be perfectly described by Newton's laws of motion. Given knowledge about the length of the pendulum and the strength of the gravitational force influencing its motion, I will be able to predict its velocities and positions at any time in the future. Now let's add a second pendulum to the bottom of the first by means of a second fulcrum. This time, the motion of the pendulums very quickly become unpredictable and chaotic, and any attempt to try to replicate the initial starting position to repeat a series of oscillations becomes an impossible task.

It is easy, then, to see how even fairly simple systems can quickly become chaotic. Slight differences in the way a snooker player might strike the cue ball could very quickly lead him to losing position. Indeed, it has even been estimated that the gravitational pull of an electron on the other side of our galaxy may have an influence, through non-linear feedback, on the outcome of a game of snooker. If we happen to be betting on him winning the frame or the match, chaos theory is something that is going to have a significant impact. In team sports, where numerous players are interacting for long durations, the potential for chaos to wreak havoc is potentially limitless. Whilst all of it may be deterministic in nature, our limited capacity to analyse the evolution of such non-linear systems essentially reduces much of what we witness to luck, even if, in a theoretical sense, Poincaré was right to insist that every phenomenon has a cause. Or was he?

The Uncertainty Principle

The Age of Reason and the scientific determinism that accompanied it were

snuffed out on the battlefields of the First World War. Until then, probability theory represented little more than an epistemological paradigm, illustrating the practical limits to analysing causality and predicting the future within a universe that nevertheless was fundamentally deterministic. All that changed in the early years of the 20th century. Already, Albert Einstein had revealed that Newton's laws were but mere approximations of a more general 'truth' about space, time and gravity, whilst quantum mechanics (the science of the very small) began to reveal that the very universe itself might behave probabilistically. The 'items of nature' that Laplace's demon was charged with studying started to behave like waves, with no fixed position. How can you predict where something is going to be in the future when you don't even know where it is right now?

It wasn't until 1926 that Werner Heisenberg, a German physicist, began to realise the implications that wave-particle duality would have for determinism. Heisenberg pointed out that you couldn't measure both the position, and the speed, of a subatomic particle exactly. The following February he published his now famous Uncertainty Principle, which stated that the more precisely the position of some particle is determined, the less precisely its momentum can be known, and *vice versa*. This was not a constraint imposed by the physical limitations of practical observation. On the contrary, it was an impossibility imposed by the very nature of matter itself.

Even Einstein himself was unhappy at such a probabilistic interpretation of the universe. In a letter to his friend and colleague Max Born, another German physicist, just before Heisenberg published his Uncertainty Principle, he expressed his dissatisfaction clearly.

"Quantum mechanics is certainly imposing. But an inner voice tells me that it is not yet the real thing. The theory says a lot, but does not really bring us any closer to the secret of the 'old one.' I, at any rate, am convinced that He does not throw dice."

Of course by 'He', he meant God. Einstein believed that the uncertainty was only provisional, and that there was an underlying reality, in which particles would have well defined positions and speeds, and would evolve according to deterministic laws in the spirit of Laplace. He was wrong. Even God is bound by the Uncertainty Principle, and cannot know both the position, and the speed, of a particle. As Stephen Hawking says, "*all the*

evidence points to Him being an inveterate gambler, who throws the dice on every possible occasion.” Moreover, He doesn’t even know what the outcomes will be.

Other scientists, however, were ready to take up the challenge. Wave functions came to represent particles which have ill-defined positions and speeds. The size of the wave function gives the probability that the particle will be found in that position, whilst the rate at which the wave function varies from point to point provides a measure of the momentum of the particle. If you know the wave function at one time, then its values at future times can be determined by what is called the Schrödinger equation, named after its inventor, the Austrian physicist, Erwin Schrödinger. This is not the sort of determinism that Laplace envisaged. Instead of being able to predict the exact positions and speeds of particles, all we can predict is the wave function, which provides only a probabilistic measure of position and momentum. According to the Schrödinger equation, the best we can do is predict only half what Laplace envisaged his demon was capable of. Perfect predictions about the future are impossible since the stuff out of which the universe is made behaves randomly.

According to Nate Silver, in his book *The Signal and the Noise*, we needn’t worry about the implications of Heisenberg’s Uncertainty Principle. Most things we are interested in predicting, like sports, the markets and the weather operate at the macroscopic level, many orders of magnitude bigger than the size of atoms. The physical stuff of reality is much too large to be discernibly influenced by quantum mechanics. While Heisenberg’s Uncertainty Principle disrupts causality at the atomic and subatomic level, it typically does not rear its head in the macroscopic world. Avogadro’s number⁷ is so large that the probabilities that influence a small number of atoms essentially collapse into virtual certainties via the law of large numbers. Not so according to Andreas Albrecht. In a paper⁸ published at the end of 2014 with co-author Daniel Phillips, both at the University of California, the quantum mechanical behaviour of atoms may very well be responsible for the probability of all actions, with far-reaching implications for theories of the universe (as well as gambling).

The connection between the subatomic quantum world and the macroscopic classical world can be seen in Brownian motion, named after

the 19th century botanist Robert Brown who first observed the random haphazard movements of small pollen grains suspended in water. Most high school students will have seen it at one time looking through a microscope during a science class. Even though they can't be seen, the water molecules are in a constant state of thermal motion, repeatedly colliding with the much larger pollen grains (up to 250,000 times in diameter) in all directions. Despite the number of collisions taking place, the randomness with which they occur ensures that there are always tiny imbalances at any given time with slightly more molecules pushing a grain on one side than the other.

The Heisenberg Uncertainty Principle dictates that the trajectory of a water molecule will have an inherent uncertainty, resulting from the uncertainties in its position and momentum. Albrecht and Phillips calculated how this uncertainty grows with each collision between molecules. In fact, the uncertainty becomes so large in the space of one collision that every single fluctuation in the property of water has a fully quantum-mechanical origin. Furthermore, the researchers claim, the quantum fluctuations manifest in the water could wholly determine the outcome of a toss of a coin. Quantum uncertainty in the position of neurotransmitter polypeptide molecules (amino acid chains) in the nervous system of a coin tosser arises because of the Brownian motion of these molecules in a fluid that is largely water. This quantum uncertainty will subsequently translate into an uncertainty in the number of times a coin turns in the air before being caught, via the molecular interactions that (non-linearly) amplify the tiny quantum fluctuations, to ultimately determine whether it lands heads or tails. Hence, Albrecht and Phillips maintain, the classical probability and randomisation associated with the tossing of a coin **emerges** from underlying quantum probabilities. As such, because the uncertainty of such a system increases non-linearly with every subsequent Brownian collision, once that uncertainty becomes large enough, its quantum effects become the dominant factor in the outcome, not classical mechanics. For a game of snooker, for example, Albrecht and Phillips calculated that it could take just 8 collisions between balls for quantum uncertainty to dominate.

In an attempt to make sense of the strange world of quantum uncertainty, Schrödinger hypothesised a cat, Schrödinger's cat, whose continued existence or swift demise would be 'determined' by the radioactive decay

of a single atom. If, during an hour, the atom decays, this triggers the release of some poison which kills the cat. On the other hand, if no decay takes place, the cat is spared. Since the decay or otherwise of the atom is governed by Heisenberg's quantum uncertainty, it is impossible to know whether the cat is alive or dead, until at some later time the cat's health is observed. Essentially, the cat's wave function describes a superposition of states during which it is paradoxically both alive and dead at the same time. Only once we observe the cat does its wave function superposition collapse to reveal whether it is dead or alive. Heisenberg's Uncertainty Principle forbids the possibility of predicting *a priori* whether the cat will live or die.

The implications of Schrödinger's thought experiment and of the emergence of classical randomness from quantum uncertainty for gamblers hoping to predict the future are considerable. Albrecht and Phillips' theory describes a kind of chaotic (but non-deterministic) system, in which the tiny fluctuations at the quantum scale – the equivalent of a butterfly flapping its wings – become amplified via countless molecular interactions until they collectively manage to have an impact at the macroscopic scale. The tossing of a coin, indeed the scoring of a goal, the serving of an ace, the electoral choice of an undecided voter, the size of dividend of a company and any number of other real world phenomena where decision making driven by neural processing plays a role are the probabilistic equivalents to Schrödinger's cat. The final state cannot be predicted until it has actually happened. In fact, there may be no physically verifiable fully classical theory of probability at all, just a quantum one, where the multitude of possible shapes and sizes of Cleopatra's nose and the futures that evolve from it may all be happening at the same time. The history that we experience may all just be an illusion. This is a very disconcerting picture for human beings designed to handle only either/or. Visualising quantum-like superposition of probabilistic states is not something that comes naturally.

Playing Games

If the horrors of the First World War and new science of quantum mechanics didn't present enough of a challenge to those still insisting that

all uncertainty could be managed, risks reduced to zero and economic stability guaranteed, along came the Great Crash of 1929 and the Great Depression that followed it. In 2008, with the financial crash, the world was again reminded that the science of prediction will never be perfect, and black swans⁹ are inevitably waiting around every corner of the future. The turmoil unleashed at the end of the Roaring Twenties, characterised by the decade's social, artistic, cultural and economic dynamism, particularly in the United States, would ensure that never again would economists insist that fluctuations in the economy were a theoretical impossibility.

The American author F. Scott Fitzgerald seems to have been intuitively aware of the events that were about to unfold (or was he just blind lucky) with his 1925 novel *The Great Gatsby*. Symbolising the reckless greed of the period, the central character, Jay Gatsby, a mysterious millionaire with shady business connections, had returned to New York to pursue his dream of winning back his former love Daisy Fay Buchanan, an attractive, if shallow and self-absorbed, socialite married to Tom Buchanan, an 'old money' millionaire. Gatsby was possessed with the idea that the past could be repeated. Indeed, for Fitzgerald, it seems he was the personification of the overconfidence, exuberance and hope that embodied the Roaring Twenties, the idea that the past was the key to the future, self-interested rationality was still king and all was right with the world. Such rationality would, in due course, be shown to be largely an illusion.

It was about this time that a new branch of mathematics began to grasp people's attention. Thus far, economists and mathematicians had really only considered the significance of isolated individual choices with regards to behavioural decision making. The rational choice (or utility) theory of Daniel Bernoulli, another nephew of Jacob, had formed the basis of how people make decisions when faced with uncertainty since the mid 18th century. How much, for example, should someone consider wagering on red-black roulette in comparison to a bet on a single number of the wheel, or horses in a race priced at odds of 2/1 versus 10/1? I'll be taking a closer look at Daniel's theory later in the book. Most gambling wagers, of course, with the exception of pure games of chance (casino and lottery), involve the interaction of at least two players, and in financial and sports betting markets, very many more. This new mathematics, by contrast, started to

investigate how these interactions influenced the behaviour of decision makers. The new mathematics was called game theory. Its premise was that the true source of uncertainty lies not in the probability of outcomes but in the intentions of others. In game theory, you make choices by anticipating the payoffs for your opponents. The correct thing to do depends upon what other people do.

Despite it having historical and philosophical motivation stretching back to the Ancient Greeks, game theory did not secure a proper mathematical grounding until 1928, with the publication by John von Neumann, a Hungarian turned American polymath, of his minimax theorem, a decision rule used for minimising the possible loss for a worst case scenario. Originally formulated for two-player zero-sum¹⁰ games covering both the cases where players take alternate moves and those where they make simultaneous moves, it has also been extended to more complex games and to general decision making in the presence of uncertainty. In zero-sum games, the minimax solution is equivalent to the Nash equilibrium, named after its formulation in 1951 by John Nash, an American mathematician and 1994 Nobel Prize winner whose moving life story was the subject of the 2001 film *A Beautiful Mind*. The Nash equilibrium is a solution concept of a non-cooperative game involving two or more players, in which each player is assumed to know the rational motivations of the other players, and no player has anything to gain by changing only their own strategy. If no player can benefit by changing strategies while the other players leave their strategies unchanged, then the mix of strategy choices and the corresponding payoffs for all players constitutes a Nash equilibrium. Real life examples of Nash equilibria include the game rock-paper-scissors, penalty shoot outs in football, the growth of forests in the Amazon and even the Israel-Palestine crisis. The classic example of a Nash Equilibrium where players refuse to cooperate and gain less than they otherwise would through mutual cooperation is also known more popularly as The Prisoner's Dilemma.

One of von Neumann's case studies involved a trivial game of penny match, conceptually a two-strategy equivalent of rock-paper-scissors. Two players simultaneously turn over a coin. If they match, player 1 wins, if they are different player 2 wins. We don't need to call on Pascal, the

Bernoullis and de Moivre to know that each has a 50-50 chance of winning each round. The crucial point of the game, however, is that if one player attempts to adopt a systematic strategy, its predictability will also be visible to his opponent. In terms of game theory, there is no pure strategy that offers the best response to a best response, but instead a Nash equilibrium involving a mixed strategy, which in this case is for both players to show heads or tails randomly. According to von Neumann, the trick to playing this game, and any zero-sum strategy game for that matter where most players are trying to act rationally, lies not in attempting to guess the intentions of your opponents but in not revealing your own intentions. Again, think of penalty shoot outs.

Game theory paints gambling games in a completely different light. Poker players, of course, will be entirely familiar with the premise that positive expectancy is only to be found through a better interpretation of what your opponents are up to and a superior means of concealing your own strategy. Often, the most profitable players are not those with the best hands, since what cards you are dealt is purely a matter of chance. On the contrary, it is those players most adept at bluffing, most skilled in knowing when to fold a poor hand (and even some good ones for that matter) and what stakes to call and raise who are most likely to do well at poker. Paradoxically, the same will be true in sports betting and investing in the stock market too. Playing in such markets offers reward expectancy, not so much through a thorough understanding of ‘true’ chances or ‘true’ value, but through the awareness of how other players are playing the game. Since the odds or price in a market are defined less by the theoretical probabilities of outcomes and more by the flow of money into that market, it is only by appreciating the opinions, motivations and intentions of others that we can truly hope to secure a long term positive expectancy. We are not playing the odds, we are playing real people.

Irrationality

16 years after von Neumann’s minimax theory, he published together with his colleague Oskar Morgenstern, a German-born economist, the *Theory of Games and Economic Behaviour* that explored more deeply into the nature

of behavioural decision making. Morgenstern, clearly, was of a similar mind to von Neumann. No one, he insisted, can know what everybody else is going to do at any given moment. Crucially however, their theory of games was premised upon one core assumption: that players behave rationally. Increasingly, students of behavioural psychology were becoming aware that this assumption might have major flaws. The work that evolved from this criticism was to culminate in a Nobel Prize in 2002 for one of its leading lights, an Israeli psychologist named Daniel Kahneman, for his research with colleague Amos Tversky into the cognitive basis for systematic human errors that arise from heuristics¹¹ and biases and the development of ‘Prospect Theory’ that describes the way people choose between probabilistic alternatives that involve uncertainty. The magnificence of prospect theory was that it finally codified what most of us probably know already: that we are all less than fully rational creatures. Prospect theory is to rational choice theory (the standard economic model since the Enlightenment) as Einstein’s Theory of Relativity is to Newton’s Laws of Motion.

Kahneman’s initial sensitisation to the systematic mistakes that human beings make first arose with the realisation that his students were completely oblivious to regression to the mean. During the 1950s, Kahneman worked for the Israeli air force, teaching flight instructors about the psychology of effective training, and in particular that rewards for improved performance work better than punishment for mistakes. One of the instructors observed that, more often than not, following praise for a successful manoeuvre by one of his trainees, the trainee would subsequently perform less well. On the other hand, criticism of bad flight execution was more typically followed by an improved performance. For Kahneman this was a Eureka moment when he saw in a new light the principle of statistics that he had been teaching for years. To quote Kahneman, “*the instructor was right – but he was also completely wrong!*” The observation was correct, but the inference he had drawn about the efficacy of reward and punishment was completely invalid. Essentially, the instructor had attached a causal interpretation to a completely random process, regression to the mean, in much the same way as followers of football believe in the new manager effect. On average, superior flying performance, being less

common than is typical, can be expected (although of course not guaranteed) to be followed by less skilled flying, and *vice versa*.

Following his discussion with Tversky about the episode, the pair quickly realised that ignorance about regression to the mean was probably not the only way people make systematic errors when trying to find causal explanations for events both now and in the future. The rest, as they say, is history. Over the next three decades until Tversky's untimely death in 1996, the two of them, along with several other economists and behavioural psychologists, set about uncovering many of the systematic ways human beings commit errors, why they commit them, and how they deviate from fully rational choices. The crowning achievement was in proving, through prospect theory, that people make decisions based on the relative value of losses and gains rather than the final outcomes, and that losses are more important than gains. Whilst game theory showed that the outcome of zero-sum competitions was dependent on the interaction of the intentions of players, prospect theory revealed that very often the intentions of players are not even rational.

The implications for gamblers of this theory are immense, and for this reason much of a later chapter is devoted to it. For now it is sufficient to say that regression to the mean is not the only systematic error, or fallacy, that gamblers exhibit. Of course, we already know about the Monte Carlo fallacy or the fallacy of the maturity of chances, the mistaken belief that, if something happens more frequently than normal during some period, it will happen less frequently in the future, and *vice versa*. Here are some more that are most relevant to gambling: the tendency to over-bet low probability outcomes and underbet high probability outcomes (the so-called favourite-longshot bias); the tendency to chase losses (paradoxically a form of loss aversion); the tendency to overestimate causality and the significance of past events despite mathematical independence; the tendency to remember long streaks more than short ones (pattern recognition); the tendency to overestimate the degree of skill involved in games involving both skill and chance (for example sports betting, poker and financial trading); the tendency to overweight and generalise the significance of small samples; and perhaps most important of all the tendency to assume that winning has causal explanations but losing arises because of bad luck.

If a lot of gambling (excluding pure games of chance) represents a game

of psychology – let's call this speculative gambling – where many of the players are behaving irrationally, does this not imply that opportunities for long term profitable expectancy exist for those ready and armed to exploit this? Theoretically, yes. In practice, however, it's not so simple. Market makers who facilitate the playing of these games demand a commission for their service; they're not charities after all. Bookmakers impose an overround on their betting odds; financial trading platforms take a cut of the action through the buy-sell spread and other transaction costs; poker rooms charge a rake. Perhaps more importantly, however, understanding that irrationality in gambling markets exists is one thing, but consistently being able to predict it for profitable ends is quite another. Unfortunately the evidence on that is pretty unambiguous: whilst gambling markets can and do behave irrationally from time to time, few players are actually capable of taking advantage of it. Most players are really engaged in little or nothing more than a random game of coin tossing. Of course, it's easy to believe, through a bias in our confidence, that this is not so, but the data just speaks for itself. We'll be looking at some of that data later in this book.

Many of these psychological errors, biases and fallacies stem from an illusion of control. We are programmed to find explanations for things that happen, whether that is the winning of a poker game, the successful prediction of a football score or the rise and fall of a stock price. Whilst in their turn probability theory, chaos theory, quantum theory, game theory and prospect theory have all revealed how uncertain, unpredictable and irrational the world and our interaction with it often is, we are biased to ignore and be fooled by this randomness in favour of explanations that provide us with meaning. If we can understand why things happen, we can begin to predict them too. And if we can do that we can achieve control. To a living creature trying to compete and survive in a world of limited resources, where every other is trying to do the same, control is everything, even if it's illusory.

In this sense gambling, and in particular the speculative gambling of competitive markets (betting, trading, poker), represent an attempt to take control of uncertainty. The paradox here is that most of us **know** that chance and luck can't be controlled. The trouble is that often we **feel** that they can be. Daniel Kahneman showed us the two sides to our brains, one slow, rational, methodical but unfortunately energy intensive and

consequently lazy, the other fast, intuitive and often emotional, designed over the course of evolution to take short cuts to achieve desired outcomes. Where the world is random, or mostly random, those short cuts can lead us down blind alleys, but it's hard to avoid them. Hence, whilst most of gambling might be considered to be irrational, the paradox is that many of us continue to do it, whether on games of pure chance or others that offer a theoretical advantage to players astute enough to find it. It is to why some of us do choose to seek control through gambling and why others choose to condemn it that we will now turn.

² Including the influence of the single zero on a European roulette wheel reduces the odds of black (or red) to 18/37 or 0.486. Consequently, whilst Wikipedia report the odds of this event as 1 in 67,108,864, the actual probability was less than half as likely.

³ The binomial theorem describes the algebraic expansion of powers of a binomial, that is to say two algebraic terms, for example x and y, or heads and tails. Expanding these powers reveal coefficients that appear as entries of Pascal's triangle, where each entry is the sum of the two above it (for example 1; 1,2,1; 1,3,3,1; 1,4,6,4,1; 1,5,10,10,5,1 and so on). The binomial theorem, for example, can be used to determine the number of possible outcomes of successive tosses of a coin.

⁴ For example, see <https://www.mathsisfun.com/data/quincunx.xhtml>

⁵ Ter Weel, B., 2011. Does Manager Turnover Improve Firm Performance? Evidence from Dutch Soccer, 1986–2004. *De Economist*, **159**(3), pp.279-303.

⁶ Anderson, C. & Sally, D., 2013. *The Numbers Game: Why Everything You Know About Football is Wrong*. New York: Viking.

⁷ Avogadro's constant defines the number of atoms or molecules contained in the amount of substance given by one mole, where a mole is the amount of pure substance containing the same number of chemical units as there are atoms in exactly 12 grams of carbon. The number is 6.023 x 10²³, or about 600 billion, trillion.

⁸ Albrecht, A. & Phillips, D., 2014. Origin of probabilities and their application to the multiverse. *Physical Review Letters*, **D90**(12), 123514.

⁹ A black swan is a metaphor, originally defined by the author Nassim Nicholas Taleb of the book with the same name, to describe an event that comes as a surprise, has a major, often ruinous effect, and is often inappropriately rationalised after the event with the benefit of hindsight.

¹⁰ In game theory and economic theory, a zero-sum game is a mathematical representation of a situation in which each player's gain (or loss) is exactly balanced by the losses (or gains) of your opponents. If the total gains of the participants are added up and the total losses are subtracted, they

will sum to zero. In purely financial terms, all of what we traditionally understand as gambling represents zero-sum games. Financial investment is considered to represent a non-zero-sum game, a feature that moral critiques of gambling use to emphasise their opposition to zero-sum games and their separation of investment as something distinct from and different to gambling. These ideas will be explored in the next chapter.

[11](#) A heuristic is an approach to learning or decision making that employs a practical methodology not guaranteed to be optimal or perfect, but sufficient for the immediate goals. It is a kind of cognitive (mental) short cut.

To GAMBLE OR NOT TO GAMBLE: IS THERE A QUESTION?

Here's another paradox: if gambling, ingrained as it is in the human psyche, manifests itself so ubiquitously through time and across cultures, why have so many so often considered it to be deviant? In addition to the secular arguments that gambling is a tax on the poor and the stupid, or more historical ones presenting it as a distraction from more noble and socially beneficial pursuits, much of this objection has been moral, and more specifically religious, in origin. Whilst God made man in his own image, he has become inherently corruptible through sin: wrath, greed, sloth, pride, lust, envy, and gluttony. For its detractors, gambling encapsulates most, if not all, of those. To make right the wrongs, condemnation and prohibition from a higher authority, whether religious or political, is usually the preferred medicine. And throughout, the facts about how humans actually behave are retrofitted to match prescribed theories of how it is believed human beings ought to behave. I want to abandon this top-down philosophical jamming of a square peg into a round hole. I will endeavour to investigate the social and evolutionary explanations for why so many people, particularly men, like to gamble in one form or another, and why they are not necessarily bad for doing so. Evidently, a useful starting point is to consider the historical context of gambling, its origins and cultural variances. However, before embarking on this undertaking, it is probably necessary to begin by defining exactly what we mean by gambling, and other activities closely related to it. You may have noticed I've already made a first attempt at the end of the previous chapter.

Gambling, Speculation & Investing

In his *Complete Guide to Gambling*, John Scarne, the American magician and foremost gambling expert, defined gambling as “*risking something one*

possesses in the hope of obtaining something better.” We would all probably agree. Yet such a definition, on the face of it, appears to differ little from the business of investing and speculating. Instinctively, we may feel there is a difference; gambling means things like roulette, craps, blackjack, bingo and lotteries; investing means things like the stocks, bonds, property and pensions. The former concerns games of pure chance, the latter arguably educated attempts to increase one’s wealth. Speculation probably lies somewhere between the two, not just subject to luck but ostensibly involving greater risk than investing. However, on closer inspection things are not quite as simple as they first appear.

Let’s start with some dictionary definitions. thefreedictionary.com lists the following possibilities to describe the act of gambling: a bet on an uncertain outcome; to play a game of chance for stakes; to take a risk in the hope of gaining an advantage or benefit. Investing, meanwhile, is described as an act of committing money in order to gain a financial return, whilst speculation is engaging in buying or selling of an asset with an element of risk on the chance of profit. Spot the difference. All three appear to involve the same thing: risking money on the chance of making more of it. Clearly, the distinction must be more nuanced than this.

As we’ve already observed, gambling and investing are often distinguished by identifying the possibility of skill and a profitable expectation, that is to say, tilting the odds in your favour. Pure games of chance operate simply according to the laws of probability. There is no element of skill involved and no chance of a positive expectation. The only way you can win is through luck. We are on safe ground it would seem in describing such games as gambling. Or are we? Whilst this might be true for online casinos that can manage their games through mathematical algorithms, bricks and mortar casinos use roulette wheels that have physical imperfections which can, in theory, be exploited through a technique known as clocking, recording thousands of observations of wheel outcomes to detect any bias and the possibility of a positive expectation. The most celebrated example of roulette clocking occurred at the Monte Carlo Casino in 1873, when Joseph Jagers, accompanied by six clerks, clocked one wheel with such significant bias that they managed to walk away with two million francs, then about £65,000 and in today’s money equivalent to over £3,000,000. Others have attempted to meticulously predict the compartment

where the roulette ball will come to rest. Famously, the Eudaemons were a small group headed by graduate physics students J. Doyne Farmer and Norman Packard at the University of California Santa Cruz, who in 1978 managed to make about \$10,000 (averaging a 44% profit for every dollar wagered) by using a video camera and a computer concealed in a shoe that interpreted the visual data by means of some fairly sophisticated mathematics. In an attempt to combat such enterprises, casinos have increased their level of maintenance and rotation of roulette wheels, reducing the window of opportunity to exploit any available bias.

It is also possible to tilt the odds of blackjack in your favour by means of card counting, which whilst technically legal, is frowned upon by most casinos who will probably ask you to leave if they catch you indulging in it. Card counting involves the tracking of cards played in previous rounds of blackjack, thereby allowing the counter to predict with greater probability what cards the dealer will hold in subsequent rounds, enabling him to bet more with less risk when the count gives an advantage as well as minimise losses during an unfavourable count. Card counting systems that track fluctuations in deck composition can yield player expectations in excess of 2%. To combat card counting, casinos make use of automatic detection systems as well as automatic deck shuffling machines and the use of a greater number of decks of cards. Online blackjack, of course, precludes the possibility of card counting, since every card is drawn randomly.

At physical casinos at least then, games of pure chance appear to offer the possibility of a positive expectation and tilting the odds in one's favour, even if the lengths one has to go to are considerable. Would we really choose to describe such activity as investing? What about activities like poker and betting? Both games undeniably offer the theoretical possibility of a profitable expectation because they involve the speculation, by competing players, about things of unknown probability. For that reason, I've already labelled this as speculative gambling. Furthermore, some players, conceivably, might be better at it than others. Having observed the various practices of people who tip on sports over the past 14 years, it is common to see them describe what they are offering as investment. In recent years the United States court system has tied itself in knots debating whether poker is a game of skill. Much of it hinged around the following question: can you deliberately lose a game of poker? Clearly the answer is

yes, although of course trying to do so won't guarantee you a loss, given the substantial amount of luck involved in the game. The same is probably true of betting, whether on horses or sports. Intentionally betting randomly on longshots will increase a negative expectation relative to a similar strategy on favourites.

Most sports bettors, however, are not the slightest bit skilled, despite beliefs to the contrary. The data on this, which I will review later in the book, is pretty unequivocal. Their pattern of profits and losses matches almost perfectly the pattern that we would predict to occur simply by chance alone. It's one thing to say that you are theoretically engaged in investing with the odds on your side; it's quite another to prove that you actually are. If your outcomes match those which are predicted by luck, it's probably safe to say that what you are really doing is gambling.

The same is true of the stock market. Traditionally, this has always been regarded as an investment arena, partly because of the function it serves as an engine of capitalism, and partly because of long-standing social and cultural differences that have tended to regard zero-sum gambling as 'bad' (profits balanced by losses) in contrast to positive-sum investing as something mutually beneficial for all of society. But again, on closer inspection, things are not so clear cut. High frequency trading, for example, which seeks to exploit tiny market inefficiencies over time scales as short as a nanosecond, would appear to represent a zero-sum game with little benefit to the wider economy other than the taxes that companies engaged in such practices will be contributing. Similarly, individual investors who trade over periods of hours to days – so-called day traders – will be superficially engaged in a zero-sum game. More generally, however, we might also question whether the longer term investment mechanisms of the financial markets really represent a positive-sum economy at all. It is true to say that long term economic growth is positive, but many have begun to question whether the social and environmental consequences of this growth, including pollution and differential poverty, have been properly costed. Indeed gross domestic product (GDP), the monetary value of all goods and services produced by a nation, makes no distinction between ones which are advantageous to individual, environmental or societal well-being versus those that detract from well-being. Trade, undeniably, is a good thing, but not at any price.

Perhaps more significantly, it is doubtful whether many ‘investors’ in the stock market really have the odds in their favour. In his eye-opening ebook *Monkey with a Pin*, Pete Comley puts forward the very convincing argument that the average investor is losing 1% per year once the charges of playing in the stock market and the effects of inflation are properly taken into account. Furthermore, there is now substantial research that reveals the majority of professional fund managers are failing to consistently beat the market as well. Since that market, most of the time, represents a random walk, this must surely bring into question whether most of us playing the financial markets game are doing anything other than throwing dice. If that is the case, can we really call this investing as we’ve defined it above?

Perhaps a better way to distinguish between gambling and investing is to consider their underlying motivations. People gamble for entertainment, for the thrill of playing and the anticipation of winning, and usually over very short time horizons with swift closure to the games. Indeed the word is believed to be derived from the old English *gamenian*, meaning to joke or play, and *gamen*, meaning to sport, joke or jest. Coincidentally, ‘happy’ is derived from the Middle English *hap* meaning luck, chance or fortune. By contrast people invest for business reasons, for their futures, for their security, usually over much longer time frames, often with no closure in mind at all. Similarly, whilst compulsive gambling is a well recognised problem, no such addiction is believed to exist for investing. In truth, this is probably more to do with a lack of research into the social impacts of recent phenomena like day trading, now accessible to anybody with a PC, some inexpensive trading software and a trading account. Having watched a member of my own family squander his entire inheritance on alternative investments like carbon credits, ‘development’ land and rare earth metals through unregulated (and sometime fraudulent) companies, it’s hard to accept that so-called stockaholics aren’t just as prone to addiction. More importantly, a new field of research called neuroeconomics is starting to reveal that the brain circuits which light up during casino gambling are the same ones that get excited when people trade on the financial markets and for that matter when people get high on cocaine, alcohol, chocolate and sex. At the heart of it is the hormone dopamine, responsible for the pleasurable feelings associated with reward anticipation; more about that later.

Clearly, then, whilst we might believe that a social and cultural

distinction between gambling and investing exists, even if simply for moral reasons, behavioural motivations and their outcomes seem to differ little across the numerous opportunities to play these games. Whether investing or gambling, for most of us the thrill lies with the anticipation of reward, with that reward mostly subject to chance alone. Whatever we choose to call it, it largely amounts to the same thing. Arguably a better distinction is between professionalism and recreation. The few (and they really are a few) who do manage to tilt the odds in their favour could be regarded as professionals, whether roulette clockers, card counters and consistent winners of betting on sports, poker and the financial markets. The rest of us, including many harbouring false confidence, are really just in it for the fun, whether we like it or not. It might not be such a stretch to call professional gamblers ‘investors’ and recreational investors ‘gamblers’. When all is said and done, perhaps the surest way to tell if you’re a gambler or investor is to ask yourself the following question: does it consistently and reliably provide my main source of income with which to pay the bills? For almost everyone who plays these games, whether poker, sports or stocks, the answer, as will become clear, surely has to be no.

So finally let’s define gambling as a speculation involving money on the future which is unable to show a consistent (risk-adjusted) return on investment superior to the market benchmark. What that market benchmark is will depend on the game. At the casino it will be defined as their house margin; for poker, it’s the rake; for sports betting it’s the bookmaker’s overround, conceivably as low as 0% if one is diligently comparing prices to find best market value; and for investments in the financial markets it will be whatever appropriate index your stock, bond or mutual fund should be measured against. Given that very few players manage to beat their benchmarks, whatever game they are playing, we should perhaps conclude that most of what we’ll be talking about in the remainder of this book is all just gambling. Perhaps what are more important in this context of definition are outcomes, not expectations.

A Brief History of Gambling

It seems *Homo sapiens* may have been infatuated with gambling for a very

long time. Archaeological evidence from prehistoric sites across Europe, Asia and into North America has uncovered cube-shaped ankle-bones called astragalia, some of them dating back as much as 40,000 years. Their purpose is a matter of speculation, but accompanying cave drawings hint at the possibility they were used as some form of entertainment and a means of prophecy. By casting them and interpreting the outcome, Stone Age man may have sought knowledge of the future and the intentions of his gods. Furthermore, it is conceivable that the playing of such games formed an integral part of a hunter-gatherer psychology that was well versed in the art of risk taking as a means of survival. When faced with uncertainty, particularly concerning matters of food availability and safety from predators, it pays to have a means of divining the future.

More recent civilisations have continued to indulge in gambling related play. In Ancient Greek mythology, the universe was even created by a game of chance. Zeus, Hades and Poseidon are said to have divided up the spoils (heaven, hell and sea) with the throw of some dice, a popular game in Ancient Greece where they used three cubes made of clay. The Minoan civilisation on the island of Crete is thought to be responsible for the origin of poker more than 3,500 years ago. The Romans, too, seemed to like playing with dice as well, but reduced the number to two, as is now common in games of craps. Pairs of dice have even turned up in the ruins of Pompeii, some of them 'loaded'. So passionate was the emperor Claudius about the game of dice that he published a book on the subject and had his carriage redesigned with a special board to keep his dice from rolling off.

The ancient Chinese were prolific inventors of gambling games. Around 4,300 years ago they created a game of chance using tiles. The game of keno, which is played with cards or tickets numbered 1 to 80 in squares, has its origins dating back at least 2,000 years. The original game was called *baige piao* meaning 'white pigeon ticket,' referring to the tickets used in a betting game involving homing pigeons. The Chinese were also the first to start using playing cards as far back as the Tang Dynasty in the 9th century. Elsewhere, ancient gambling artefacts have been uncovered as far afield as Egypt, India and Japan. Native Americans also gambled, believing both that their gods invented games of chance using coloured stones called plum stones and divined their outcome as well.

Whilst often scorned by the great monotheistic religions, one form of

gambling makes a regular appearance in their texts: the casting of lots. Given its purpose as a means of divination that is perhaps not surprising. No one was appealing to chance when lots were drawn, but to the will of God. The origin of the word lot can be found in the old English word *hlot* and its Germanic precursor *hleut*, meaning pebble, although other objects such as dice, straw and wood chips would have been used. The practice is mentioned in the Old Testament as many as 70 times, and a further 7 times in the New Testament. References also appear in the Talmud and the Qur'an. The drawing of lots during religious rituals was used to discover God's will in decisions concerning a number of issues, including the election of kings, the identification of sacrilegious offenders and the settlement of disputes. Typically, however, the practice was used in the division of land and property, most notably the tribal allotments of Israel under Joshua. The Gospel of John even describes the casting of lots by Roman soldiers for the Seamless Robe of Jesus after his Crucifixion.

Today's lottery, as a descendant of the practice of casting lots, still uses the drawing of numbered balls to award prizes, although of course Divine Providence is no longer considered to play a leading role. The lottery as a game of chance rather than a system of godly decision making appears to have been prevalent during the reign of Augustus, Rome's first emperor. Lottery tickets were sold to fund repairs in the City of Rome, and the winners were given prizes in the form of articles of unequal value. The first recorded examples of lotteries in Renaissance Europe date from the mid 15th century in Holland and Belgium. By the 16th century, the Italians, French and English held them too. In 1569, Queen Elizabeth I established the first English lottery, when she offered 400,000 tickets for sale. Prizes included china, tapestries and cash. This and subsequent lotteries were designed to raise money to help fund England's colonial endeavours and finance the nation's growing debt. The first London lottery of 1612 during the reign of King James I, for example, funded the building of the Jamestown colony in Virginia, the first English colony in America. Lotteries in colonial America later played a significant part in the financing of both private and public ventures, including the French and Indian Wars and the War of Independence, before legislation outlawing them took effect at the end of the 19th century.

The origin of 'casino' is rooted in the Italian word 'casa' meaning small

house or recreational place. The oldest casinos date from the early 17th century, the most famous of which was the Casino di Venezia, still operating today. The function of the casino was to act as a focal point for social gathering, bringing together people of similar interests and skills. Through the 18th century, their popularity spread across Europe, and in particular to Monte Carlo, which positioned itself as Europe's capital for legalised casino gambling. The game of blackjack, or twenty-one as it was known, probably evolved in the French casinos around 1700. Roulette, literally meaning 'little wheel', probably also evolved in France's casinos, a century or so after Pascal built a primitive wheel during his quest to discover a perpetual motion machine, although its roots as a game may date back much further to the ancient Egyptians. Just as the lotteries beforehand, the concept of the casino was exported to the New World. Saloons, as they were initially known, quickly appeared in the major cities of America, including New Orleans, St. Louis, Chicago and San Francisco. With them came the traditional games, and some new ones offering simpler versions of their European progenitors, for example the dice game craps, which developed from the early English game of hazard.

Apart from forerunners in ancient Rome and Greece, organised and sanctioned betting dates back to the 18th century. Bookmaking and betting on sports, including racing, became a very English pastime. The Jockey Club, believed to have been founded around 1750, allowed gentlemen gathering to socialise and watch horse races to place stakes on the winner, and the noble art of bookmaking was born. The word 'bookmaker' arises from the literal meaning 'a compiler of books,' with people accepting bets recording their details in a book. At about the same time, pedestrianism, essentially a form of race walking, allowed people to wager on how much time a competitor would take to complete a predefined distance. As the bets became larger, so did the distances. In 1789, an Irishman won himself £20,000 by walking to Constantinople (today Istanbul in Turkey) and back in less than a year. Jules Verne's 1872 novel *Around the World in 80 Days* was inspired by this craze for distance bets. Towards the end of the 19th century, sports betting took on a more organised approach. To ensure bookmakers could take a profit without the necessity to cheat, overround¹² betting was introduced in the early 19th century. This standardisation of the

bookmaker's profit margin according to mathematical principles effectively professionalised the betting industry. Towards the end of the 19th century team sports in England were also being professionalised. The Football Association was formed in 1863, followed by the Rugby Football Union in 1871. Test cricket was introduced in 1877, followed 18 years later by Rugby League. During the 1880s, newspapers started offering fixed prizes for correctly predicting the outcome of football games. These prizes became known as 'fixed odds'.

In his book *Gambling: A Story of Triumph and Disaster*, Michael Atherton (the former England cricket captain) provides a fascinating account of the synergy that existed between gambling and financial speculation during the 17th and 18th century Enlightenment, at a time when gentlemen were beginning to understand the new mathematics of chance and put the theories to good use. At the forefront of this explosion in financial risk taking was a Scottish economist named John Law. Law believed that money was only a means of exchange that did not constitute wealth in itself and that national wealth depended on trade and the law of supply and demand, a foresight that pre-dated Adam Smith and his *Wealth of Nations*, and the replacement of mercantilism as the dominant economic theory, by almost a century. It is not a stretch of the imagination to suppose that he garnered such ideas in the early gambling rooms of Europe where he amassed a fortune, in large part due to his understanding of the new theories on probability. In 1716 Law's newfound wealth and fame helped him to create his own private bank – Banque Générale Privée – which introduced the use of paper money. Most of the capital, however, consisted of French government bills and government-accepted notes, effectively making it the first central bank of France. A year later he bought the Mississippi Company, initially to control trading rights in Louisiana but later also the Indies, China and Africa and finally France's entire national debt as he sought to exaggerate the wealth of the company. Sadly for its creditors and the whole of France, Law funded the conveyor belt of share issues to meet the demand from ever more irrational investors by printing more money. The inflationary consequences of this irrational exuberance were inevitable. By the end of 1720 the bubble had burst and Law fled to Venice where he continued to gamble.

Such speculative bubbles were not just limited to France. More than 80

years earlier, Tulip mania had swept Holland in 1637, when at its height some single tulip bulbs sold for more than 10 times the annual income of a skilled craftsman. Meanwhile, across the Channel in England at the same time as the Mississippi Bubble was bursting, investors were pouring money into the South Sea Company, a joint-stock public-private partnership company founded in 1711 to consolidate and reduce the cost of national debt as England sought to expand its trading routes to South America. During 1720, the company's directors pumped and dumped their stock price with fraudulent expectations. Suckers fooled by the frenzy included Sir Isaac Newton, who lost £20,000, equivalent to about £3 million in today's money. Of the episode he is famously quoted as saying, "*I can calculate the movement of the stars, but not the madness of men.*"

If nothing else, the financial bubbles of the 17th and 18th centuries confirmed one thing: that people liked to speculate just as much as they liked to gamble. Indeed, their histories were very much intertwined. Much as the early casinos of the period were seen as places to socialise, London's stock exchange began life in the 1690s coffee houses of Exchange Alley, in particular Jonathan's Coffee House, where stock dealers, or stockjobbers as they were called, would grease the wheels of market liquidity. Then, as possibly again today, they were viewed as nothing more than gamblers. And then, as now, these market makers made use of much the same trading tools: buying on margin (or borrowing), leverage, options to buy and futures contracts. Many of these investment vehicles represent the precursors of modern day derivatives, where secondary value derives from and is dependent on the primary value of an underlying asset, such as a commodity, currency, or security. It was a derivatives bubble that essentially lay at the heart of the 2008 global financial crash, a bubble, as in the 17th century, fuelled by exuberance, greed and herd irrationality. Little, it seems, has changed in 300 years; once gamblers, always gamblers, well, some of us at least.

Condemnation & Prohibition

The condemnation of gambling is probably almost as old as gambling itself; when we consider that it was essentially competing against religion, it is

fairly self evident why. Reuven and Gabrielle Brenner, in *Gambling and Speculation: A Theory, a History, and a Future of Some Human Decisions*, make the compelling case that much of this condemnation has been linked with the idea that people's hopes could (and by extension should not) be ritualised around the idea of chance, embodied in market institutions, rather than Providence, embodied in religious institutions. Both religion and gambling offer means, or rather hopes, of predicting the future. But whilst the latter, for most of human culture, has represented little more than a social pastime, the former, by contrast, embodies an entire codification of ethical principles. Whilst both religion and gambling have competed to ritualise that hope, it was inevitable which direction the condemnation would operate. You'll probably never hear a gambler formally criticising Divine Providence as a mischievous misinterpretation of the laws of probability. Before I delve deeper into explanations for this religious opposition to gambling, I will first present a brief, if inevitably incomplete, timeline of some examples of prohibition through the ages.

For both Ancient Greeks and Romans, a belief in luck was perceived by many as weakening moral fibre. Whilst the Greeks (and their gods) were partial to a game of dice, most authors and philosophers of the time condemned gambling, labelling it a plague and encouraged governments to outlaw the practice. The Romans, too, placed restrictions on when one could gamble. Despite his love of lotteries to raise money, the Emperor Augustus prohibited gambling except during the festival of Saturnalia, in honour of the God of Saturn. At other times gamblers would face heavy fines (equivalent to four times the stake being wagered) and even jail time if caught gambling. Needless to say, such prohibition failed to have the desired effect. Dice playing moved underground and indoors to so-called private clubs.

A Very English Tradition

The prohibition of gambling in England has a long history. Prior to the major cultural and social changes that took place during the Enlightenment, gambling was not so much prohibited on moral grounds but rather

disapproved of in view of its negative consequences, and in particular its influence on military preparedness. In 1388, King Richard II passed a statute to prevent the common folk spending money on idle pastimes like dice, casting of stones, tennis and football, presumably to encourage alternative investment in things like bows and arrows. Nearly a century later (1477) King Edward IV banned the use of houses for games of chance. King Henry VIII, despite being an inveterate gambler himself, echoed his ancestors' distrust of gambling activity when he discovered his soldiers spent more time gambling than improving their archery skills. Nothing much changed until Oliver Cromwell, a notoriously puritanical Protestant, displaced the monarchy during the 1650s and established a republic known as the Commonwealth of England. Amongst other prohibitions including theatres, ale houses and brothels, horse racing and cockfights were banned and gambling dens were closed. Mirroring the Romans, legislation passed in 1657 permitted any loser in a gambling transaction to sue for the recovery of twice the sum lost, whilst gambling debts dating back to 1647 would be declared null and void. Although the Commonwealth of England would last just three more years, much of the new Puritanism foretold what was to follow under the Victorians a couple of centuries later.

During the intervening years, what condemnation and prohibition existed, including the Gaming Acts of Charles II (1664) and Queen Anne (1710), were largely for the benefit of protecting the ruling classes from the ruinous effects and 'enchanting witchery'¹³ of gambling, to prevent the redistribution of their wealth to the rest of society. This was not moral condemnation as such, but more like class protectionism. Nobody likes to be overtaken in status, particularly when only luck plays a hand in that. Neither, however, did it stop those lower on the social ladder from trying, nor prevent those at the top from being complicit in letting it happen. As the new private gambling clubs, recently introduced from continental Europe, began to flourish during the 18th century (for example White's Chocolate House, Almack's and Crockford's in London), so they became the battleground between old and new money, with sharp gambling entrepreneurs ready to fleece wealthy aristocrats happy to engage in irrational 'deep play'. There appeared to be no bounds to how much the ruling classes were willing to engage in this foolishness. Conceivably, the

more a gentleman lost, the more he could demonstrate his aristocratic privilege. For those like William Molyneux, the 2nd Earl of Sefton, and George Stanhope, the 6th Earl of Chesterfield (both ruined by William Crockford, owner of the aforementioned club), demonstrating an aloofness to the value of money was indicative of status and a regard for old fashioned values. Times may have been changing as the Industrial Revolution moved through the gears but the ruling classes perceived themselves to be above all that.

All of this was to change with the ascent to the throne of Queen Victoria in 1837. Already, 11 years earlier, state lotteries, whose control had passed from crown to Parliament in 1699, were finally outlawed. Now the pace of moral condemnation quickened. Increasingly, concern was being expressed about the damaging social effects of gambling. Whilst the aristocracy gambled in private clubs, the growing populations of the working classes were increasingly engaging themselves with gambling pastimes during periods of leisure and recreation that were now becoming separated from work. The new social mobility and the disappearance of former customs were perceived as threats to the status quo. The gathering of large numbers of people attending lottery draws, for example, was feared as providing the opportunities to spark riots and revolutions. Drinking, sex and gambling were perceived as threats to the core Victorian values of parsimony, hard work and abstinence. What better way to impose discriminatory class control than through an appeal to moral indoctrination?

The 1845 Gaming Act deemed a wager unenforceable as a legal contract. Incredibly, this legislation was not finally repealed until 2007. Of course, every prohibition has an equal and opposite unintended consequence, in this case an explosion of betting houses as gambling moved from credit to cash. The 1853 Betting Houses Act outlawed the use of any small property for the purposes of betting, and pushed it onto the street. Finally, after previous attempts by the Victorians to penalise those congregating for the purposes of betting, the practice was finally banned outright in 1906 through the Street Betting Act. Now all but the very wealthy were prohibited from gambling. Whilst such legislation was largely ignored by the working classes bent on gambling, it was not until the 1960 Betting and Gaming Act that gambling by the masses was again officially recognised as something that could be carried out legally in the United Kingdom.

Sadly, the perception that idleness and moral fog had descended upon the working classes of Victorian Britain missed the bigger structural picture. The latter half of the period marked the beginning of the decline of the nation as the global economic powerhouse, firstly due to the deteriorating entrepreneurial spirit amongst the more traditional colonialist countries, and secondly because of decreasing birth rates and increased life expectancy as a result of improvements in healthcare, with the inevitable consequences for relative economic output. John Maynard Keynes, the 20th century economist, has observed that such gradual change escapes the attention of statesmen and policy makers, to the extent that they would rather attribute foolish and fanatical causes. Far easier to identify, criticise and denigrate a class of society as lacking belief in the legitimacy of the existing social order than to consider the wider global forces at work. Contemporary prejudices towards immigration and its perception as a threat to social and economic stability might today be similarly viewed as lacking sound judgement. Protectionism and prohibition are rarely found to have provided solutions to cultural phenomena. They're not likely to in the future either.

American Schizophrenia

In many ways, the history and evolution of gambling and its prohibition in the New World, and specifically the United States, mirrored that of its motherland. As in the United Kingdom, societal standards and legislation related to gambling have tended to oscillate back and forth from prohibition to regulation. Even today, Americans still can't quite seem to make up their minds about gambling. Professor I. Nelson Rose, recognised as one of the world's leading experts on gambling and gaming law, has identified three waves of gambling regulation during the history of the colonies and the United States. The first wave, lasting for about 250 years, began in the early 1600s with the arrival of the first colonialists, including both Puritans and other English settlers. With the Puritans, who settled largely in New England and Pennsylvania, came the traditional values of abstinence and prohibition. Massachusetts Bay Colony, for example, outlawed the possession of cards, dice, and gaming tables, even in private homes, on the grounds that these habits promoted idleness. In other colonies, English

attitudes towards gambling and recreation prevailed. In the South, the 1710 Statute passed by Queen Anne was adopted, protecting losers of more than £10 but not prohibiting gambling *per se*. Here and elsewhere, gambling was largely considered a harmless diversion. As was mentioned earlier, lotteries began to play a significant part in the financial support of all 13 original colonies to raise money that became a civic responsibility. Indeed, the proceeds helped establish some of the nation's most prestigious universities, including Harvard, Yale and Princeton. Horse racing and its associated betting also flourished (the first racetrack in North America was built on Long Island in 1665), and later casino gambling. By the early 1800s, however, gambling and professional gamblers were coming under increased scrutiny from all the usual accusations: their economic impacts, association with crime and the debasement of morality. Much of it focused on the lotteries and, by 1840, most states had banned them.

The second wave coincided with the onset of the Californian Gold Rush around 1848 to 1855. As prospectors spread westward so gambling habits went with them. Yet public opinion quickly turned against the practice, and the state soon caught up with the rest of the nation, with the California Legislature outlawing most forms of gambling. Typically, however, the prohibition simply drove it underground. The first slot machine was invented and premiered in San Francisco in 1895 and was not specifically outlawed until 1911. Meanwhile, in the South lotteries attempted a comeback, albeit one that was short-lived on account of fraud, bribery and scandal that plagued them. Other forms of gambling suffered the same fate. Horse racing, in particular, was afflicted by the fixing of races by fraudulent bookmakers who sometimes owned the runners and manipulated the odds and the payouts. Under sustained attack from a spread of Victorian values, by 1910 virtually all forms of gambling were prohibited in the United States. As usual, however, those insisting on gambling found other ways.

The Great Depression of the 1930s witnessed a new explosion in gambling fever, and this third wave, ongoing today, led to much greater legalisation. Legalised gambling was viewed as a means of stimulating the economy. Bingo, usually for charitable purposes, was legal in 11 states by the 1950s. In 1931, motivated largely by the economic benefits of tourism, gambling was legalised in Nevada State and Las Vegas, paving the way for the development of super casinos. Horse racing, too, began to make a

comeback. In 1933, Michigan, New Hampshire, Ohio, and California legalised parimutuel (tote) betting. At the same time there was a major effort to crack down on illegal gambling, with many of the new Nevada casinos being financed by the mob. Finally, in 1964, in response to growing opposition, legalised lotteries made another comeback, the first in New Hampshire followed three years later by one in New York and, in 1971, the first financially successful one in New Jersey. In 1978, the state became the second, after Nevada, to legalise casino gambling in an attempt to provide an economic stimulus to Atlantic City.

Yet despite this roller coaster ride of prohibition and regulation, the United States continues to exhibit a kind of schizophrenic attitude to gambling, recognising on the one hand its economic benefits but fearing on the other its moral repercussions. Nowhere can this be seen more clearly than with the Unlawful Internet Gambling Enforcement Act of 2006 which “*prohibits gambling businesses from knowingly accepting payments in connection with the participation of another person in a bet or wager that involves the use of the Internet and that is unlawful under any federal or state law.*” Whilst attempting to present itself as some kind of moral policeman of gambling, its underlying function was surely a form of economic protectionism aimed at looking after the interests of state casinos, race tracks and lotteries, thereby preventing revenues leaving the country. Critics of the Act, such as Michael Shackleford, better known as *The Wizard of Odds*, have observed that it has failed in its primary objective as “*there are ways of funding accounts without using US banks, and millions of players know that.*” Fundamentally, regulation of gambling is always a better alternative to prohibition.

In telling people of lower social and economic standing what they should do with their time and money, the Establishment class perhaps misses a more fundamental point. Whilst we’ll look more closely at the behavioural, psychoanalytical and evolutionary reasons why some of us choose to gamble a little later in the chapter, for Reuven and Gabrielle Brenner, those with limited wealth do so not because they are idle, reckless or immoral, but because other avenues for social and economic mobility have been restricted, frequently by the very class imposing prohibition. Whilst smaller wagers for smaller regular prizes may be made purely for entertainment,

larger ones or those gambled for low probability jackpot prizes are struck because people don't perceive any alternative means of achieving their main concerns: better lives for themselves and their children. This is particularly so when accompanied by the perception of falling behind. Everyone likes to keep up with the Joneses. It's relative, not absolute, wealth that matters to most people.

It is perhaps the financially more disadvantaged, then, who are better able to appreciate the value of money. Certainly, the 'deep gambling' foolhardiness of the English aristocracy of the late 18th and early 19th centuries would appear to support that view. The paternalistic, puritanical attitude of the intelligentsia towards those less fortunate than themselves remains a very typical custom, even more than a century after the Victorians. Why? Such self-serving moral subjugation seeks to endorse a hierarchy of privilege that arises through little more than accidents of history defining the world into which we are born. Prohibition of this kind offers a means of controlling the emotions of envy, frustration and of being left behind, by freeing the envious 'have-nots' from envy whilst absolving the 'haves' from any sense of guilt. To quote the German sociologist and philosopher Max Weber:

"The fortunate is seldom satisfied with the fact of being fortunate. Beyond this, he needs to know that he has a right to his good fortune. He wants to be convinced that he 'deserves' it, and above all that he deserves it in comparison with others. He wishes to be allowed the belief that the less fortunate also merely experiences his due. Good fortune thus wants to be 'legitimate' fortune."

This view of relative, rather than absolute, wealth as providing the motivational engine which drives people to seek and maintain success and standing, even at the expense of others, is something that the work of behavioural psychologists, in particular Daniel Kahneman and Amos Tversky, laid bare in the final decades of the 20th century. We have already seen in the last chapter how it leads people to make systematic errors of judgement when faced with uncertainty, as we are when gambling. It seems that this psychological relativity may account for the very attitudes that we choose to express about gambling as well, whether of acceptance or condemnation. From this, a more general observation then arises: gambling should be seen as a symptom, not a disease (if it is to be seen as such at all). And from that a more general social, political and economic question

presents itself: if gambling and similar means of wealth creation based on chance are to be condemned, what should the fortunate do with those who are not? Of course, the answer to that is beyond the scope of this book, but at the very least this examination should give us pause for thought when we criticise people for merely hoping for a better life for themselves and their children. Not all of us are born to be leaders or astronauts or brain surgeons or Premiership footballers or J.K. Rowling or Sir Richard Branson, and arguably those who are have been just as lucky as a lottery jackpot winner. Not all of us have been blessed with the necessary genes, upbringing and experience that might help us engineer a better life for ourselves, particularly when barriers to social mobility are so prevalent. If games that appeal to chance offer a short cut to scaling the social ladder, who are we to judge that playing them is wrong? If the argument is that so few manage to win and it's immoral to encourage so many to hope then offer people something better instead. Of course, that usually costs money, and those who have it frequently don't like to spend it on those who don't.

Providence versus Fortuna

Nowhere is the self-serving, moral-high-ground, anti-gambling attitude more prevalent than amongst the religious class. That might seem curious given the church's association with bingo, particularly in the United States and with some Catholic denominations, justified on the grounds that it represents but a harmless diversion to help meet the church budget. Yet the condemnation of gambling by religion has a long history that has its roots in theological determinism, a form of determinism which states that all events that happen are pre-ordained, or predestined to happen, by a monotheistic deity. The examples of lots casting that we came across earlier exemplify the point. This was not gambling, an appeal to chance; this was submission to the will of God. Of course, for human beings craving a need to understand the world they live in, demanding causal explanations for things that happen is as natural as it gets. Prior to the Enlightenment, and the birth of probability theory and the science of randomness, pre-ordination by God was the obvious choice. Even now, in the 21st century, armed with the mathematical tools to explain uncertainty, from the law of large numbers

and regression to the mean through to quantum mechanics and game theory, people today still prefer a Divine causality to the spiritless emptiness of unexplained chance. If something happens, it is argued, it must happen for a reason. When more earthly explanations cannot be found, what better reason than God?

All the major monotheistic religions – Judaism, Christianity and Islam – have expressed negative attitudes to gambling and speculation. Ancient Jewish law, for example, held that excessive gambling, in as much as it involved the thoughtless and deliberate redistribution of money, was akin to larceny, and that blind fate should not be the governing force of human destiny. Islam, too, condemns gambling. When the Holy Prophet Muhammad was asked about gambling, he proclaimed that in it is a great sin and some benefit for men, but the sin is greater than the benefit. Early Islamic scholars argued that chess was invented to counter the idea that success could not be achieved through anything other than capriciousness, and condemned gambling games like nard, an early form of backgammon, on the grounds that the Prophet's vision revealed a world with a definite purpose, completely determined by God.

The Bible, whilst specifically never referring to the practice of gambling, is pretty unequivocal in its emphasis on the importance of thrift and hard work and its counsel against covetousness, greed and illegitimate transfers of wealth. Specifically on that last point, the Bible authorises only three morally justifiable ways for money or possessions to pass from one owner to another – labour, exchange and giving – and gambling fits none of them. Indeed, it is often considered to be closer in nature to theft. According to Christian philosophy, one may be paid as compensation for work done to produce goods or services that benefit other people, one may simply agree to exchange possessions with someone else, or one may knowingly choose freely and unconditionally to give something away as an expression of goodwill or kindness with no obligation for the receiver to offer any compensation in return. Evidently, such transfers of wealth are positive-sum, with benefits accrued mutually, either in goods or in kind. Gambling, by contrast, is viewed as zero-sum, where one party gets something for nothing at the expense of another who is obliged to provide it. To be sure, it is argued that each gambler hopes other people will lose so he can take their property, while at the same time he hopes no one will take his property. As

such, this violates the Christian law of exchange. To further emphasise the point, distinctions between things that are gambling and things that are not are frequently made. Stock investing, insurance, owning a business and other types of risk taking like driving or crossing the road are, according to this thinking, manifestly not gambling since no one necessarily wants uncompensated losses to occur.

Condemnation by the Christian church goes as far back as Roman times. Early church leaders threatened excommunication of both clergy and laity found gambling, condemning it on the grounds that it reflected an interest in material things, at the expense of a more blissful spirituality. Yet it was not until after the Protestant Reformation and the subsequent gradual decline in the influence of the church during the Age of Reason that denial of the possibility of chance or fate really became prominent. Ironically, with the new sciences now instructing how and why things happened, it was reason to which many of these Protestant theologians turned to make their case. Whilst such theologians have made frequent criticisms of unproductiveness and the zero-sum mentality of gambling, they reserve their strongest objections to the supplanting of Providence by chance. Appeals to chance, it is argued, lying beyond the realm of reason, must therefore be immoral.

A number of significant texts from around the end of the 19th and early part of the 20th centuries on the morality of gambling exist, including a couple of papers published in the *International Journal of Ethics*¹⁴ and two particularly influential theological works. These are *The Ethics of Gambling* (1893) by William Douglas MacKenzie, an American theologian belonging to the Congregational (Protestant) Church, and *Gambling and Betting: A study dealing with their origin and their relation to morality and religion* (1924) by Robert Henry Charles, a Church of England scholar and theologian who, in 1919, became the archdeacon of Westminster. All four texts seek to argue that gaining property at the expense of another through skill and knowledge amounts to cheating or fraud. Curiously, if that was valid, many individuals engaged in the act of professional competition, including all sportsmen and women, could be considered to be committing crimes¹⁵. Taken as read it would then be irrational for one party with inferior skill relative to another to accept a wager unless some form of handicap was applied to even up the chances of success. Thus, as

MacKenzie concludes, the definition of gambling may be described as a bet, through which “*property is transferred from one to another upon the occurrence of an event which...was a matter of complete chance.*” Having considered the ‘real nature’ of gambling, MacKenzie next inquires into its ‘moral quality.’ Accordingly, since the gambler deliberately chooses to lay aside reason and conscience for the purposes of enjoying the uncertainty concerning some transference of property, such an act must therefore be immoral. Robert Charles, similarly, purports that transferring property by gambling is “*essentially immoral, seeing that it is based on the repudiation of all reason.*”

Readers familiar with David Hume’s¹⁶ guillotine will immediately recognise the fallacy that has been committed. Almost imperceptibly, MacKenzie has made the philosophical leap from describing what **is** to an interpretation of what **ought** to be. Why? What quality of chance, which MacKenzie describes as non-moral, renders its use, as a means of deciding how property should be transferred, immoral? MacKenzie provides a hint: the immorality of appealing to chance lies in the fact that it involves the false proposition that property is itself non-moral. Yet what are the qualities of property that cause it to be something more than non-moral? More generally, George Edward Moore’s¹⁷ naturalistic fallacy articulates the difficulty of trying to extrapolate moral qualities from natural properties. An argument that something is bad because it is large, or heavy, or yellow or random exposes itself to such a fallacy and Hume’s guillotine. Religious critics have argued that the ‘is–ought’ paradox threatens the validity of secular ethics, by rendering them subjective and arbitrary. Of course, as determinists, they would say that, wouldn’t they?

The same critique can be made about the abandonment of reason as the arbiter of decision making. What special quality does it possess that ensures the ‘right’ outcomes for transference of property? Moreover, would those individuals who have either lost the power of reason (such as victims of neurological damage) or are yet to fully acquire it (for example newborns) be behaving immorally? Given that MacKenzie formulated his moral thesis of gambling at a time before the pillars of scientific and philosophical rationalism were yet to be shaken by the tumultuous events of the 20th century that revealed the world to be far less certain and predictable than

had been previously thought, it is perhaps understandable that he adopted such an entrenched view of the power of reason and its ability to sustain moral arguments. Presumably, Immanuel Kant's¹⁸ Categorical Imperative which implied that morality was based on reason alone, which once understood would ensure that acting morally is the same as acting rationally, must have heavily influenced MacKenzie's thinking. Furthermore, MacKenzie's interpretation of 'property', and for that matter the interpretation by most Christian philosophy, is surely profoundly influenced by the concept of stewardship. The Bible makes repeated reference to the idea that we are stewards and hence temporary owners of our Master's property. As Robert Charles articulates, property, for a man, "*is nothing more than a trust committed to him by God.*"

Robert Charles is even more forthright in his condemnation. For him, since our "*early ancestors lived in a world where animism, lawlessness, and un-reason prevailed, and where life was largely non-moral..., to discover the future by an appeal to chance, or to secure his neighbour's goods by a like appeal was undoubtedly natural.*" It follows then, that since "*gambling is essentially an appeal to chance, or the element of the irrational and unknown in life, it...belongs intrinsically to the savage of uncivilised character.*" Of course, more recent work in the fields of socio-anthropology and evolutionary psychology have demonstrated Charles' original premise to be completely flawed. Far from being savage, our human ancestors had been behaving morally for a very long time, behaviour manifestly a consequence of our evolution as a social primate functioning not individually but as part of groups. Game theory, furthermore, has revealed how and why cooperating, behaving altruistically and finally by extension morally works for such social groups, indeed why it is the best behavioural strategy. Contrary to most religious dogma, human morality evolved long before religion, which probably only began codifying moral precepts into moral rules in the last 10,000 years with the invention of written script, the shift from hunter-gathering to more settled agriculture and the formation of larger chiefdoms and states. People have probably been playing gambling games, and trying to divine the future from them, a lot longer than that.

Echoing many other Protestant and Puritanical opinions that have their

roots in the importance of a proper work ethic, Charles sees a fundamental difference between ‘legitimate business’ and gambling, between prudent investment and reckless speculation. “[T]he former seeks to eliminate chance, to use judgment and the rest of man’s best powers; the latter makes its main appeal to chance.” Yet the irony here is that not only is most ‘legitimate business’ at the mercy of the goddess Fortuna but also that gamblers engaging themselves in ‘illegitimate business’ are, erroneously, hoping to subdue her powers, to establish causality and explain why the future happens as it does, not succumb to them. Far from appealing to chance, gamblers are seeking to control it. This is as true for roulette players who have no positive expectation as it is for sports bettors who might believe, almost always mistakenly, that they have one. Charles is quite wrong to insist that gamblers choose to “eliminate so far as is possible the element of reason.” On the contrary, gamblers believe they are exercising it. Their error is not a moral one in appealing to chance, but a psychological one in believing uncertainty like this can be controlled in the first place.

Charles does at least seem to understand that many gamblers subject themselves to the various fallacies of causality, connecting past and future outcomes and seeing patterns where only randomness exists. They are hardly doing this intentionally, however, but rather through ignorance. Presumably for Charles the distinction is irrelevant, since willing or otherwise it demonstrates an ‘inhibition of reason’. Yet if he is to condemn gamblers on the grounds of ignorance or irrationality, we would need to cast the net of moral retribution a lot wider than that. Daniel Kahneman and his colleagues have shown us that being rational, in many walks of life, is slow and tiresome. We are still wired to take faster, intuitive, emotional short cuts to find solutions to problems. Often those short cuts evoke an absence of rational judgement and lead to mistakes, particularly in environments high in uncertainty. Does that make us bad or unreasonable? I don’t think so; it just makes us human.

The departure of reason and a belief in fallacious causality, Charles contends, makes gamblers slaves of superstition and a belief in luck. “It is beyond the wit of man to determine the intellectual havoc wrought by these...superstitions.” But superstition is at least as old as civilisation itself, and probably a lot older to boot. Indeed, Burrhus Frederic Skinner, an

American psychologist and behaviourist, once said, “*If we want to understand the basis of superstition in humans, the best place to start is by looking at the behaviour of pigeons.*” In 1947, Skinner conducted a now famous experiment¹⁹ revealing that hungry pigeons would adopt superstitious behaviour in an attempt to control the pursuit of food. One bird was conditioned to turn counter-clockwise about the cage, making two or three turns between reinforcements. Another repeatedly thrust its head into one of the upper corners of the cage. A third developed a ‘tossing’ response, as if placing its head beneath an invisible bar and lifting it repeatedly. Of course, unbeknown to the pigeons, Skinner all the while was just delivering food at purely random intervals. Skinner’s pigeons were seeing patterns where none actually existed, in the mistaken belief that their behaviour could control the outcome.

Other examples of superstitious pattern recognition in animals have been found in orang-utans and dogs. Rhesus monkeys have even been shown to believe in the hot hand fallacy²⁰, possibly the first instance to be found in non-humans, and will gamble for food accordingly. Researchers at the University of Rochester²¹ devised a fast-paced task in which each monkey could choose right or left and receive a reward when they guessed correctly. The researchers created three types of play, two with clear patterns (the correct answer tended to repeat on one side or to alternate from side to side) and a third in which the lucky pick was completely random. Where clear patterns existed, the rhesus monkeys in the study quickly guessed the correct sequence. But in the random scenarios, the monkeys continued to make choices as if they expected a ‘streak’. In other words, even when rewards were random, the monkeys favoured one side, exhibiting a hot-hand bias consistently over weeks of play.

For the existence of such cognitive pattern recognition, whether deceptive or otherwise, to be so pervasive must surely demonstrate its evolutionary adaptiveness. Conceivably, belief in illusory control might well be beneficial, particularly for food procurement and defence against threats. As Benjamin Hayden, one of the researchers explains, “*if you find a nice juicy beetle on the underside of a log, this is pretty good evidence that there might be a beetle in a similar location nearby.*” Such superstitious behaviour, especially in low probability domains, could lead to enhanced

motivation, more effective performance and greater long term success. Superstitions really represent nothing more than a flawed pattern recognition engine. We are primed to find patterns everywhere, but are not always as good at understanding how or why they are related. According to Glenn Croston, author of *The Real Story of Risk: Adventures in a Hazardous World*, both science and superstition are a consequence of our pattern recognition at work. The difference is that science is a formalised system of harnessing the pattern recognition and putting it to good use, compared to superstition, which dispenses with the formal part, operating all on its own. For a rational atheist submitting to the laws of probability, game theory and Heisenberg's Uncertainty Principle, his conjecture might be that the nonsense espoused by religious determinists in promotion of Divine Providence, and of religious belief in general, itself represents little more than lazy untested faith in superstition. In God, some of us see the most important pattern of all. To others, He just plays dice.

For Charles, “[b]elief in luck or chance is incompatible with belief in God who rules all things according to His Divine will.” Indeed, “the Christian Religion denies that there is such a thing as chance,” and our lives are marked out for us by God. Apparently, we are justified in using the word ‘chance’ only as a substitute for mankind’s ignorance of the ‘infinite realm’. In making such a declaration, Charles was unmistakably mirroring Laplace’s philosophical standpoint, that with sufficient data about the present we could know everything there is to know about the past and the future, and that really there was no such thing as luck. One can but wonder what he would have made of Heisenberg’s Uncertainty Principle published just two years later. Presumably, that too would have to yield to the will of God. If this was all so, one might then wonder why He even bothered to bestow upon human beings the psychology of free will. Thankfully, Werner Heisenberg suggested an alternative world, a world of randomness, where nothing can be known or predicted with absolute certainty, in which probability constructs the very fabric of reality, and where God must surrender to chance, not the other way around. Some might find that world bleak; I find it quite beautiful.

Whether secular or religious in origin, condemnation of gambling arises from the view that its practitioners lack belief in the legitimacy of the existing social order. For governments and the ruling classes that implies a

fear of the undermining of a work ethic and the elevation of money and the quest for material gain at the expense of concern for the common good. For the religious classes moreover, it also implies a fear of the subversion of Divine Providence with the immorality of luck. In both cases, these fears, condemnations and prohibitions will be more pronounced during periods of underlying structural changes in society; when long-established beliefs are challenged, all hierarchies may feel threatened. As such, this represents little more than attempts at controlling the status quo, in maintaining a social pecking order where people should know their place. Given our predisposition to psychological relativity, there is perhaps nothing more natural than ensuring that those beneath us stay where they are. That someone could take our place simply through the draw of a lottery ball, the throw of a dice, the play of a hand of cards or the outcome of a race is surely an affront to the sensibilities of the human spirit. Isn't it?

The question, then, in the face of all this opposition, from governments, from the church and from the wealthy bourgeoisie with their holier-than-thou superiority, is why some of us really feel the urge to gamble at all. To explore the reasons, we must widen the scope of investigation to consider gambling from a wider psychological and evolutionary perspective, beyond the largely subjective and prescriptive theorising of moral and religious philosophy.

Gambling for Control

Reuven and Gabrielle Brenner have made a persuasive case that the less financially comfortable amongst us may turn to gambling in search of a better life. Yet this alone cannot account for the myriad of motivations that might drive such individuals to consider gambling as a behavioural norm. Most gambles, moreover, will not offer potentially life-changing prospects. Indeed, it is probably also true to say that, whilst many who do gamble believe in the ability to manipulate their fortunes, whether through skill or fallacious reasoning, they nonetheless understand that such games come with a negative expected value. Casinos, bookmakers, poker rooms, lotteries, bingo halls, financial trading platforms and mutual fund managers don't, after all, offer their services for free. Surely then it is a paradox why

such gambling exists at all. For *Homo economicus*, perhaps so; we've already seen, however, that most players are not fully rational, and we'll learn a lot more about why in the next chapter.

Various scientific disciplines have had a go at explaining the phenomenon of gambling, including psychiatry, economics, sociology and psychology. Between them a myriad of reasons has been proposed to account for gambling motivations: entertainment, greed, competition, impulsive control disorder, anxiety, depression, loss compulsion, psychosis, personality disorder, void filling; there are probably many others. Clearly, some of these explanations have a concern for the pathological or problem gambler, whilst others will have their origins in the biases and prejudices of those against the practice. Sigmund Freud, for example, reckoned that gambling was a form of self-punishment, arising from an oedipal complex. Another theory, known as the 'Four Es', identifies four psychological factors that put people at increased risk of becoming a problem gambler: esteem, excitement, excess and escape. Most players, however, are not reckless and do not gamble to excess.

For some it's purely about the money. If it wasn't, why would we bother to keep score with it? For others, it's more simply about the excitement, the thrill and the anticipation of winning. For those playing competitive gambling games where success is influenced by the behaviour of people (for example, poker, sports betting and the financial markets), it may be neither about the thrill of the win nor the money acquired, but a demonstration of intelligence and skill, a so-called 'winning with wits,' an expression of 'I'm better than you.' The common theme for all, however, is a sense of goal direction, the desire to bring about a state of mind and being that is more pleasurable or acceptable than was previously the case. Seen in a behavioural context like this gambling, like any other goal-directed action, is then simply the expression of an underlying (genetic) predisposition to seek out life-rewarding outcomes. For most living creatures this has traditionally included food, safety and sex. For human beings with an evolved culture and society, with a much richer and deeper hierarchy of needs, it can mean so much more. A sense of belonging, self-esteem and self-actualisation (achieving individual potential) are all higher level needs proposed by the psychologist Abraham Maslow, although ultimately those secondary goals are fundamentally serving the primary

physiological needs underpinning them.

Sigmund Freud called this the pleasure-pain principle, the instinctual seeking of pleasure and avoidance of pain in order to satisfy biological and psychological needs. He was clearly an admirer of Charles Darwin and understood the theory of natural selection. At the heart of this principle is the need for a sense of control over the maintenance of psychological homeostasis, our state of well-being, organising how we seek out pleasure and avoid pain and how successful we are at achieving that. A sense of control sustains feelings of certainty, helps us understand causality and why things happen, and enables us to predict what will happen next. Perhaps it is unsurprising, then, that the decision to gamble – to predict the future – will be for so many overtly linked with a yearning to feel in control. Gambling to make money; gambling to have fun; gambling to become socially mobile; gambling to escape; gambling to win; gambling to prove intelligence and self-worth; all of them in their own way are expressions of goal-directed behaviour to bring about more satisfying psychological states and a more comfortable existence. Most of the time, and for most people, gambling behaviour is egosyntonic, in harmony with or acceptable to our needs and goals, that is to say, moderate, satisfying and rewarding. It becomes egodystonic, for example through impulsivity and excess, only when our gambling habits come into conflict with our needs and goals, that is to say, self-defeating, maladaptive and pathological. The transition from harmony to dissonance inevitably involves a loss of control. Paradoxically, that arises because the craving for a sense of control becomes the problem rather than the solution.

The need for a sense of control is manifestly an evolutionary survival mechanism, and one that probably goes back hundreds of millions of years. Living creatures possess a strong instinct to survive and pass on copies of their (selfish) genes; taking control, or at least feeling in control, helps them achieve this. According to Glen Croston (*The Real Story of Risk*), human expressions of attempts at control might even include ancient cave paintings. Our hunter-gathering ancestors may well have considered the storytelling such artwork represented as a means of predicting and controlling future outcomes, for example the success of a hunt. To be sure, storytelling and pattern recognition more generally are forms of virtual world simplification to aid with problem solving which is far less taxing for

our brains than a full-blown statistical analysis of data. It's the reason why Garry Kasparov, capable of evaluating just three different positions per second, was initially more than a match for Deep Blue, the IBM computer capable of analysing 200 million positions per second. Even when the stories we tell and patterns we find create an illusory sense of causality and control, this may not be entirely maladaptive. Feeling in control reduces anxiety and a sense of helplessness, enhances motivation, even in low probability or random environments, and potentially leads to more effective performance and greater long term success. Evidently, evolution concluded that the payoff of winning big through luck outweighs the costs of failure and not trying. The randomness of gambling games may not consistently reward luck, but at least we now have an evolutionary explanation for why so many people choose to play them. What cognitive neuroscientist Michael Gazzaniga calls the (left brain) 'interpreter' acts as a kind of 'belief-engine' creating causal narratives about the world which, even when fallacious, help to reduce stress and imbue a sense of control. Natural selection may very well favour strategies that make many incorrect causal associations, including all those associated with gambling, in order to establish those that are essential for survival. So much for the superstitious savage.

Our 'interpreter' may also explain the prevalence, particularly amongst gamblers, of attribution bias, a cognitive bias that refers to the systematic errors made when people evaluate or try to find explanations for their own and others' behaviours. Attributing causes to the events and outcomes we experience provides us with a greater sense of control. When we make a mistake we'll prefer external attribution, attributing causes to situational factors like luck, rather than blaming ourselves. Conversely, when we've made the right call, we'll invoke internal attribution, saying it is due to internal personality factors, in particular our skill and good judgement. The opposite is true when attributing explanations for the behaviour and influence of others. When sports bettors win, the natural inclination is to feel that something they did, for example team or player research, 'caused' their success. By contrast, when they lose, they're more likely to feel that the rub of the green was against them. More generally, attribution error also explains why so many people choose to believe they can win even under conditions of negative expectation, or believe that winning itself is evidence that they have engineered their own positive expectation. Squares playing

games of pure chance, and who suffer from typical fallacies like the maturity of chances and regression to the mean, undoubtedly overestimate the influence of internal attribution when winning, but even games that theoretically offer the possibility of positive expectation encourage their players to do likewise. The psychology of winning breeds an overconfidence that is sadly unwarranted according to real world data. In speculative gambling markets like sports betting and trading, most of what happens, happens for absolutely no consistently predictable reason at all.

Research from an international team of scientists²² suggests that basic survival techniques adapted by early humans may influence the decisions gamblers make when placing bets, and specifically how they rely on their past experiences to predict what might happen in the future. Observing two targets being rapidly illuminated at random, participants were asked in a first experiment to move their index finger to the illuminated target, to test their response times, and secondly to place bets on which target would light up. Consistent with the phenomenon of inhibition of return, in the first instance participants were slower to initiate their movements on a subsequent trial when the target was the same as the previous one. It has been suggested that this effect is an evolutionary adaptation, originating in neural processing, that serves to prevent the return of an organism to a previously explored, and presumably now inadequate, location in space. For example, having collected all the fallen apples from one tree, which strategy serves the maximisation of apple collection: staying at the same tree or finding another one? In the second experiment, instances where participants won were more likely to be followed by a change of target for their next wager, with the likelihood of switching correlating positively with the strength of their inhibition of return. The researchers suggest that early humans developed specialised attentional systems to deal with linear non-random environments, and that these automatic processes are sometimes maladaptive in artificial complex, non-linear and random equivalents. Such findings might indicate that gambling behaviour may be related to hard-wired, basic neurobiological factors concerning how we direct our attention. Put more simply, our predisposition to committing the gambler's fallacy, and presumably other systematic errors involving the illusion of causality and control, may very well be evolved and genetic. In games of chance

where the outcome is completely random, these evolved strategies don't work.

It's easy to see how such overconfidence in predictive ability can arise. If we can predict what will happen in the future, this gives us a lot better chance to control our environment and well-being. In an evolutionary context, failure to predict dangers accurately and associate causes with effects is a threat to our continued existence. Living things genetically better equipped to do so are more likely to survive, reproduce and pass on their abilities; those less so die out. In simple linear environments such adaptive advantage is obvious. In complex non-linear and largely random ones, like betting markets, where predictable signals are weak and the random noise deafening, it can be a hindrance, causing systematic errors of judgement. Possessing certainty is comforting, whilst doubt causes confusion and anxiety. Any wonder, then, that most people are hopelessly fooled by randomness, preferring the psychological reassurance of determinism to the disconcerting insecurity of uncertainty. Perhaps John F Kennedy was one of the exceptions: "*There is nothing more certain and unchanging than uncertainty and change.*"

Igor Kusyszyn²³, a Professor of Psychology at York University, Canada, reflects on this evolutionary interpretation of gambling as a means of securing a sense of control. He concludes that gambling behaviour is an expression of the need to search for meaning in life, even at the risk of financial and social losses. Irving Kenneth Zola²⁴, furthermore, has described how racegoers frequently perceive themselves as engaged in a process of 'beating the system.' Clearly, 'system' means more than beating the odds, but life and fate as a whole. By 'beating the system', outsmarting it by rational means, we exercise control, and become 'winners'. This view of gambling is perhaps best illustrated by the response one successful punter gave Zola to the question of how he picked his winners. "*What do you think I am; a nobody?*" In that single line, we capture what, for many, it means to be a gambler. Gambling confirms our existence and our self-worth. Gambling confirms we are in control of both now and the future. Gambling, through hope and the repeated cycle of reward anticipation and disappointment, confirms that we are alive. Winning, moreover, confirms that we have left the realm of the 'nobody' for the realm of the 'somebody'.

Winning reinforces the feeling that we have influenced our future. Given our need for control, that is a hard feeling to abandon when told that gambling, including sports betting and financial trading and investing, is largely just a matter of chance. When we win it's perfectly natural to assume we must have had something to do with that.

In a complex world full of the harsh realities of life, finding meaning and purpose can, for some at least, be quite a challenge. Irvin Yalom, the existential psychiatrist and emeritus professor of psychiatry at Stanford University, believes that 'existential angst' or 'existence pain', arising from a perceived inner emptiness, emerges from a person's endeavours to cope with the 'givens' of existence, in particular the absence of any obvious meaning or sense to life. If we feel that our lives are subject to a set of forces over which we have relatively little control, behaviours like gambling might conceivably help deny such futility. In so doing it acts as a kind of psychological displacement, an unconscious defence mechanism whereby the mind substitutes a new goal-directed behaviour in place of original ones perceived to be unacceptable or intractable. If we can't find the answers to life, the universe and everything, reassuring ourselves that we know how to pick winners offers a rewarding substitute. That gambling, for almost everybody, offers a futile means of explaining why things happen, is irrelevant. As we've seen, even just the sense that we are in control is beneficial. Perhaps more paradoxical is the need to displace our existence pain with anything else in the first place. If there is no meaning to existence why bother trying to fill that void that it creates?

As an afterthought, one might also hypothesise that those with a more deterministic outlook on life will be less prone to existence pain. Indeed, numerous studies have been carried out to show that people with faith, and presumably a belief in Divine order, are more likely to be happier, have increased life satisfaction and be less prone to gambling addiction. Of course, the latter might just as easily be a consequence of a greater willingness to accept the condemnatory messages, or a fear of the repercussions if they don't.

Risks and the Anticipation of Uncertain Rewards

Taking gambles is a form of risk taking. And taking risks is part of a wider family of decision making behaviour that serves our hierarchy of needs. Evidently, we don't all respond to risk in the same way. Some of us are more prone to seeking risks for the rewards that they bring, whilst others prefer to be more risk averse. The most significant difference in this respect is one of gender. Men tend to take more risks, particularly recreational and financial risks, than women. Check the audience demographics for any online bookmaker, casino or poker room and you will find a vastly disproportionate number of male visitors. From my own personal experience, of the hundreds of sports advisory services that I have followed and verified since 2001, only one was managed by a woman. The fairer sex, furthermore, tend to outperform their male counterparts when it comes to financial investment, on account of their lower propensity to chase riskier stocks and unwillingness to trade as frequently. Trading incurs costs; those who do it more tend to show poorer outcomes. According to Andreas Wilke²⁵, an evolutionary psychologist at Clarkson University and an expert on risk taking and decision making, besides social and cultural reasons there is a biological underpinning that in part drives this sex difference: sexual competition. Essentially, men must advertise their sexual fitness through daring exploits more overtly. Women, by contrast, can be choosier in this context and so are likely to be more risk-averse. Presumably, one way our male hunter-gathering ancestors could demonstrate fitness was through prowess and success at big game kills. That, of course, would entail considerably more risk than the task of gathering nuts and berries, more usually the domain of the female.

Rather than being generally risk-seeking or risk-avoiding, however, people are a complicated blend. A person might be a recreational thrill-seeking BASE jumper but be appalled at the thought of spending money in a casino. Elke Weber²⁶, professor of management and psychology at Columbia University, has attempted to account for this subjectivity with a model called 'domain-specific risk propensity'. Her theory proposes that everyone has a unique risk propensity in each of five categories: financial, health and safety, recreational, ethical and social. Gambling might very well occupy more than one category. A person's risk propensity in one category says little about his or her propensity in another. Furthermore, within a

certain domain, a person's tendency to take risks correlates with how much he or she expects to benefit from the outcome. A skydiver, for example, perceives a greater pleasurable reward from jumping out of a plane than does someone suffering from a fear of heights. Similarly, gamblers with 'winning systems' will choose to risk money that others, having rationally calculated the pointlessness of negative expectation, will avoid. According to Wilke, when people are optimistic about the outcome of their behaviour, they actually perceive it as not being risky. One might even say that they feel like they are in control. The engine for much of this risk-reward analysis is the limbic system, our reptilian brain.

So called because it evolved hundreds of millions of years ago with the first reptiles, the reptilian limbic system is a collection of structures, including the hippocampus, amygdala, hypothalamus and a number of other areas intimately connected to it, for example the ventral tegmental area and nucleus accumbens, that lie at the centre of our brain surrounded by a large cortex. It is the limbic system which directs motivational goal-directed behaviour by sending the sensory inputs necessary for survival to the nucleus accumbens for processing, via the mesolimbic pathway. The limbic system tells us what we need; the nucleus accumbens then directs us to go out and get it. Together, they constitute the 'reward centre' of the brain. Its functioning is controlled by the release of hormones, most significantly dopamine and serotonin. It is the interplay of these hormones and how they are influenced by both experience and genetics that shapes much of how and why people choose to take risks.

As explained by Ronald Ruden in his book *The Craving Brain*, the driving force behind this motivational system is pain. Hormonally, this takes the form of cortisol, a steroid released in response to stress (for example, anxiety due to perceived threats) and low blood glucose (hunger). Stress disturbs homeostasis, our state of well-being, and sparks the limbic system and nucleus accumbens into action, by releasing the neurochemical dopamine, engaging us in goal-directed behaviour to relieve the stress. The greater the stress, the greater the dopamine released. In fight-or-flight situations, it is converted to adrenaline. Dopamine motivates you to get what you need, even when it takes a lot of effort. Dopamine levels are at their highest during the anticipatory, 'trying to get it' phase of goal-directed behaviour. Dopamine, then, is not actually about happiness and pleasure *per*

se, but about the anticipation of pleasure and the pursuit of happiness. Robert Sapolsky, professor of biology, neuroscience and neurosurgery at Stanford University, provides an account of this anticipatory signal in an online video for Fora.tv²⁷. A hungry monkey that spots a juicy fruit at the top of a tree, for example, will experience a dopamine surge as it makes its attempt at acquiring it. Likewise, a craps player will feel a dopamine rush as he throws the dice in anticipation of rolling a 7 or 11. Traditionally, this feeling has been called the craving response.

What happens when a goal is realised? Whilst dopamine levels remain elevated, concentrations of another hormone called serotonin, responsible for regulating mood, begin to rise. As Ronald Ruden says, if dopamine is the ‘gotta have it’ hormone, serotonin is the ‘got it’ variety. A sense of satisfaction is experienced, and with the sense of craving abated, both dopamine and serotonin are in a state of biobalance. These two hormones are complementary, the yin and yang of motivational neurochemistry. In humans, every time our reward centre is activated it builds new neural circuits in the brain. We are not born with circuits for goal-directed behaviour like most other animals. Since our distant evolutionary ancestors started to walk upright about three to five million years ago, the space available for the birth canal shrank. Consequently, newborn babies born today essentially arrive into the world prematurely. To compensate, our brains have become more neurally plastic relative to earlier primates, allowing neurons (the brain’s nerve cells) to form new connections more easily in response to behaviour, in other words to build as we go along. It’s the reason our experiences, and not just our genes, are so fundamental in shaping the person we are. This is particularly so in the earliest years of life. This process of neural circuit building takes place through associative learning, in which a new response becomes associated with a particular stimulus. Not only does dopamine encourage reward seeking, but it helps to store information that can lead you to another similar reward in the future: ‘gotta have more of it.’ A child, for example, will easily learn to associate the sweet taste of an ice cream with feeling good. When ice cream is available again in the future, the learnt association or pattern will be recalled and its dopamine will surge in anticipation of receiving it. Sometimes our responses will be further conditioned by a secondary

stimulus. If the child hears the chime of the ice cream van beforehand, it will quickly learn to associate the pleasure with the sound, and dopamine will begin to surge at the sound of the chime.

Dopamine and serotonin, however, are transitory. When you succeed in triggering these neurotransmitters, the spurt is soon over: the monkey gets its fruit; the child gets its ice cream; the craps player rolls his 7. To get more, you have to do more, but as Loretta Breuning reminds us in *Meet Your Happy Chemicals*, they did not evolve to be on all the time, to create constant ecstasy. They were meant to steer us toward things that promote survival. When an adult tastes an ice cream, it won't experience the same 'wow factor' in the way that a child does. The taste is no longer rare – the adult has tasted ice cream many times – and consequently there isn't the same surge in dopamine. Dopamine responds to new rewards more relevant to the task of survival instead of wasting time on things that are easily available. If we demand constant rewards from them, disappointment is the likely outcome.

The best way to avoid such disappointment, counterintuitively, is to increase the uncertainty of reward. The disappointment lies not so much in the failure to achieve the reward but in the habituation of a guaranteed one. It's why new things: tastes, music, experiences, relationships, you name it, always seem the best the first time round. Professor Sapolsky explains what happens to monkeys where rewards for doing work are only given 50% of the time: their dopamine levels "*go through the roof.*" When rewards are guaranteed, there is no longer any need for the brain to elicit a craving response, since reward expectation can be taken for granted. When rewards are uncertain, a craving response ensures that you keep trying, even after the reward has failed to materialise. Wolfram Schultz²⁸, Professor of Neuroscience at Cambridge University, revealed much the same pattern. Tracking dopamine production in the brains of monkeys when given small squirts of apple juice, he found that when the monkeys received unpredictable squirts or larger ones than they were anticipating, dopamine production surged. In contrast, when they received the juice they were expecting, dopamine-producing neurons remained inactive. According to Schultz, a reward that's unpredictable typically counts three or four times as much. Dopamine neurons are in effect operating as 'prediction neurons.'

It's easy to see how the unpredictability of gambling can hijack the brain's reward centre. Games of chance prey on this neural system. Whenever you win some money, the reward centre constructs a predictive pattern for the purposes of anticipating future rewards; the higher the dopamine, the stronger the predictive pattern. Of course, most gambling games are purely random. Roulette, dice and slots are purposely designed that way. Sports, poker and finance pretty much behave as if they are. Humans, however, like other animals, don't usually want to see it that way. Instead, we prefer causal and predictable explanations for things, and now we have a neural explanation for why. Instead of switching off, of getting bored by the haphazard payouts, our dopamine neurons become obsessed. The random rewards of gambling are much more seductive than a more predictable or even guaranteed reward cycle. Playing noughts and crosses, for example, soon becomes boring because the outcome for anyone beyond a certain age is fairly predictable. Remember Skinner's pattern-recognising pigeons? They were experiencing the same dopamine surges that we do in anticipation of uncertain rewards. In fact, the greatest dopamine response occurs where probability of payoffs for uncertain outcomes is close to 50%. It's surely no coincidence, then, that many games of chance offer 50-50 propositions. Clearly, the market setters for football Asian handicaps and American point spreads knew they were on to a good thing.

Neurologically speaking, then, gamblers are playing to play rather than playing to win, and monetary gains are perceived as opportunities to extend the duration of play. According to Patrick Anselme²⁹, a behavioural neuroscientist, paradoxically it also seems that the disappointment of losing is more attractive than the thrill of winning. Near misses enhance the motivation to gamble and stimulate the reward centre into repeated action just as much as winning³⁰. Such a mechanism might conceivably provide an explanation for the prevalence of loss chasing. The opportunity to predict, to explain and to control outcomes evidently matters more than the outcomes themselves. Sapolsky calls this the addictive power of 'maybe.' Some people have even considered the dopamine responsible for it to be the most evil chemical in the world³¹. Perhaps some of us don't become addicted to gambling, but to dopamine instead.

Pathological gambling is certainly maladaptive behaviour, but the

attractiveness of uncertain rewards is so widespread in the animal world that this tendency should surely have an adaptive origin. Anselme calls this evolutionary explanation for our inherent gambling predisposition the compensatory hypothesis. According to this hypothesis, if reward uncertainty was not a source of motivation most predictive behaviours would extinguish because of the high failure rate. In other words, allowing an animal to persevere in a task is only possible if its behaviour is motivated by a lack of predictability rather than the reward itself. Such an explanation should also account for the attraction of losses. Without the possibility of losing, taking risks to acquire rewards becomes dull. Assuming this interpretation is correct, gambling behaviour in humans exists because the motivation of reward uncertainty was an evolutionary winning strategy. Far from being deviant, gambling and risk taking more generally is a behaviour expressed by the winners of evolution. Perhaps the biggest paradox of all is that whilst we crave certainty, explanation and meaning, it is actually uncertainty and chance which ensures we keep demanding it.

In the comfortable world many of us live in today, the motivation of reward uncertainty is no longer that essential, and the numerous opportunities available to gamble may be hijacking an evolutionary designed adaptation to aid survival. Whether, as for most, our gambling remains egosyntonic and rewarding, or for a few becomes egodystonic and pathological, may in part be down to our genes. Much interest in this field stems from the pioneering work of Dean Hamer, an American geneticist, who in 1996 published an investigation into a possible connection between a particular gene responsible for dopamine regulation, DRD4, located on the 11th chromosome, and thrill seeking³². Specifically, DRD4 gene polymorphism (variations to typical DNA sequences that are more common than mutations) that affects the number of its copies translates into lower dopamine sensitivity. Those individuals showing DRD4 polymorphism are consequently found to engage in more thrill-seeking behaviour to get their required dopamine fix. A word of caution: the association between dopamine receptor DRD4 polymorphism and thrill seeking accounted for just 4% of the observed variation in actual behaviour. Furthermore, whilst a 2006 twin-study into sensation seeking provided additional evidence of its

heritability³³, further research has not been able to fully replicate Hamer's original findings. Other investigations have looked at different genes involved in regulating dopamine, whilst still others have investigated genes responsible for the transport of serotonin and its influence on how we choose to gamble. Finally, the propensity to make optimal financial decisions under conditions of uncertainty might even be linked to the so-called 'warrior gene'³⁴, a monoamine oxidase A gene variant (MAOA-L) that has been linked with psychopathy. The possibility that 'bastards', as Paul Zak, author of *The Moral Molecule*, calls them, are more likely to gamble and invest aggressively will probably not come as a surprise to many. In general, however, as for Hamer's work, the genetic mechanisms proposed to influence gambling behaviour usually only manage to explain a small proportion of behavioural variation. In summary, whilst it is highly probable that the propensity to gamble is hereditary, our environment is just as likely to influence whether and to what extent we choose to do it. Remember, our brain's neural connections are being constructed and reinforced by learning experiences all the time. We might be born with predispositions to behave in certain ways, but the events of life ultimately determine much of how we actually choose to behave.

Markets for Opinions

At the start of this chapter, I set out to investigate why it is that gambling has so often been condemned despite being part of human culture, in one form or another, for such a long time. Such condemnation has taken many forms but essentially reduces to two commonalities: something for nothing and appeal to chance. Both are considered morally wrong and form an intrinsic part of religious philosophy. Both, in my opinion, are fallacious. We have learnt that gambling has often been perceived, by both governments and religions, as posing a threat to the natural social order. The relativity of social hierarchy comes naturally to people; those at the top wish to remain there; those at the bottom yearn to move up. For the latter, gambling, whether rationally or otherwise, may be perceived as a goal-directed objective for controlling the future and possibly achieving a better life. For the former, gambling's prohibition offers a means of controlling

and oppressing such idle dreams.

For hundreds, possibly thousands of years, theologians and religious philosophers have made God an integral part of the moral process, creating and sanctioning moral principles that acquire a transcendental or absolute quality. By the Enlightenment, however, philosophers like David Hume started to take God out of the moral equation altogether, grounding moral principles in natural law. The danger with such an approach is to risk committing the naturalistic fallacy, by reducing things that morally ought to be purely to things that are, that is to say, explaining moral facts simply in terms of facts about nature. Hume understood the paradox this created, as I think did many of the evolutionary theorists, beginning with Darwin, who followed over the next couple of centuries. They solved this paradox by revealing that the act of behaving morally rather than the stuff of morals *per se* is but one of many evolutionary adaptive behaviours adopted by humans to solve problems and conflicts associated with survival.

Moral sentiments, according to Michael Shermer, author of *The Science of Good and Evil* and *The Mind of the Market*, do not originate in the transcendental nature of God but evolved through the process of natural selection. The moral sense of doing or being ‘good’ evolved out of behaviours that were selected because they were beneficial; the immoral sense of doing or being ‘bad’ evolved out of behaviours that were selected because they were detrimental. This evolution of a moral sensibility, beginning with cooperation and reciprocal altruism – you scratch my back and I’ll scratch yours – took place within the confines of groups, since from the very start our human ancestors have been essentially social creatures. 20th century game theory has since showed why cooperation was such an evolutionary stable and winning strategy. If life is considered to be an iterative series of dilemmas involving problem solving by individuals in a group, then cooperating and, by extension, behaving altruistically and ultimately morally will deliver the best ‘scores.’ Curiously, our selfish genes have been the source of something inherently very selfless.

So what does all of this have to do with gambling? In *The Origin of Virtue*, Matt Ridley champions the idea that cooperation and by extension moral sentiments evolved via the division of labour and egalitarian food sharing. Within hunter-gathering groups, individuals performed different roles according to their expertise, and subsequently shared the spoils

communally. Why? For hunter-gatherers, food, and particularly meat, represents a scarce resource and its acquisition is often a matter of luck. Even the most skilful of hunters will often come back empty handed. Sharing the workload as well as the spoils thus spreads the risk as well as the rewards of hunting. The sharing of food represents a kind of reciprocity in which one individual trades his current good luck for insurance against **future** bad luck. The basis of such reciprocity can still be seen in the most universal of moral maxims, the Golden Rule: treat others as one would like others to treat oneself. Such behaviour is not unique to humans, or even other primates. Even vampire bats do it when granting their neighbours a share of their blood meals. Like humans, they have sophisticated mechanisms for detecting and punishing free-riders. Unsurprisingly, this theory is called the risk-reduction hypothesis of food sharing. Spreading risks and rewards like this, of course, is a means of controlling uncertainty, and making future outcomes more favourable. And attempting to control uncertainty is fundamentally the rationale behind risk taking and gambling.

Matt Ridley sees in this behaviour of our ancestral hunter-gatherers distant echoes of the origins of modern derivatives markets, with its futures, options and swaps, and the phenomenon of hedging. Sophisticated and incomprehensible as they may seem, derivatives trading is merely a system for managing risk. The futures contract, for example, evolved many centuries ago from the interaction of wheat farmers and dealers looking to secure a fair price. The farmer would agree with the dealer to sell his crop six months hence at a fixed price. Such a contract was beneficial to both parties. The farmer knew how much he would be paid for his wheat, and the dealer knew his costs in advance. The farmer might miss out on some profit if wheat prices rose, but the contract would protect him against the risk of falling prices. Conversely, the dealer, facing the opposite risk, has hedged against the possibility of rising prices. Such a contract involves the sharing or trading of risks: the seller passes along the risk of lower prices to the buyer; the buyer transfers the risk of higher prices to the seller. By engaging in such cooperation, the overall risk to both parties is actually reduced. Whilst in recent times the ‘casino gambling’ of the global derivatives market has come to symbolise for many the destructive side of capitalism, we would do well to remember that its *raison d'être* is simply one of minimising the effects of uncertainty and the hedging of future bad luck;

surely that's a good thing?

Seen from a perspective of evolutionary problem solving, Mary Midgley, the English philosopher, imagines the business of morality to be a system of conflict resolution. Without conflicts, she argues, morality could never have arisen. Ironically, it is moral absolutism, and in particular religious absolutism of the kind espoused by MacKenzie and Charles, which leaves us with conflicts unresolved and no moral compass. Indeed, absolute moral laws probably tell us more about the law maker than the law breaker. On the contrary, morals, like good scientific hypotheses, should always be seen as provisional, temporary, falsifiable and, borrowing from statistics, Bayesian. Bayesian inference is a statistical method which begins with the specification of some prior probability; this is then updated in the light of new evidence. In a Bayesian world, there is no such thing as absolute truth or certainty, right or wrong. These things are merely provisional; to be discussed, debated, disputed and ultimately improved upon iteratively as we acquire more information, but never to be concluded. The founding fathers of probability theory, like Jacob Bernoulli, clearly understood morality in these terms. The Latin phrase *moralis certitudo* (literally, moral certainty) was first used by the French philosopher Jean Gerson around 1400, to provide a basis for moral action that falls short of absolute or mathematical certainty. For Jacob Bernoulli, if one cannot have absolute certainty in a decision, it may be possible to have a very high degree nonetheless. A morally certain event is an event whose probability is nearly that of a certain event. For Jacob, moral certainty meant a probability of 999 in 1000.

Fundamentally, then, the business of morality does not stem from absolute and Divine decree, but represents a market for opinions, where views about how things are and should be can be debated and traded. Religion, as a codification of moral sentiments, and the invention of a deterministic God were the predictable end products of a biological and cultural evolution of a species that became aware of its future, learnt to interpret the world through causality and discovered, through the sharing and spreading of risks, how to better cope with the unpredictable. God did not make man in his own image; it was the other way around. Gambling, too, is a market for opinions, where views about the future – the price of wheat, the value of a company, the winning of a game, the spin of a wheel or the throw of a dice – can be traded amongst players. Gambling just

happens to make the quantification of opinions, with money, explicit. So both our moral sense and our urge to contemplate, wonder, speculate and finally gamble represent evolutionary beneficial adaptations that help us express our beliefs about things, manage risks that we face and ultimately control uncertainty. What is more natural than attempting to take care of our own future?

Of the criticism that gambling indulges a something-for-nothing ideology, we can now say this: far from being zero-sum, a market of opinions represents a positive-sum exchange. Granted, in a financial sense at least, gambling necessarily involves a redistribution of monetary advantage with winners paid for by losers. Yet, evidently, there is so much more to gambling than this. Physiologically at least, we now know that to be true. The feelings of hope, anticipation, confirmation, success (and even failure) and control are all things that can be measured neurochemically (via dopamine). Through the consensual exchange of opinions, parties can engage in a positive-sum trade of beliefs about how they think the future will be. Ultimately, some parties will be wrong, but the emotional gains from this experience are potentially no less rewarding. If the purpose of gambling is to achieve authority over uncertainty, to feel in control of one's destiny, surely everyone who plays sensibly and reasonably is a winner.

[12](#) The overround, vigorish, juice, cut, take or margin, depending on your origin of culture, is the amount charged, via the shortening of odds relative to their 'true' probabilities, by a bookmaker for taking a bet. I have discussed the nature of the overround at length in my first two books *Fixed Odds Sports Betting* and *How to Find a Black Cat in a Coal Cellar*, and refer any reader still not familiar with the concept to these.

[13](#) A phrase used by Charles Cotton in his 1674 book *Compleat Gamester* with instructions on how to play gentle games such as cards, dice and racing to name but a few.

[14](#) Hobson, J.A., 1905. The Ethics of Gambling. *International Journal of Ethics*, **15**(2), pp.135-148 and Freeman, F.N., 1907. The Ethics of Gambling. *International Journal of Ethics*, **18**(1), pp.76-91.

[15](#) Perhaps Marxist ideologists who prioritise equality of outcome over equality of opportunity might agree.

[16](#) David Hume (1711-1776) was an 18th century Scottish philosopher, historian and economist, known for his radical philosophical empiricism and scepticism.

[17](#) George Edward Moore (1873-1958) was an English philosopher and one of the founders of the

analytic tradition in philosophy.

18 Immanuel Kant (1724 to 1804) was a German philosopher whose major work, the *Critique of Pure Reason* (1781) aimed to explain the relationship between reason and human experience, and in doing so hoped to put an end to speculative and sceptical theories of thinkers such as David Hume.

19 Skinner, B.F., 1948. Superstition in the Pigeon. *Journal of Experimental Psychology*, **38**(2), pp.168-172.

20 The hot-hand fallacy is the fallacious belief that a person who has experienced success with a random event has a greater chance of further success in additional attempts. It was first discovered and proposed in 1985 as a phenomenon by Amos Tversky and colleagues during a statistical analysis of basketball shooting success.

21 Blanchard, T. C., Wilke, A. & Hayden, B. Y., 2014. Hot-hand bias in rhesus monkeys. *Journal of Experimental Psychology: Animal Learning and Cognition*, **40**(3), pp.280-286.

22 Lyons J., Weeks D.J. and Elliott D., 2013. The gambler's fallacy: a basic inhibitory process? *Frontiers in Psychology*, **4**:72.

23 Kusyszyn, I., 1977. How gambling saved me from a misspent sabbatical. *Journal of Humanistic Psychology*, **17**(3), pp.19-34.

24 Zola, I. K., 1963. Observations on gambling in a lower-class setting. *Social Problem*, **10**(4), pp.353-361.

25 Wilke A., Hutchinson J.M.C., Todd P.M., Kruger D.J., 2006. Is risk taking used as a cue in mate choice? *Evolutionary Psychology*, **4**, pp.367-393.

26 Weber, E.U., Blais, A-R., & Betz, N., 2002. A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviours. *Journal of Behavioral Decision Making*, **15**, pp.263-290.

27 For the full video visit: http://fora.tv/2011/02/15/Robert_Sapolsky_Are_Humans_Just_Another_Primate For the relevant excerpt visit: <https://www.youtube.com/watch?v=axrywDP9Ii0>

28 Schultz W., Dayan P., Montague R.R., 1997. A neural substrate of prediction and reward. *Science*, **275**, pp.1593-1599.

29 Anselme, P., 2013. Dopamine, motivation, and the evolutionary significance of gambling-like behaviour. *Behavioural Brain Research*, **256**, pp.1-4.

30 Chase, H.W. & Clark, L., 2010. Gambling severity predicts midbrain response to near-miss outcomes. *Journal of Neuroscience*, **30**(18), pp.6180-6187.

31 http://www.chemistryviews.org/details/ezine/1340629/Could_Dopamine_be_the_Most_Evil_Chemical_in_the_World.xhtml

32 Benjamin, J., Li L., Patterson C., Greenberg B.D., Murphy D.L. & Hamer D.H., 1996. Population and familial association between the D4 dopamine receptor gene and measures of Novelty Seeking. *Nature Genetics*, **12**, pp.81-84.

33 Stoel, R.D., De Geus E.J. & Boomsma D.I., 2006. Genetic analysis of sensation seeking with an

extended twin design. *Behaviour Genetics*, **36**(2), pp.229-237.

³⁴ Frydman, C., Camerer, C., Bossaerts, P. & Rangel, A., 2010. MAOA-L carriers are better at making optimal financial decisions under risk. *Proceedings of the Royal Society B: Biological Sciences*, **278**, pp.2053-2059.

THE THREE Rs: RISK, REWARD AND RATIONALITY

How much is something worth? More particularly, how much is something worth if we have to risk something to get it? This is the question at the heart of all decision making under conditions of uncertainty. There is no such thing as a free lunch. To get something out, we have to put something in. For the sky diver, the thrill of jumping out of a plane must be weighed against the risk of his parachute not opening. For the smoker, the pleasure and relaxation given by the nicotine must be measured against the future possibility of cancer. For the stock investor, the hope that his investment will increase in value can be balanced against the risk that it will fall. For the roulette player betting red, his risk is that it will land black. Trying to measure the ability of something to satisfy needs or wants, particularly when faced with the prospect of losing, is the business of utility.

In the simplest sense, economists consider utility to be a measure of how much someone is prepared to pay for something. For 250 years the science of utility was grounded in rationality because, it was argued, human beings are profoundly rational agents. You wouldn't pay more for a banana than an apple if you liked apples more than bananas. If you did, you would be considered neurotic. It turns out that we are actually less rational than we first thought, but more about that later. Rationality, or rather the perceived lack of it, formed the basis of religious condemnation of gambling. For economists more generally, the zero-sum nature of most gambling games presented a headache. If gambling has an expected value of zero or worse, why is it that people still choose to gamble? Of course, we've seen that the utility, that is to say, the perceived likability, of gambling means more than simply money. If, for example, it was measured in dopamine, losing as well as winning would have a positive utility expectation. Evidently, expected value and expected utility don't mean quite the same thing.

Expected Value

The concept of expected value was first considered by Pascal and Fermat when attempting to solve the problem of the game of points for their compatriot Chevalier de Méré, who wanted to know how a prize pool should be divided amongst players in the event of an unfinished game (see the chapter ‘Cleopatra’s Nose’). The starting insight for Pascal and Fermat was that the division should not depend so much on the history of the game played thus far, but rather on the possible ways the game might have continued had it not been interrupted. In other words, what was important was not the number of rounds each player had already won, but the number of rounds each player still needed to win in order to achieve overall victory. In thinking about future events that had yet to take place, the idea of mathematical expectation was founded.

If I roll a dice, what should I expect to happen? For one that is unbiased there are six equally possible outcomes, each with a one sixth probability. Multiplying each outcome by its probability and summing those products returns a value of 3.5³⁵. In other words 3.5 is the weighted average of all the possible outcomes, where each outcome is weighted by its probability of occurrence. This weighted average is known as the expected value, value expectation or sometimes even mathematical expectation. More importantly, once we start attaching monetary values to each possible outcome, we can calculate the expected value of profits or losses when gambling. We might call this profit expectation. In European roulette, for example, where the wheel has 37 numbered pockets each with a payout value of $35/1$, the expected value or profit expectation of such a game will be -0.027. Imagine betting \$1 over every pocket, including the zero. 36 of them would lose and just one, returning \$35, would win. For a total outlay of \$37 you have lost \$1 overall, equivalent to 2.7 cents for every dollar you bet. We can calculate expected value for any number of games in the same way, although usually the calculation of probabilities for each outcome will be a little more involved. Richard Epstein’s *Theory of Gambling and Statistical Logic* does just that for all manner of games of chance, including games with dice, cards, numbers and coins, although the mathematics is not for the faint-hearted.

Such games of chance deal with examples of **risk**. The probability distribution of outcomes is known, even if the outcome itself is not. Frequently in gambling, we neither know what the outcome will be nor the probability distribution of possible outcomes. This is the domain of **uncertainty**. Calculating expected value *a priori* for games of psychology like poker, sports betting and financial trading is obviously not possible, since we don't know the true probabilities of outcomes like those which can be calculated for pure games of chance. However, by examining outcomes *a posteriori* (that is to say, after the event) we can estimate what those true probabilities will have been. For example, suppose a sports handicapper bets even money on all his propositions, and after 100 wagers of \$1 each he's returned a profit of \$10 from 55 winners and 45 losers. Assuming good and bad luck to have cancelled out, the implication is that the handicapper has a value expectation of 0.1 or 10% (in other words 10/100). Alternatively, we could say that with a 55% probability of winning any wager, his profit expectation will be given by $55 - 45 / 100 = 0.1$. More generally, profit expectancy (PE) in betting can be calculated by multiplying your probability of winning (p) with the amount you could win per bet, and subtracting the probability of losing multiplied by the amount lost per bet. Since the probability of losing is equivalent to 1 (or 100%) minus the probability of winning, we arrive at the following simplification:

$$PE = po - 1$$

‘o’ represents the European decimal odds made available by the bookmaker. For even money propositions, as in the example here, $PE = 2p - 1$. Sometimes, bettors talk of returns rather than profits. When one wins a bet, the stake is returned with the profit. For unit stakes, the return will simply be the profit plus 1. Hence, we can now also define the return expectation, $RE = po$. Evidently, positive profit expectation can only be achieved where the product of the probability of winning and the betting odds (po) is greater than 1.

‘Odds’ is really just another word for probability. For odds quoted in decimal format, the implied probability of outcome is simply the inverse of the odds. Even money odds of 2.00, for example, imply a 50% probability (or 0.5). Those more familiar with fractional notation more typically used in

the UK would quote this as 1/1. So really what we are doing here is comparing the true probability of outcome with the probability implied by the bookmaker's odds. If the true probability is greater than that implied by the bookmaker's odds, the bettor has achieved a positive value expectation. Unsurprisingly, this is the principle of value betting. Of course, to reiterate, in things like sports and racing we never actually know in advance what the true probability for an outcome is; we can only estimate this retrospectively once many of our wagers have been settled and the role of chance has been minimised. To be clear, because we can never be entirely sure whether good luck has outweighed bad or *vice versa*, estimations for true probabilities are all that we'll ever have.

Profit expectation is the most important number for any gambler, for it informs him about whether he can expect to make or lose money in the long run. The analogous problem in investing is to find investments with excess (risk-adjusted) expected rates of return. Positive expectation is a necessity for gamblers wishing to make a profit once the vagaries of good and bad luck have cancelled out. Of course, actually having positive expectation is quite different to believing one has it. One of the most pervasive cognitive biases gamblers typically suffer from is overconfidence. We'll meet this little devil later in the chapter. If you can convince yourself that you have positive expectation, something that might otherwise be considered irrational can suddenly appear quite reasonable. In the last chapter, we learnt why gamblers can be particularly adept at denying the absence of positive expectation. Having evolved to see patterns in random (or mostly random) domains, and attributing internal causes for them (i.e. explaining outcomes by things that they did), encourages a perception of control that, despite being unwarranted, is nonetheless psychologically advantageous. Gamblers, in the main, vastly underestimate the influence of luck when looking at their history of wagers. Where they see losses they tend to see those as unfortunate, unexplainable and temporary (indeed I've even seen one individual describe them as 'impossible'). Where they see profits, they tend to assume that their predictive skill had something to do with it. Did you know that about 14% of people playing fair even money games could do better than a 10% profit expectation after 100 wagers purely by chance? If you find that surprising, you are not thinking enough about the influence of luck. If your positive profit expectation has been built on good luck, it's

not really a positive profit expectation at all because good luck, in the end, runs out.

Expected Utility

Once the gambler has found, or more usually believes he has found, positive expected value, he must decide how much of his capital to bet. This problem has been of interest since at least the 18th century when it dawned on Daniel Bernoulli, nephew of Jacob (and the inspiration behind moral certainty), that only the foolhardy make decisions about how much to risk based on the **objective** expected value without regard to the **subjective** consequences of the gamble. As Peter Bernstein in *Against the Gods: The Remarkable Story of Risk* puts it, any decision relating to risk must surely involve two distinct and yet inseparable elements: the objective facts (about probability and mathematical expectation) and a subjective view about the desirability of what is to be gained, or lost, by the decision. Today, this subjective desirability is known as utility. Daniel Bernoulli called it moral expectation.

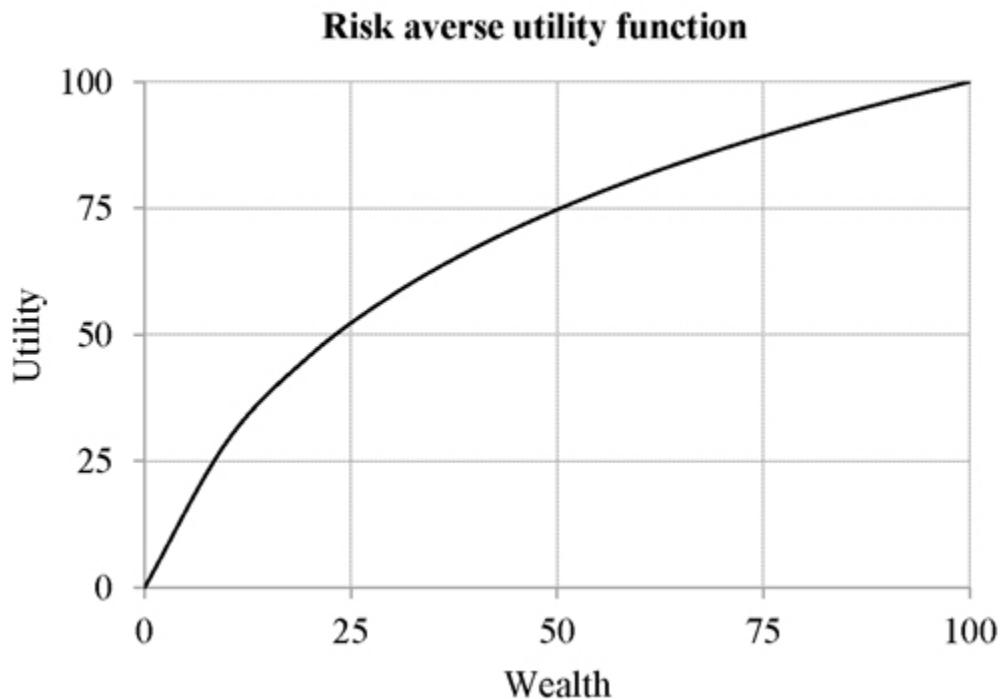
Suppose we are presented with two chests. The first one contains \$10,000 in cash. We know it's in there because we saw it. The second chest contains either \$20,000 in cash or nothing. We don't know which but have been told that both options are equally likely, that is to say, 50:50. You are now allowed to take one of the chests, without touching or weighing either before making your selection. Which one would you choose? This is a classic utility puzzle. Mathematically, both of these chests have the same expected value, that is to say, \$10,000. Assuming you could repeat this game over and over again *ad infinitum* Jacob Bernoulli's law of large numbers would ensure that picking the second chest all the time would grant you an average of \$10,000, exactly the same as the certain outcome of picking the first chest. However, in this game you are only allowed to play once. The law of large numbers does not apply. If you take the first chest, you're certain to gain \$10,000. If you choose the second, what you receive is a matter of chance: be lucky and you'll be \$20,000 richer; unlucky, and you'll receive nothing. Unsurprisingly, given these sums of money, most people choose the certainty of the first chest. From a utility perspective, the

certainty of \$10,000 is surely a lot better than a punch in the face, which is what it might feel like if one had gambled on the second chest and lost, knowing that there was a chance to guarantee a profit of \$10,000. People who find greater utility in certainties than in gambles with the same mathematical expectation are demonstrating aversion to risk.

Daniel Bernoulli reasoned that the standard **rational** behaviour of people when making decisions under uncertainty is risk aversion. Specifically, he hypothesised as people acquire wealth each additional incremental increase is less useful than the one before. He even managed to quantify his hypothesis. “[T]he utility resulting from any small increase in wealth will be inversely proportionate to the quantity of goods previously possessed.” [Bold and underlining mine.] Today, this idea is known as the diminishing marginal utility of wealth. Daniel Bernoulli, living during the period known as the Enlightenment or Age of Reason, was thus heavily influenced by the pervading culture of rationality. The theory of diminishing marginal utility is part of a wider framework of rational choice theory which posits that individuals balance costs (or risks) against benefits with the intention of maximising personal advantage; acting rationally, in this context, means wanting more rather than less of something.

In economics, the marginal utility of a good or service is the gain from an increase, or loss from a decrease, in the consumption of that good or service. The law of diminishing marginal utility implies that the first unit of consumption of a good or service yields more utility than the second and subsequent units, with a continuing reduction for greater amounts. In 1738 when David Bernoulli first formulated his thesis, he concluded that the marginal desirability of wealth, that is to say, that amount by which it changes, decreases in inverse proportion to the wealth already possessed. To a mathematician, that's the same thing as saying that the utility function of wealth is logarithmic. The chart below illustrates a version of the utility function that he calculated. The units for both wealth and utility are arbitrary. Wealth might be measured in money, or apples, or elephants (assuming someone likes having a lot of elephants). Utility might be measured in pleasure, or dopamine, or utils. Similarly, the scaling of 1 to 100 for both wealth and utility is purely for convenience, to help with clarification; we could choose any scale. For each increment of objective wealth that is added, the subjective utility that this brings is smaller and

increasingly so (according to the natural logarithm) each time. In this example, moving from nothing to a wealth of 25 increases utility by about 50; a further increase in wealth of 25 then increases utility by only about half that.



Intuitively, Daniel Bernoulli's theory of diminishing marginal utility feels like it makes sense. Imagine winning \$1 million on the lottery. It would change your life. Imagine, then, that lightning strikes twice and you win it again. Another \$1 million is nice, but not nearly as life-changing as the first. His logarithmic utility function of wealth also provides a mathematical explanation for risk aversion. Looking at the chart, we can easily see how. Suppose we are offered the choice between receiving 50 elephants for sure, or a 50:50 gamble for either 25 or 75 elephants. The mathematical expectation for both is the same, namely 50 elephants. The moral expectation, however, is different. In the first case, the desirability of 50 elephants, according to this utility function, is about 75. For the second case, the desirability of having either 25 or 75 elephants according to an even money gamble will be given by averaging the utilities for those two propositions (in exactly the same way as we calculate mathematical expected value). This gives us an expected utility of about 70, which is less

than that for the sure thing. In other words, it is more desirable to take the guaranteed 50 elephants than to risk only receiving half of them just for the hope of receiving half as many again. Of course, most people probably don't gain much utility in possessing elephants, but you get the picture.

Daniel Bernoulli's theory of the diminishing marginal utility of wealth helped him resolve the St. Petersburg paradox³⁶, a thought experiment concerning a theoretical lottery game with infinite expected value that nonetheless seems to be worth only finite amount to the participants. The St. Petersburg game is played by flipping a fair coin until it comes up tails. When it does, the player wins \$2 and the game ends. If it comes up heads, the coin is tossed again. If on the second toss it comes up tails, the player wins \$4 and the game ends. And so it goes, with the game ending at the first tails, but the available prize money doubling after each successive heads. If, for example, tails came up after five consecutive heads, the player would receive \$64. More generally, the prize available will be given by $\$2^n$, where n is the total number of tosses required to end the game. Obviously $n-1$ heads will have to precede a tails. The expected value of the game is the sum of the expected payoffs of all the possible outcomes. Since the expected payoff of each possible outcome is \$1 (i.e. 50% chance of \$2 on the first toss; 25% chance of \$4 on the second toss; and so on), and there are an infinite number of them, this sum is an infinite number of dollars. The teaser for this thought experiment was as follows: what would be a fair price to pay for entering the game? Appealing to mathematical expectation only, the answer must surely be infinite, hence the paradox; no one, surely, in their right mind would do such a thing, even if they had infinite wealth to indulge such a fancy. Appealing to moral expectation, however, it is obvious that most people would pay only a small sum to play the game. Bernoulli thought it would be about 20 ducats.

Gamblers unfamiliar with the theory of diminishing marginal utility, and that's probably most, may well have come across one of its more practical applications using a money management strategy known as the Kelly criterion. Developed by John Kelly while working at AT&T's Bell Labs in 1956 on solving a problem concerning long distance telephone noise³⁷, it was quickly adopted by gamblers and investors as a means of optimising money management and profits growth. Whilst Kelly's motivation was

entirely different to Bernoulli's, his criterion was mathematically equivalent to the logarithmic utility function. Practically, it directs the gambler to risk a percentage of his overall wealth on a proposition that is both directly proportional to the mathematical expectation and inversely proportional to the probability of success. Recalling that mathematical or profit expectation is equal to $po - 1$ (where p is the 'true' probability of success and o the European or decimal odds received on the wager), we can calculate the Kelly stake percentage (K) as follows:

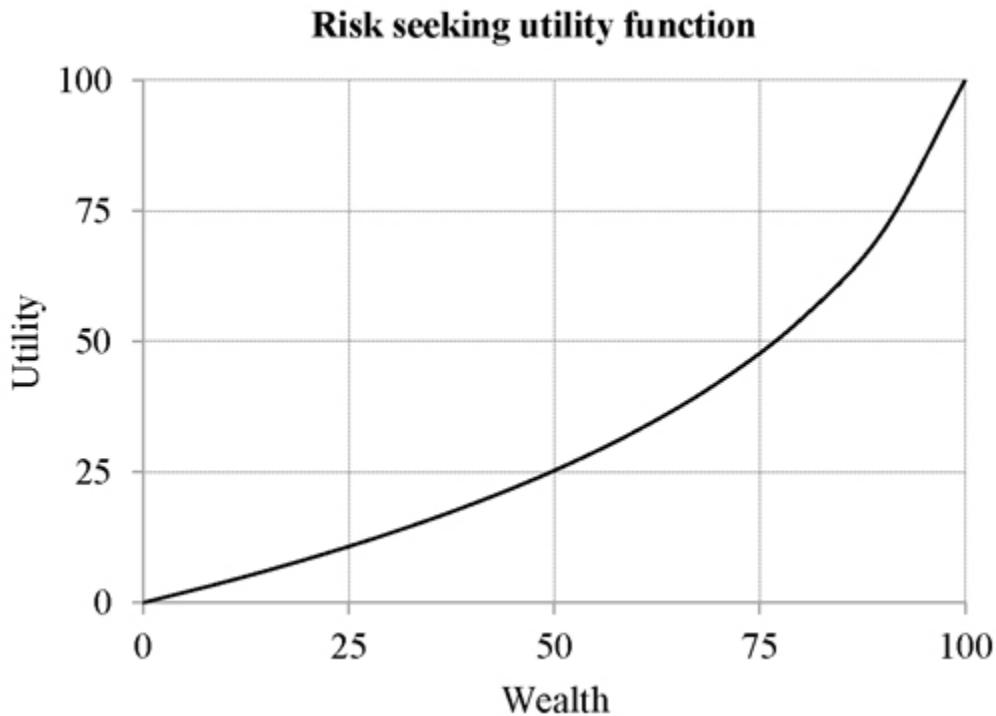
$$K = \frac{po - 1}{o - 1}$$

Essentially, the Kelly criterion maximises expected logarithmic utility. For even money propositions, K reduces to $2p - 1$. A win rate of 55%, for example would imply a Kelly stake of 10% of one's existing wealth. Needless to say, the Kelly criterion only has value under conditions of positive profit expectation. Of course, that would be true for any money management strategy, since losing is not usually in the interests of rational players. Furthermore, optimal growth rate is typically accompanied by a significant volatility in outcomes, a feature that may not best serve everyone's utility. To compensate, some have advocated fractional Kelly percentages as more reasonable options. Additionally, the practical constraints of making many simultaneous wagers need to be given some consideration. How, for example, is one meant to place 15 even money bets each with Kelly stakes of 10%? Nevertheless, Kelly's approach does technically enable winning sharps to maximise the size of their bankroll over the long term. And if it's good enough for Warren Buffett, the American investor and billionaire, it's surely good enough for others aspiring to be as successful as him.

Yet something doesn't seem quite right. Bernoulli's theory of utility implies that zero-sum gambling is a loser's game in utility terms. Qualitatively, that is obvious just by looking at the chart for the risk-averse utility function. Drops in wealth hurt more than equivalently-sized gains. Why would any rationally minded person choose to gamble faced with such grim prospects? Given that some people evidently do, are we to conclude that Bernoulli's underlying assumption that people behave rationally is not

quite correct?

To investigate this, let's return to our original game, with the two chests containing a guaranteed \$10,000 and either \$20,000 or nothing. What if we change the values? Granted, most of us faced with a choice between receiving \$10,000 guaranteed or gambling for double or nothing would opt for the first. \$10,000 is a lot of money. Suppose, instead, this time the choice is between the certainty of gaining \$1 and the 50:50 proposition of winning either \$2 or nothing. Would it make a difference? Probably; for most people \$1 is not a significant sum of money, and one that would not cause that much angst if the certainty of receiving it was forgone. Instead, gambling on winning \$2 would, for some at least, be perceived to be the most desirable option. What about \$10 versus \$20/\$0, or \$100 versus \$200/\$0 or \$1,000 versus \$2,000/\$0? At some point there will exist a crossover point at which the utility of accepting the certainty outweighs the utility of choosing the gamble. Quite where the crossover point will be will depend upon the initial wealth of the person playing this game. For someone on minimum wage with no savings, that might be a three-figure value. For someone like Kerry Packer, the late Australian media mogul and legendary gambler who was worth billions and had an insatiable gambling habit, conceivably we would have to add another 5 zeros at least. Evidently, we aren't always risk averse all the time. For small amounts of money relative to overall wealth, some people will be willing to take a gamble; these people are risk seeking. An example of their utility function is shown below.



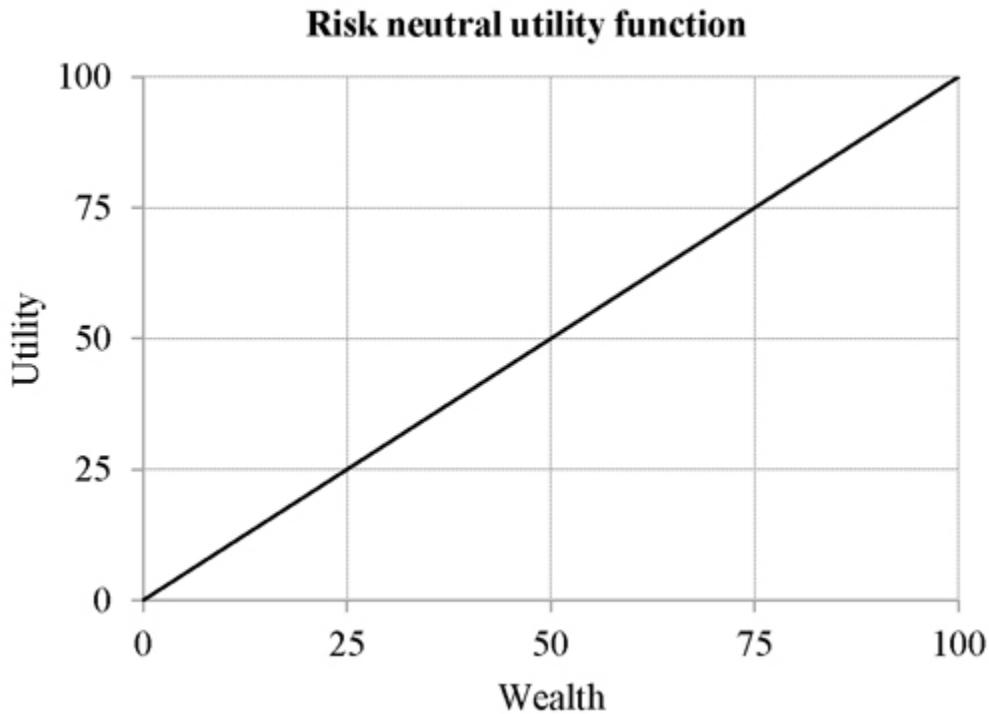
This time, increases in wealth are accompanied by increasing marginal utility. This might seem irrational to some, but clearly such a function can help account for the gambler who prefers the chance to win \$20 or \$200 over the guarantee of receiving \$10 or \$100. This would certainly seem plausible if such a person was already worth \$10,000. \$100 would just add 1% to his wealth. He might as well take a chance on doubling that gain. Similarly, for him the utility of the gamble for either 25 or 75 units of wealth is worth about 30 utils, compared to 25 for the sure thing. Faced with such prospects, the risk seeker will choose to take a gamble.

A useful way to examine risk preferences is via the use of what economists call certainty equivalents. A certainty equivalent is the minimum guaranteed amount of money that an individual would consider as equally desirable to a gamble usually, although not necessarily, with the same mathematical expectation. For example, what is the minimum amount of money you would pay to toss a fair coin and gamble for \$100? The utility charts above help to answer that question. For everyone playing this game, the expected utility should be 50 (utils), since the utility of winning is 100 and the utility of losing is 0, and either is just as likely. For the risk averter (see the first chart) such a game is only worth about \$25. He expects

a premium, in the form of additional expected value, to play this game. It ensures that whatever happens, he'll guarantee that he is at least \$25 ahead in terms of mathematical expectation. In other words, his certainty equivalent is \$25, and represents the minimum amount of money he would rather have for certain, instead of taking some risk. Such premiums, of course, are the stuff of insurance, which is really just a form of risk aversion. Furthermore, Bernoulli's utility function for risk aversion explains how an insurance market can exist. A loss of an asset will cause a greater loss of utility to the buyers than the sellers, since those selling insurance will invariably be richer than those buying it. Buyers, in effect, are transferring risk to the sellers, who are more able to accommodate it. Evidently, insurance, like all gambling, is just another market for risk and opinions about risk.

The risk seeker, by contrast (see second chart), is happy to pay about \$75 for the right to gamble for \$100. His certainty equivalent is thus \$75. Clearly, such a player then has negative net profit expectation, since on average he can expect to win \$50 from the gamble. Rationally, this might seem like a bad bet, but if he's only playing the once he still has a 50% chance of making a profit. Of course, most of gambling is really nothing more than a series of bad bets since the market operators demand a commission for allowing you to play. Yet people buy lottery tickets, play casino games, bet on sports and trade stocks, all of which arguably have negative mathematical expectation. Does that make us irrational for playing, or is the utility from seeking control over uncertainty and our addiction to dopamine such that risk seeking of this kind becomes inevitable? Arguably, because of the illusion of control and the overconfidence it breeds, most gamblers (ignoring those who might actually have a compulsion to lose) may still perceive themselves as holding positive mathematical expectation, being risk averse and therefore behaving rationally, when their unprofitable outcomes suggest something to the contrary. Perhaps Bernoulli was right; perhaps we have mostly risk averse intentions; any expressions of risk seeking might be nothing more than a denial of reality and a refusal to admit that all gambling has, financially speaking at least, negative expectation. We might say that we know this but our behaviour and attitudes to uncertainty suggest that many of us don't really believe it.

Finally, the risk neutral player who is concerned only with mathematical expectation has a certainty equivalent that matches the expected value of the game, in this case \$50. His utility function is a straight line where changes in wealth are linearly proportional to perceived changes in utility. In summary, then, the more risk-averse a person is, the lower his certainty equivalent.



When constructing utility puzzles, as well as changing the size of rewards we can also change the balance of mathematical expectation such that the gamble has a greater expected value than the sure thing. Suppose I give you a choice of receiving \$1 for sure or tossing a fair coin with heads winning \$10 and tails winning nothing. What would you do? Arguably, far more people would gamble for \$10 since the expected value, \$5, is 5 times that of the sure thing. One would have to have a pretty strong risk aversion to decline such an offer. What if the gamble was worth \$100? Really, that's a no-brainer. But what if the choice was \$1,000,000 for sure or gamble for \$100,000,000? Not so obvious now. Alternatively, we could make gambles less attractive, that is to say, give them negative expected value. How about paying \$1 for a raffle ticket with a prize of \$10 and a 1-in-100 chance of winning? I suppose most people wouldn't waste their time on such a small

prize with such a poor expected value (-\\$89). What if the prize was \$1,000,000 instead with a 1-in-10 million possibility of winning? This is the stuff of lotteries. They all have terrible expected values, but the prize available is life-changing. Faced with such a prospect, far more people are willing to take risks than Bernoulli's utility theory had accounted for. Essentially these little thought experiments demonstrate that the business of utility and risk aversion is a very subjective thing depending on the real life consequences for those making such decisions. Inherently, Bernoulli did understand this. His error was in failing to appreciate how irrational, and risk seeking, we can sometimes be. It is true that humans are hardwired to crave certainty and explanation for things, shunning risk and uncertainty, but sometimes that hard wiring can lead to some fundamentally poor judgement, most particularly in random environments.

Despite Bernoulli's errors, his concept of utility has made an immense contribution, not just to economics but also to other ideas about behavioural decision making and theories of rational choice, so much so that it took nearly 250 years before anyone seriously began to question its assumptions. Even as late as the mid 20th century, game theory, the study of strategic decision making under conditions of conflict and cooperation between intelligent rational agents, made utility an integral part of its mathematics. Von Neumann and Morgenstern's utility theorem, published in 1947 as part of *Theory of Games and Economic Behavior*, essentially reproduced Bernoulli's expected utility hypothesis and the principle that marginal increase in utility with increases in wealth is inversely proportional to the amount of wealth already possessed.

The Utility of Gains and Losses

In 1979, two Israeli scientists, Daniel Kahneman and Amos Tversky, published a paper³⁸ that was to fundamentally change the way we think about utility. It was the crowning achievement of a body of work the two of them had collaborated on for the best part of two decades. We met them earlier in the book during our discussion on regression to the mean, and specifically how we're happier to be fooled by it and see patterns of causality that don't really exist. The examination of this and countless other

cognitive errors that we make formed the basis of a large research framework that began to expose the weakness of the underlying assumption of Daniel Bernoulli's expected utility hypothesis: that we are rational creatures.

Rationality can mean many things depending on the context. In purely economic terms, it has tended to imply logical consistency with how we choose to express our preferences and beliefs, specifically with a view to achieving the best possible outcomes. As psychologists, both Kahneman and Tversky had been aware that people frequently opt to behave in ways that are less than fully rational. As Einstein had done with his theories of relativity, they tore up the script and began fitting theory to observations rather than the other way around. If it was no longer possible to make the way people actually behave fit with ideas about how they should behave, it was time to build a new theory. This new big idea was called Prospect Theory; it dealt with how people consider their prospects when making decisions under uncertainty. The most important aspect it considered, something that Daniel Bernoulli seemingly never saw fit to talk about, was the business of losses.

For some of the thought experiments I presented above, it is already apparent that we aren't fully consistent in the way we choose to evaluate risk; most of the time we avoid it but sometimes we seek it, depending on the particular circumstances. Kahneman and Tversky set about demolishing rational utility theory with some more of their own, specifically looking at how we think about losses compared to gains. In Kahneman's book, *Thinking Fast and Slow*, he asks us to consider the following:

For both problems, which do you choose?

- 1) Get \$900 for sure, or a 90% chance to get \$1,000
- 2) Lose \$900 for sure, or 90% chance to lose \$1,000

For each problem the mathematical expectation is the same: a gain of \$900 in 1) and a loss of \$900 in 2). Most people, unsurprisingly, are risk averse in problem 1 but risk seeking in problem 2. Essentially, your response in problem 2 is the mirror image of that in problem 1. Risk averse responses to problem 1) follow classical rational utility theory. Clearly the risk seeking

response in problem 2) is not. Evidently we don't treat gains and losses in the same way. Kahneman's next thought experiment, however, was far more decisive.

3) You have been given \$1,000 in addition to your existing wealth. You are now asked to choose one of two options:

- a) 50% chance to win \$1,000
- b) Get \$500 for sure

4) You have been given \$2,000 in addition to your existing wealth. You are now asked to choose one of two options:

- a) 50% chance to lose \$1,000
- b) Lose \$500 for sure

In terms of Bernoulli's states of wealth, the outcomes for problems 3) and 4) are identical. If you choose the sure thing in either 3) or 4) you will end up with \$1,500 (in addition to your existing wealth). If you choose to gamble, you will end up with either \$2,000 or \$1,000, depending on the outcome. Which did you choose in 3) and 4)? When Kahneman and Tversky experimented with this teaser they found that the majority of respondents preferred risk aversion (and took the sure thing) when faced with the gain in 3) and risk seeking (and took the gamble) when faced with a loss in 4).

For Bernoulli, utility of wealth was all that mattered. Evidently he was wrong, since equivalent statements of the same decision making problem should yield identical choices. Since in this example they don't, respondents were obviously not behaving rationally. The explanation, of course, is that problems 3) and 4) have different starting, or reference points. In 3) it was existing wealth + \$1,000; in 4) it was existing wealth + \$2,000. Kahneman contends that, since few of us pay much attention to these reference points, our attitudes to gains and losses are not, as Bernoulli would have argued, derived from our evaluation of wealth, but simply from the fact that we dislike losing more than we like winning. We are motivated more by the utility of gains and losses than the utility of absolute states of wealth. To use

Kahneman's phraseology, gains and losses are the 'carriers' of psychological value in prospect theory. What really matters is relative, not absolute wealth.

In *Mind of the Market*, Michael Shermer reviews the evolutionary reasons why this should be so. According to the homeostatic model of emotions, we are motivated to behave in ways that aim to restore balance to internal bodily states that have been pushed out of equilibrium. Simply put, we can feel when we are thirsty, hungry, too hot or too cold. We feel too much or too little of something rather than the absolute amount of that thing, and are motivated to respond to the sensations that this relativity elicits. Our neural circuits are finely tuned to detect small (relative) changes in stimuli rather than absolute levels. You can confirm this yourself using three glasses of water, one hot, one cold and the other with a temperature in between. For a minute or so, leave your left hand in the hot glass and your right hand in the cold, before immersing both simultaneously into the one in between. Despite both hands experiencing the same absolute temperature, your left hand will feel colder and your right hand warmer, by virtue of the different reference points each hand started at.

For monetary states of wealth and decisions that concern them, the same emotional circuitry will kick in. As an example, consider this question: would you drive 10 minutes out of your way to save \$10 on a \$25 t-shirt that you had wanted to buy? Now consider another: would you drive 10 minutes out of your way to save \$10 on a \$125 jacket? If you're like many people, you'd be more willing to drive to save the \$10 in the first example than in the second. Why? Isn't \$10 worth \$10? Absolutely, yes, of course it is; relatively speaking, however, not at all. In this example, people determine the value of savings relative to the cost of the items. \$10 seems to be worth a lot more in comparison to \$25 than to \$125. People evaluate choices in relative rather than in absolute terms. More generally, people focus on relative rather than absolute states of wealth, assigning value to things by comparing one thing to another. As for temperature, people do not possess an innate value gauge that determines absolute value.

Wealth relativity provides the evolutionary basis for keeping up with the Joneses, and why most people would rather earn \$100,000 per year whilst everyone else was earning \$50,000, instead of \$200,000 per year whilst others took home \$400,000. Presumably it also explains the (political and

religious) motivations we considered in the last chapter that try to ensure people relatively less wealthy than the rule makers are sometimes treated as rule breakers for finding means of changing the status quo. And finally it might also help people to understand the futility of the hedonic treadmill and attempts to secure ever increasing amounts of wealth. Because we evaluate it relatively rather than absolutely, no amount will ever be enough. Making more money raises expectations and desires, resulting in no permanent gain in happiness. Really that is just another way of expressing the idea of diminishing marginal utility.

When it comes to gains and losses, evolution has also had a part to play. As Kahneman explains, living things that evaluate threats more urgently than opportunities have a better chance of surviving and reproducing. Since we represent the winners in the line of evolution (we are here after all), it necessarily implies that loss aversion is a preferentially selected adaptation according to natural selection. Sure enough, humans are not the only species to display loss aversion. Laurie Santos and colleagues at the University of Yale have spent considerable time investigating the trading and gambling behaviour of capuchin monkeys³⁹. Not only do the capuchins show the same sensitivity to changes in supply, demand and prices but display classic loss aversion when faced with gambles (in this case preferring the certainty of receiving one apple slice compared to gambling for two). In particular, when faced with the possibility of a sure loss, the capuchins preferred to take the riskier option to avoid it. Given the capuchins' inexperience with money and trade, they conclude that loss aversion is probably an innate rather than learned behaviour.

Intrinsically, we are more sensitive to losses than we are to gains. A wonderful demonstration of our aversion to losses was found from the world of professional golf. After studying over two and half million putts from the PGA Tour between 2004 and 2009, researchers Devin Pope and Maurice Schweitzer⁴⁰ observed that a disproportionate number of those for par were completed compared with attempts for birdie. Of course, a putt for making birdie is invariably a trickier proposition than that for saving par – it's usually much longer. Nevertheless, even after accounting for the effects of distance, the golfers still putted 3.7% more shots for par than for birdies. Pope and Schweitzer speculated that the reason was loss aversion. Making

birdie is considered a gain, understandably so because it is one shot better than what the golfer should be achieving. Similarly, shooting a bogey will psychologically be seen as a loss. Given what we know about loss aversion, these findings were probably inevitable.

Intriguingly, we might even speculate that loss aversion accounts for some of the home advantage seen in football. Typically, at the professional level of the game, home teams win close to twice as often as they do away. In psychological terms, points dropped at home will be experienced as a loss, whilst points gained away from home will be seen as a gain. Since we are more sensitised to avoiding losses than receiving gains, the home team will, on average, try harder than their opposition. Of course, loss aversion won't necessarily explain why home teams should feel that dropping points at home should hurt more than dropping points away, although saving face in front of home fans might have something to do with it. Presumably other factors like ground familiarity and increased travel distance for the away team could also account for why such a bias arises in the first place. But once established, loss aversion might offer a credible explanation for why it has continued to persist as strongly as it does.

Loss aversion and the relativity and sensitivity of outcome evaluation are considered by Kahneman to represent the core cognitive features at the heart of prospect theory, and play an essential role in any financial decision making under conditions of uncertainty. They represent operating characteristics of what he calls the fast thinking (System 1) brain, analogous to but not entirely synonymous with our emotional intuitive limbic system. Indeed, research⁴¹ looking at the neural basis of loss aversion has demonstrated an elevated response in the mesolimbic dopamine system as the potential for gains increases, in contrast to a depressed response as the potential for losses increases.

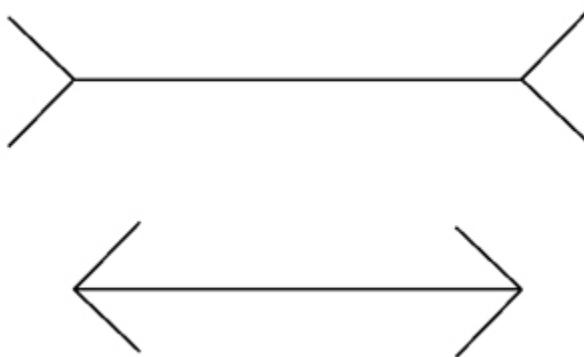
Earlier, we considered whether the disappointment of losing might be more attractive than the thrill of winning. Seen in the light of prospect theory, perhaps 'attractive' was the wrong word. Cognitively speaking, loss can be acknowledged as failure or error, particularly for a species designed to seek control. Whilst losing evidently seems to hurt more than winning, it is the elevated sensitivity to it that motivates us to try again, and to sometimes take riskier decisions in the first place to try to avoid it. Perhaps

people are not so perturbed by uncertainty as much as they hate losing. Evidently, that would fit with the evolutionary theory that we are neurochemically motivated more by the uncertainty of rewards (and losses) than sure things; this is the addictive power of ‘maybe’. We are programmed to gamble, but just don’t like losing, and when we do, we are motivated even further to regain control.

It turns out that we are about twice as sensitive to losses than we are to gains. When offered 50:50 propositions to win some money, players typically require the reward to be about twice the size of the potential loss to make the gamble worthwhile. That’s a pretty significant mathematical expectation (0.5), and one you won’t find in too many gambling outlets. Why is it then that so many people still choose to gamble? Well, partly because those who do may not be entirely typical, but partly also because other cognitive biases operating via System 1 create illusions that encourage us to believe we have more control over outcomes than we really do; illusions that Kahneman’s System 2 brain (analogous to the pre-frontal cortex) is too lazy to overcome, despite its capacity for rationality.

Cognitive Illusions and Biases

Take a look at the next picture and ask yourself which horizontal line is the longest.



Intuitively, you would probably guess the top one. Now measure them both with a ruler, and you can confirm that they are actually the same length. This little psychological test is known as the Müller-Lyer illusion, named

after the German psychiatrist Franz Müller-Lyer who created it in 1889. Now that you know the top line is longer, take another look. Try as one might to reason that both lines are the same length, one of them still ‘looks’ longer than the other. It’s just a perception that won’t go away. There are several explanations for why this illusion occurs. Whilst none is definitive they all share something in common: there are essentially two parts of the brain at work providing an account of how a phenomenon can occur in two different ways. In psychology this is called dual process theory.

The two processes consist of an implicit (automatic), unconscious process, what we might regard as fast, instinctive and emotional, and an explicit (controlled), conscious process that is slower, more deliberative, and more logical. Daniel Kahneman has labelled these two processes or parts of the mind System 1 (the automatic fast one) and System 2 (the controlled slow one). Implicit System 1 processes are usually deeply ingrained, and some of them we may well be born with. Explicit, conscious System 2 processes, by contrast, may change with learning into unconscious, automatic System 1 ones. Such learning processes may be relatively fast, for example via the kind of Pavlovian conditioning we discussed in the last chapter in the context of dopamine reward circuitry building. A lot of this emotional leaning takes place early in life, when neural circuits are plastic and primed for building and rebuilding. Alternatively, developing automatic responses from previously conscious ones can arise via the slower repeated learning process of pattern recognition. That’s how we can all learn to intuitively and almost instantaneously solve problems of basic arithmetic (for example $2 + 2$; compare that to trying to solve 43×89), and do things like ride a bike and drive a car. At first they take mental effort, but through repeated practice they eventually become second nature. For the more professionally minded, that’s how more deliberate or purposeful practice (some have called this the 10,000-hour rule) can help footballers like Cristiano Ronaldo score a goal when the lights are switched out⁴² or tennis players like Andy Murray return balls served at 150mph with such precision. Such pattern recognition in more simple linear environments creates feedback where successful outcomes that can easily be attributed to causes (things that we did) become self-sustaining.

With the Müller-Lyer illusion, our fast automatic System 1 makes the assumption that the lines are of different length. It takes our slower more controlled System 2 to reason, through the use of a ruler, that this is a fallacy. Dual processes are not restricted merely to visual illusions. They can be found throughout social, personality, cognitive, and clinical psychology. Illusions of thought are called cognitive illusions. System 1 is essentially a storyteller. We are primed from birth to construct coherent narratives, and through pattern recognition and associative memory recall we seek causal explanations for what happens and why. This process is automatic and unconscious. It then offers the causal explanations to System 2 which will accept them, often unquestioningly. In 1944, psychologists Fritz Heider & Marianne Simmel demonstrated how even the most abstract schemes can be constructed into stories. Making a short soundless video showing two triangles, one large and one small, a circle and a box with a door, the shapes were animated in such a way to enable viewers to construct a causal narrative of bullying and retaliation. You can watch the video for yourself via YouTube⁴³ and see what story you build. Whilst System 2 informs us that there aren't really shapes engaged in psychological warfare, System 1 creates this emotional illusion nonetheless. As Kahneman explains, Heider's triangles and circles are not really agents – it's just much easier to think of them that way, a kind of mental economy where we interpret the world as a series of intentions. Evidently interpreting the world through storytelling has been evolutionarily advantageous. Narratives that helped our ancestors identify survivalist threats on the savannahs of east Africa presumably provided a better mechanism for not ending up as something else's lunch. Slow objective statistical reasoning that requires the efforts of System 2 would not be much use in such high-risk environments. If better intuition aids survival, then intuitive storytelling will be the naturally selected winner.

Sometimes these cognitive illusions can lead to biases and systematic errors because people are prone to applying causal thinking inappropriately. These biases are bad news for gamblers, where objectivity is essential. Just as it is hard to convince yourself that the two lines above are really the same length, it is hard for people to accept that it's no more likely that a heads will follow 10 tails than it is following just 1. Decisions that gamblers are

faced with making – which team to back, which hand to fold or raise, which stock to buy or sell – require statistical reasoning that takes time and mental effort. Our intuitive System 1 has never developed a capacity for it. System 2 can learn to think statistically and to become aware of such fallacies, but that can take a lot of effort and practice and System 1 will always be running in the background waiting to trip you up. System 2 is inherently lazy when faced with doing a lot of unnecessary work. It often prefers to delegate to System 1 which then engages its own heuristics, associatively learnt or evolutionary hard-wired short cuts taken to reach decisions about how to perceive, judge and behave. Worse still, gambling inhabits a world that is inherently random and unpredictable (whether the casino, the race track or the stock exchange), and opportunities may not even exist for deliberate practice and the acquisition of skill to exploit them. Kahneman labels these ‘zero-validity’ environments. If little or nothing that we can learn about them will have any influence on our interaction with them, then all we are left with are errors, errors that for gamblers cost money.

Despite uncertainty, System 1 is designed to bet on an answer, guided by patterns stored in its associative memory, to offer a causal explanation. It perceives the world in black and white, drawing definitive conclusions. Any wonder, then, that it struggles to cope with Heisenberg’s Uncertainty Principle. Furthermore, System 1 does not keep track of alternatives. It will not even be aware that ambiguity might exist. Minimising the mental baggage that has to be carried saves energy; things that do that are evolutionarily adaptive. System 1 is automatically primed to ‘believe.’ System 2 can be ‘unbelieving’, critical and sceptical, but that takes work. If System 2 is otherwise engaged, we may believe almost anything.

“Swimming in Australia is dangerous.”

Given our fear of sharks and our knowledge of where many of the fiercest species live, many of us would reasonably believe this statement to be true. System 2 must choose to engage in a mentally taxing analysis of shark attack statistics to reveal that there have been barely 1,000 since records began in 1791, with fewer than a quarter of them fatal. And yet, even armed with such information, would some of us still be fearful to swim in an Australian sea? For sure; try as one might, the intuitive fear is too hard to

rationalise away. Kahneman explains this confirmatory bias of System 1 as a preference for uncritical acceptance of suggestions and exaggeration of the likelihood of extreme and improbable events. We prefer to stick to what we believe, rather than change our opinions based on mentally taxing and time consuming evidence-collecting and hypothesis testing. Instead of objectively analysing new information we discover, we instinctively pay attention to certain sources that confirm pre-existing ideas and disregard others that challenge our existing perceptions. This is more bad news for gamblers; any increase in subjectivity can move people away from the most accurate prediction of outcomes and statistical assessments of expected value. Foolishly, I bet Germany would fail to win the 2014 World Cup. Why? Because defensively I thought they were inept, largely off the back of a 5-3 demolition at the hands of Switzerland in 2012. Such a result confirmed my underlying belief that the Germans were still not back to their glory days of past generations, and probably also supported my attitude that Englishmen never like to see Germans winning. They may well have been defensively inept, but going forward they were world beaters, as Brazil learnt to their cost in the semi final.

Basing my judgement on recent information like this was exploiting an availability heuristic, and giving rise to a bias. The availability heuristic is a mental short cut that relies on immediate examples that come to a person's mind when evaluating a specific topic, concept, method or decision. Imagine being asked if swimming in Australia was dangerous in the same week as news networks were covering a story of a shark attack. We exhibit a tendency to overestimate the likelihood of events with greater availability in memory, which can be influenced by how recent the memories are or how unusual or emotionally charged they may be.

Now, what if I said that I was merely talking about swimming pools? You see how easy it is for System 1 to construct explanatory narratives? What matters for System 1 is coherency and consistency, not completeness. You didn't need additional data about shark attack numbers to construct a coherent story in your mind about the dangers of swimming in Australia. And you did it without information about where we were swimming, since stored associations between the words 'swimming', 'Australia' and 'dangerous' inevitably drew you to the conclusion that we were probably talking about sharks. When information is lacking, System 1 substitutes

heuristic short cuts for the purposes of cognitive ease, to jump to conclusions which may often be found to be erroneous. Kahneman calls this ‘*what you see is all there is*’ or WYSIATI for short. More formally, it is an example of attribute substitution. His Israeli Air Force pilot trainers were cognitively short-cutting when failing to recognise regression to the mean, an example of a representativeness heuristic where people assess the probability of a particular event based solely on the generalisation of previous similar events. Other representativeness heuristics include the base rate fallacy, where the mind focuses on specific rather than general information⁴⁴, the conjunction fallacy, where it is assumed that specific conditions are more probable than a general one⁴⁵, and the gambler’s fallacy, which we know all about.

Such lazy belief reinforcement is the basis of conspiracy theory. The ‘Face on Mars’ conspiracy is a case in point. Located in a region on Mars called Cydonia, the Viking 1 spacecraft first photographed it in 1976, depicting what looked like the face of a person, about 2km long. Seeing, as they say, is believing. Its low resolution image was naturally interpreted by scientists as a striking pattern of light and shadow cast by hills in the area. Conspiracists interpreted it as a sculpture carved by intelligent Martians. What they saw was all there was. Later imagery taken 20 years later improved upon the early quality and resolution to reveal its features in far greater detail. The ‘face’ has been near-universally accepted as an optical illusion, an example of the psychological phenomenon of pareidolia where a random stimulus may be perceived as significant. The philosopher Karl Popper argued that the fallacy of conspiracy theories lies in their tendency to describe every event as intentional and planned, thereby underestimating their underlying random nature. In fact, Popper was describing a cognitive bias that psychologists now commonly refer to as the fundamental attribution error: the tendency to overestimate the actions of others as being intentional rather than the product of random situational circumstances.

Gamblers who like to play around with large data sets on the hunt for statistical associations can similarly suffer such illusions of pattern recognition. It’s easy to believe that a statistically significant association between variables of data (for example shots, corners, possession and goals in football) are grounded in causal mechanisms that offer predictive

potential. Two related illusions that are derived from the representativeness heuristic are apophenia, the experience of perceiving patterns or connections in random or meaningless data, and clustering, the tendency to erroneously consider the inevitable streaks or clusters arising in small samples from random distributions to be meaningful. By way of example, consider the following series of wins (W) and losses (L). One series is a random pattern of results from the world of sport whilst the other one is a fake; which is which?

- 1) LWLLLLLWLWLWLLLWWLLWLWWLLWWWWWWWW
- 2) WLLWWLLLWLWLWWLWLWWLLWWWWLLWLWLWW

Did you believe that the long sequences of the same result in series 1) look manufactured? Did you think the shorter sequences in 2) make that one look more random? If so, then you've got the two the wrong way around. In fact, series 1) represents the 1973 to 2004 results for Cambridge University in the Boat Race. Series 2 is just made up. It might look more random on account of its shorter sequences of wins and losses, but in fact it was artificially constructed to be so. If asked to create random sequences like this many people will switch from W to L or *vice versa* if they feel that one of them is happening too often. Long sequences of the same outcome are perceived as being non-random. The hot hand fallacy, where people believe sequences of events have causal explanations that make success more likely to be followed by further success, is another similar example. For gamblers with long winning records who falsely attribute their wins to things that they did, being fooled by randomness like this can prove to be a costly illusion over the longer term.

Many of us also exhibit a tendency to assert that outcomes after they are known were actually predictable all along. This 'knew-it-all-along' effect, or creeping determinism, is known as hindsight bias. Politicians and those in the media are frequently guilty of hindsight bias. In the run up to the 2015 UK General Election, Ed Miliband, the leader of the Labour Party in opposition at the time foolishly accused the Prime Minister and leader of the Conservative Party, David Cameron, of not planning for the deleterious effects of the removal of Colonel Gaddafi as Libyan head of state in 2011 with the assistance of British and French air strikes. Of course, if Mr.

Miliband was truly blessed with such powers of prediction, one might wonder why he chose to vote with the Government to allow air strikes to take place. Similarly, after Germany won the 2014 World Cup it was much easier to convince myself that such an outcome was obvious from the start, given the way they played the tournament. In truth, the confirmation errors which I committed prior to them winning, and the hindsight bias that I exhibited to rationalise my wounded ego in defence against my friend's mockery, were all just examples of poor judgement.

Hindsight bias is similarly related to the *post hoc* fallacy, in which a person erroneously assumes that one event caused another simply because the first one occurred before the second. For Michael Mauboussin, author of *The Success Equation*, the *post hoc* fallacy and hindsight bias ensure that, in complex environments, we have a tendency to overestimate our skill and a tendency to forget about luck. Whilst most people recognise that many more things can happen than actually do, an inevitability of history is formed by our 'belief-engine', through which the full range of possibilities is fused into a single outcome. As Michael Shermer says in *Mind of the Market*, we "connect the dots from our complex and seemingly chaotic world and construct narratives based on connections we think we have found. Whether the patterns are real or not is a separate issue entirely." Building a clear sense of cause and effect, we start to believe that what happened was inevitable and predictable by the existence of our own skill.

Nassim Taleb, trader, statistician, risk analyst, philosopher and author of *The Black Swan* and *Fooled by Randomness* labels this kind of biased storytelling the narrative fallacy. Inevitably, the explanations we put forward to account for events assign a larger role to talent and intention, whilst largely disregarding the influence of luck. Perceiving causality and the illusion of control that it confers may well have an evolutionary explanation: it probably helped save our lives in environments of extreme danger. Jackpot predictions arguably have more utility than average success. As Kahneman argues, being pattern seekers in which regularities appear not by chance but by design was part of the general threat detection and avoidance mechanisms we have inherited from our ancestors. "*Lions may appear on the plain at random times, but it would be safer to notice and respond to apparent increases in the rate of appearance of prides of lions, even if it is actually due to the fluctuations of a random process.*"

Evolution did a good job catering for fast moving environments where choices are made from experience, where linear cause and effect might be discernible and where feedback can reinforce perceptions of accuracy and reliability. This illusion of understanding, furthermore, reinforces the conviction that the future is predictable and controllable, and control, as we know, is comforting, reducing anxiety about uncertainties. In slow moving, quasi-random markets like finance, betting and poker, where feedback is far more limited, however, decisions must be made by description, and these illusions tend to lead to systematic errors which make most of us financially poorer. In such environments, the past is much less knowable than we believe it to be, and so, therefore, is the future. When faced with random food delivery, pigeons resort to superstitious expressions of control. Despite a more evolved System 2 pre-frontal cortex equipped with the capacity to override System 1, we nonetheless still find ourselves committing the same mistakes.

Arising from this illusion of causality is perhaps the most powerful cognitive bias of all: overconfidence. It is one that is particularly detrimental to gamblers. Sometimes known as the Lake Wobegon effect, named after a fictional town in Minnesota, it describes the natural human self-serving tendency to overestimate one's capabilities. In Lake Wobegon, all the women are strong, all the men are good looking, and all the children are above average. The Lake Wobegon effect, where the majority of a group claims to be above average, has been observed in many domains including social popularity, intelligence and driving skill. In one incredible example rating the ability to get along with others, not a single one of the 829,000 school students rated themselves as below average. Talk about overconfidence!

The optimism that people have in their beliefs about the world depends mostly on the quality of the story, its coherence (does it make sense?) based on what they can see. When you win a bet, what do you see? It's easy to let System 1 build the narrative: I picked Manchester United to beat Liverpool; I bet on Manchester United; Manchester United **did** beat Liverpool; I won some money; my prediction caused me to win some money; I have predictive talent. You don't see the potentially limitless number of tiny chance events that dictated how the game evolved and ultimately determined how the result ended. Any of them could have delivered a

different result (remember Cleopatra's nose?). Did your predictive ability really cause you to win money? Or was what you saw all there was? System 1 prefers to settle on a coherent pattern of explanation and suppresses doubt and ambiguity. This self-serving bias leads to exaggerated explanations for success and unwarranted attribution to internal factors like skill, whilst blaming failures on external bad luck. Uncertainty may be neurochemically motivational but ultimately our minds prefer to draw definitive conclusions. We are not coded to think probabilistically, but causally. As Kahneman declares, causes trump statistics, even when they are illusory.

Placing excessive focus on results like this is known as outcome bias, an error made in evaluating the quality of a decision when the outcome of that decision is already known. Specifically, negative outcomes produce more condemnation, typically directed at external agents, whilst positive outcomes produce more praise, usually self-serving praise, even if the outcome is determined by chance. What would you see if Manchester United lost? Conceivably, something completely different; yet the process by which you arrived at the decision to back them was exactly the same. Outcome bias is a particularly powerful one for sports bettors, poker players and casino gamblers to overcome because those games experience fairly swift closure, in contrast to financial investing which is potentially more open-ended. For the most part, you either win or lose, but what contributes to that outcome can be largely a matter of luck. For example, standard blackjack strategy dictates that if you have 16 you should hit rather than stand. Do so and get a 5 and you're a genius; get a 6 and you're a dumb-ass. In both cases the decision to hit was correct, just that they resulted in two completely different outcomes. In the final chapter of this book I'll spend some time examining why it's far more important to focus on the process of decision making, rather than the outcomes themselves. Gamblers who do so are invariably much better gamblers.

Mutual fund managers have been shown to become overconfident following superior investment returns. Subsequently, a self-serving bias has been followed by poorer performance. All that's really going on is regression to the mean, with attribution errors ensuring that the traders assign excessive weight to skill and not enough to luck. I am not aware of any published research into such self-serving bias in the world of sports betting with which I am more familiar, but anecdotally from speaking with

players on public forums it is all too obvious how so many continue to live in a sea of denial about randomness, where most insist that trends exist which work in their favour, but yet where luck and skill are essentially indistinguishable for almost all of them. Later in the book, I'll be reviewing the data that support this viewpoint. When presented with the contradictory evidence that lucky is almost certainly all they have been, it is far easier to massage the cognitive dissonance or mental angst that it creates by simply rejecting the new information that challenges existing beliefs. There will be those reading this who, being profitable from their betting, will continue to deny: "look at my winnings, that's all the proof I need. *What do you think I am; a nobody?*" Such a refutation misses the point. I've never said it isn't possible to make money from betting (quite a lot do, actually), just that it's very hard to perform consistently better than luck will allow, and it's your overconfidence that may have led you to believe you have something more.

Manifestly, by their nature it seems that loss aversion and overconfidence are often in conflict. Indeed, optimism in general may operate as a kind of evolutionary defence mechanism against loss aversion, since otherwise we would never risk anything. The dopamine rushes we receive from believing we have a greater chance of success than we actually do presumably help to increase the chances of having a go. As the saying goes, you've got to be 'in it to win it.' If you don't try, your probability of success is zero. As a consequence of this balance, we tend to make bold forecasts but timid choices. "Germany won't win the World Cup, not in a million years!" "So how much will you offer me?" "Ah, well, it's not in my interests to offer you a bet where I might have to pay out 4 or 5 times what you will offer me in return. For it to be worthwhile to me I'd want a decent 3-figure return at least, but then I risk having to pay you 4." Our cognitive biases work in mysterious ways.

The Utility of Possibility & Certainty

Prior to Daniel Bernoulli's theory of utility, gambles were measured according to their expected value. Recall that this is the average of all possible outcomes, weighted by their individual probabilities. Bernoulli retained this expectation principle, applying it to subjective (psychological)

rather than objective values of outcomes, which gamblers and investors embraced via Kelly staking. This method of evaluation hinges on an underlying principle: the evaluators are rational agents. The work of Kahneman and Tversky into heuristics and biases, however, has revealed that the rationality assumption rests on very shaky ground. In particular, as Kahneman explains, “[t]he expectation principle does not correctly describe how you think about probabilities related to risky prospects.” In his book, *Thinking Fast and Slow*, he invites us to consider how we might feel about having the chances of receiving \$1 million improved by 5%.

- 1) From 0% to 5%
- 2) From 5% to 10%
- 3) From 60% to 65%
- 4) From 95% to 100%

Evidently, the news for each case is not equivalent. Consider the 4th case. Learning that we are certain to become millionaires is fantastic news indeed. The jump from 95% to 100% in terms of attractiveness of the prospect is significant. Consider this from the point of view of a gamble. How much would you be prepared to pay to increase the probability of winning from 95% to 100%, to thereby guarantee a certainty and avoid the small chance of failure? The theory of expected value would suggest \$50,000 to be a fair price, and no more. But suppose that sum was refused, and you were asked for \$100,000. Would you pay? Almost certainly you would. Essentially, outcomes that are nearly (but not quite) certain are given less weight, compared to absolute certainty, than their probability justifies. In such an example, to avoid the fear of disappointment we are risk averse and accept a settlement with negative value expectation. Kahneman and Tversky described this phenomenon as the ‘certainty effect.’ Intriguingly, the certainty effect applies to losses as well, only this time it encourages us to be risk seeking. Imagine, having spent a lifetime saving \$100,000, you were faced with a 95% chance of losing it, or losing \$95,000 for sure. Which would you choose? The sure loss is very unpleasant and faced with such a choice most people would choose to take the risk. After all, if I lose, I lose only \$5,000 more than for the sure loss, so what the hell? In hoping to avoid a large loss we reject a favourable settlement with a positive

expectation and gamble.

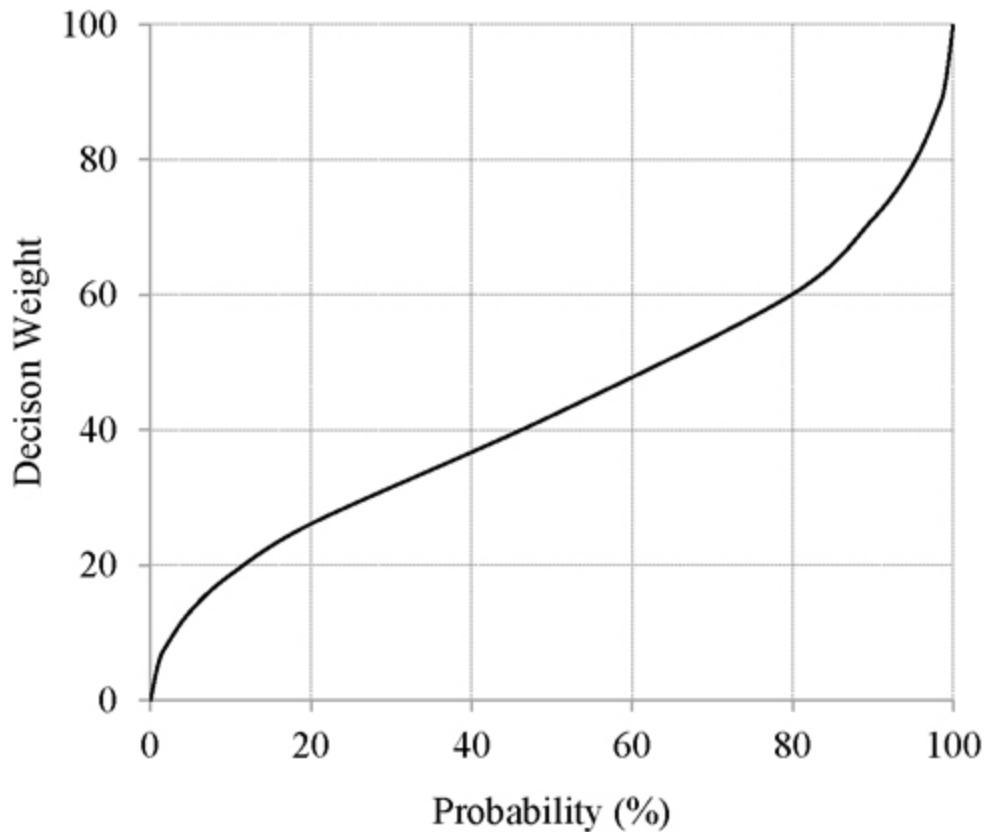
At the other end of the probability scale we encounter the ‘possibility effect.’ The change from a 5% chance to a 10% chance of winning \$1 million involves a doubling of expected value, but evidently not expected utility. In contrast, the same absolute percentage increase from 0% to 5% transforms something that was previously impossible into something that has a possibility. No one will buy a lottery ticket if there is no chance of winning, but introduce a tiny possibility of a jackpot and millions will flock to take part, even under conditions of extreme negative expectation. Completing what Kahneman and Tversky called the ‘fourfold pattern’ is fear of and risk aversion to a small probability of a large loss. This is the domain of insurance. To avoid losing your house in a fire, you accept an unfavourable settlement in the form of premium that has negative expectation. Consider how much insurance premiums you might actually pay over the course of your home ownership lifetime, and compare it to the probability of really losing it in a fire or some similar ruinous disaster. Then consider whether you’d be prepared to take the risk of not paying them.

The overweighting and overestimation of improbable events, in particular, has been the subject of much study since the development of prospect theory. Kahneman accounts for its occurrence as a consequence of our confirmatory bias of memory, the availability heuristic and cognitive ease. System 1 works automatically and unconsciously to retrieve immediate evidence and examples (that easily spring to mind) suggesting the outcome will come true; evidence that invariably will be biased towards pre-existing belief systems. Those examples might even be imagined. System 2, being lazy and uninterested in contemplating the real statistics of the problem, accepts the decision unquestioningly. Try to imagine England winning the World Cup. It’s far easier to imagine possible scenarios why England might manage it than to imagine the scenarios for all the other teams as well. Preferring to avoid the baggage of excessive statistical calculation we arrive, as Kahneman says, at “*a plausible scenario that conforms to the constraints of reality... Your judgment of the probability was ultimately determined by the cognitive ease, or fluency, with which a plausible scenario came to mind.*” Every four years (assuming they have qualified) English sports bettors wager far more money on England ending the decades of hurt than a realistic assessment of their probabilities might

suggest. Simply playing the ‘Three Lions’ song, the official anthem of the England football team for the 1996 European Championships that lamented the lack of success for the team, would conceivably be enough to set punters down the cognitive path of imagining a ‘vivid’ fantasy of an unlikely outcome that ultimately leads to an over-betting of a slim probability. Consequently, English bookmakers usually have shorter odds for England than bookmakers elsewhere do.

The certainty and possibility effects mean that we don’t weight probabilities of outcome exactly according to their numerical values. Improbable outcomes are over-weighted – the possibility effect – whilst near-certain outcomes are underweighted – the certainty effect. Kahneman and Tversky attempted to measure these decision weights in the context of gambles for modest monetary value. The chart below provides an illustration for the data which they collected. Between the extremes of possibility and near-certainty is a demonstrable insensitivity to intermediate probabilities. Between outcome probabilities of about 10% and 90%, people’s range of decision weights is about two-thirds of what would be predicted by rational decision theory, and between 20% and 80% it’s barely half. As the example of betting on England to win the World Cup illustrates, this insensitivity has been exploited by market setters in the domain of sports betting and horse racing.

Possibility & Certainty Effects



The Favourite–Longshot Bias

The over-betting of low probability outcomes in sports and racing, with the corollary that their associated odds are shorter than they should be, has been a well observed and frequently studied phenomenon for many decades. In a typical betting market, the odds quoted by a bookmaker provide a measure of the implied probabilities of all the possible outcomes. They add up to more than 100%, and the amount by which they do, known as the overround, provides a measure of the profit margin the bookmaker typically expects to make. For example, if the ‘true’ odds for a 50-50 proposition are 2.00 for each side, he might offer them at 1.90, implying a probability of 52.6% and an overround of 105.3% (or margin of 5.3%, or 0.053 expressed as a decimal). In such an example, his advantage is spread evenly across both sides of the proposition.

What happens in cases where one side is a heavy favourite? Consider, for

example, a market where the true odds were 1.25 (implying a probability of 80%) and 5.00 (implying a probability of 20%). If the bookmaker spreads a 5% margin equally across these two propositions he'll be offering odds of 1.19 and 4.76 respectively. That, however, is not what happens. On the contrary he'll more usually offer something like 1.21 and 4.44. The favourite has only been shortened by about 3% whilst the underdog or longshot has been shortened by about 12%. This disproportionate application of the bookmaker's profit margin gives rise to what is called the favourite-longshot bias.

Whilst the bookmakers will never tell us exactly how they calculate their odds, I have attempted to model them by means of a simple algorithm which takes into account this disproportionate weighting of their profit margin. Specifically, I have assumed that the weights they apply are in inverse proportion to the probabilities of the outcomes, that is to say, bigger weights for longer odds. Hence, for a book with n runners and overall profit margin M , the differential margin applied to the fair odds for the i^{th} runner (O_i) will be given by:

$$M_i = \frac{MO_i}{n}$$

For example, for a two-player book ($n = 2$) with fair odds $O_1 = 1.25$ and $O_2 = 5.00$, and book margin $M = 0.05$ (5% or in other words an overround of 105%), we find $M_1 = 0.03125$ and $M_2 = 0.125$. To calculate the actual odds we then simply divide each by their appropriate margin plus 1, in this case arriving at 1.212 ($1.25 \div 1.03125$) and 4.444 ($5 \div 1.125$) respectively. We can repeat this type of calculation for any price and any size of margin. We can also use this model to estimate what a bookmaker has estimated the fair odds (O_i) for a book with n runners to be as implied by their published ones, by using the following equation.

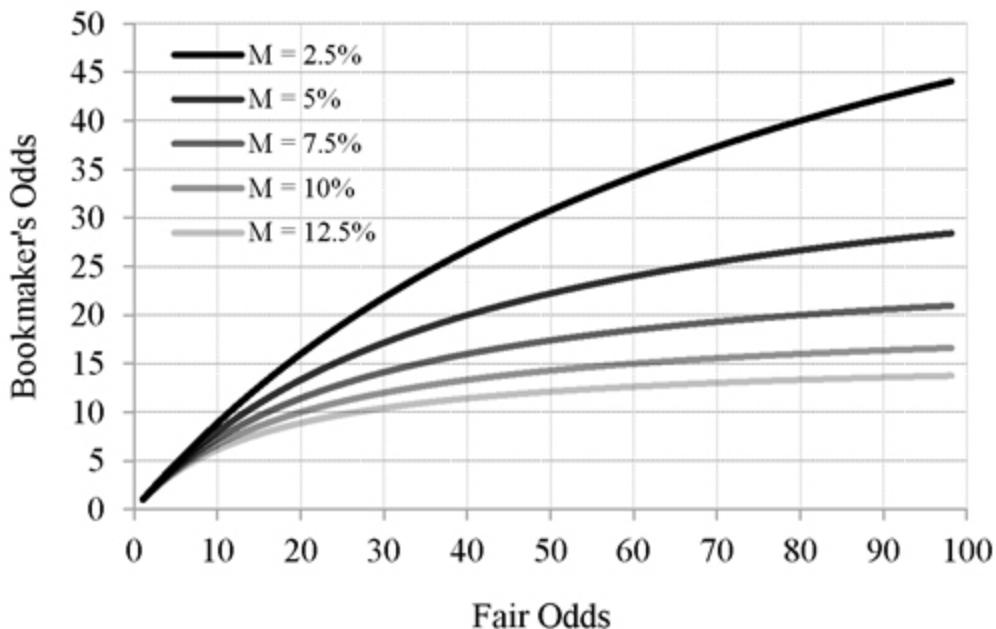
$$O_i = \frac{nO_{\text{bookmaker}}}{n - MO_{\text{bookmaker}}}$$

For example, bet365 posted 1.05 and 11 ($M = 0.043$ ⁴⁶) respectively for

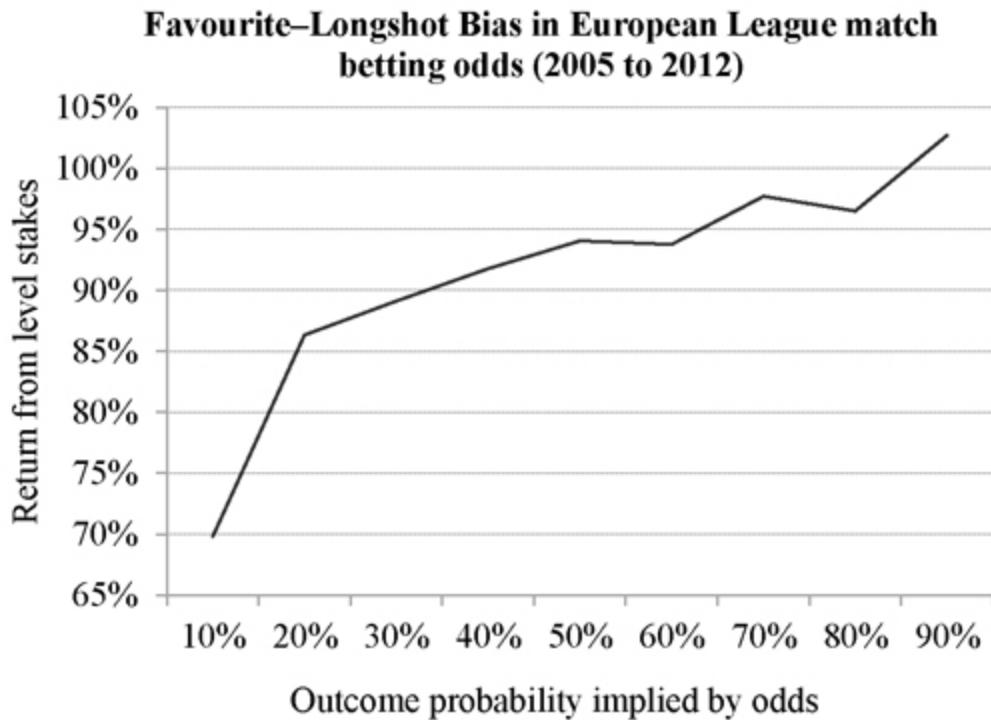
tennis players Andy Murray and Kevin Anderson in their Wimbledon 4th round match played on 30 June 2014. Hence the implied fair odds would be 1.074 and 14.44. Whilst the basis for this simple odds model is just conjecture, it does appear to closely reflect the betting prices for many of the major brands. In this example, the best market prices available for this match were 1.071 (Ladbrokes) and 14.85 (5Dimes), close to the fair odds implied by bet365.

Clearly, as the odds get longer, the differential margin weight that must be applied to them gets larger. Ultimately, there will be a limit to how high those odds will go, determined by the size of the bookmaker's margin. The chart below illustrates the evolution of actual odds that might be offered by bookmakers with varying profit margin for a 2-player book. As the fair odds increase, so do the actual odds offered by a bookmaker but at ever decreasing margins. The theoretical limit to the highest odds a bookmaker will offer according to this model will actually be given by $2/M$. For example, with a margin of 10% or 0.1, the highest odds you would see would be 20; for a margin of 0.05 it would be 40; for a margin of 0.02, 100; and so on. Evidently less generous bookmakers get close to their limits much more quickly. More generically, for a book with n runners, this model would imply maximum odds given by n/M .

Influence of differential margin weights on bookmakers' odds for a 2-player book



What is the consequence of this differential loading of margin by the bookmakers? Evidently, sports bettors should experience poorer percentage returns betting randomly on longshots than they will from favourites. That, in fact, is exactly what happens. Whilst the bookmakers won't provide profits data broken down according to betting price, we can infer what is happening by comparing odds to their actual outcomes. In my previous book *How to Find a Black Cat in a Coal Cellar: the Truth about Sports Tipsters* I reviewed a number of examples where this favourite–longshot bias makes an appearance. The bias in domestic European league football, based on market average match betting odds collected from the 2005/06 to 2011/12 seasons, is illustrated below.



Returns from randomly betting football match odds shorter than about 1.50 (or greater than 66% outcome probability) are close to break-even. By contrast, randomly betting prices longer than 5.00 (or less than 20% outcome probability) can be 80% or lower. Similar biases can be found in horse racing, tennis, basketball, darts and snooker.

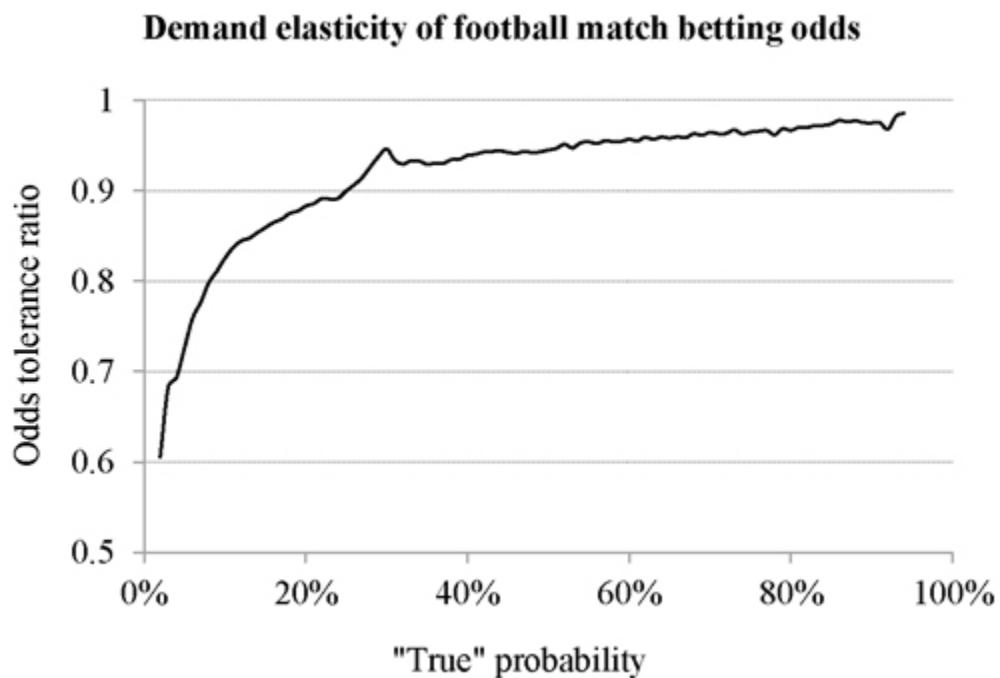
The longer the odds you are betting the poorer the theoretical returns⁴⁷, because the greater the margin that has been applied by the bookmaker. The really interesting question is why does this happen? The favourite–longshot bias has been a source of fascination for academics for a long time. Numerous explanations for its existence have been proposed, including demand-side (risk-seeking) utility preferences of bettors, supply side responses to insider information (more usually associated with longshots) and even an expression of the gambler's fallacy⁴⁸. Probably many of them are working together but the one that perhaps seems most compelling is that offered by prospect theory and specifically the misjudgement or misperception of probabilities as predicted by the possibility and certainty effects and bettors insensitivity to changes in probability. Numerous research papers⁴⁹ have now concluded that bettors express risk-seeking utility towards longshots (the possibility effect) whilst expressing risk

aversion towards favourites (the certainty effect), and that the non-linear probability weights that they apply in formulating these specific utility preferences arise from misperceptions of the probabilities involved.

One significant question remains, however: why do betting exchanges largely fail to show evidence of the favourite–longshot bias, with quoted odds closely reflecting ‘true’ probabilities as inferred from *a posteriori* analysis of results? In attempting to find answers to this question, David McDonald and colleagues from the Centre for Risk Research at the University of Southampton⁵⁰ considered their differences with traditional bookmakers. Whilst bookmakers set the odds, must manage risks and therefore have higher operating costs, at exchanges odds are set by the players themselves, there is no risk management and hence operating costs are minimal. Naturally, the higher operating costs for bookmakers result in less competitive pricing; but why more significantly so for longshots? Indeed, prices offered by exchanges and bookmakers alike (even the least generous bookmakers) are largely the same for the shortest-priced favourites. Given the superior prices for longshots available at the exchanges, explanations for bookmaker favourite–longshot bias that have relied on insider or informed traders knowing more than the bookmaker can surely be discounted. Moreover, in an attempt to manage the liabilities, bookmakers can and do refuse business from ‘smarts’. Instead, McDonald has proposed that a bookmaker’s optimal policy will be to impose a favourite–longshot bias in their prices to reflect biased preferences of bettors as predicted by Kahneman and Tversky’s fourfold pattern.

Essentially, McDonald is suggesting that bookmakers manipulate their odds for the purposes of risk management simply because they can. Since the customers are less sensitive to changes in lower probabilities, or as McDonald says exhibit greater demand elasticity with respect to favourites, bookmakers can exploit these non-rational utility preferences without anyone much caring or noticing. We can actually see how elastic bettors’ price demand is by comparing average and best market betting odds that are available to them. The chart below illustrates how the odds tolerance ratio varies with the ‘true’ probability of outcome for football match betting, based on 50,000 domestic European league games played between July 2005 and March 2012. The best market odds are assumed to provide a

reasonable measure of the true outcome probability⁵¹, whilst the odds tolerance ratio is simply the ‘true’ probability divided by the probability implied by the average betting odds. For example, if maximum and average betting odds were 2.00 and 1.90 respectively, the ‘true’ probability would be 50% and the odds tolerance ratio 0.95.



Quite clearly, as the chance of a result falls, bettors’ demand for a good price falls too. For example, when best market odds are 1.10 (implied probability 91%), the average is only marginally shorter at 1.07 (93%), giving an odds tolerance ratio of 0.97. By contrast, for best odds of 30 (3.33%), the average price is typically about 20 (5%) and the odds tolerance ratio is far lower (0.67). Presumably, prices shorter than the best available are being backed by bettors. If they weren’t, the bookmakers offering them would not be able to sustain them for very long, since a disproportionate amount of action would be focused on the other available options for each particular book. Evidently, the longer the odds, the more a bookmaker can shorten them without his customers knowing or caring. If correct, such a hypothesis would mean that favourite–longshot bias would be largely a supply side phenomenon. Rather than bookmakers passively responding to the demand preferences of their customers, McDonald is arguing that

bookmakers are in fact encouraging their customers to act irrationally simply because they will.

One might even go further and suggest that bookmakers' customers are actually influenced by the value of the betting odds themselves. The tendency to rely too heavily on the first piece of information offered when making subsequent decisions is known as anchoring. The anchoring heuristic was yet another cognitive mechanism first investigated by Kahneman and Tversky. In one of their experiments, participants observed a wheel of fortune numbered 0 to 100 but which was predetermined to stop on either 10 or 65. They were then asked to guess the percentage of the United Nations that was comprised of African countries. Those whose wheel stopped on 10 guessed about 25%; those whose wheel stopped at 65 guessed around 45%. The respondents' judgements had been anchored by a preceding yet completely unrelated piece of information.

This trickery of first perceptions that linger in your mind affecting later perceptions and decisions is used by retailers all the time. When we go shopping, do any of us really have the slightest idea how much things are intrinsically worth? When we see something in the sale, are we sure whether that represents a bargain? Probably not, but the mechanism of anchoring ensures that we can guess at the true answers even though it doesn't feel like we're guessing. Why? Because System 1 is unconsciously, automatically doing all the work, and giving System 2 an answer that it just can't be bothered to think about most of the time. Bookmakers similarly can use the anchoring trick. Knowing that bettors are insensitive to odds and like to back a longshot they just cut them, say for a 10/1 shot down to 5/1. Then, anchored to the shorter price, the bettor might go looking elsewhere for something slightly better. Another brand might be offering 7/1. Great, that looks like good value. 5/1 was probably fair price; I'll take the 7/1 since it will give me some expected value. Of course, probabilistically even the 7/1 is still negative expectation, but with the player anchored to initial price inspection, he won't have the slightest idea about the error that he's committing. Bookmakers truly can take their players for fools, at least for longshots anyway.

If the majority of bettors' misperceptions are resulting in what economists call an inefficient market, where the odds available do not properly reflect the 'true' probabilities of outcomes, surely it should then be

possible for ‘contrarians’ to exploit such a bias for profitable ends. Undoubtedly, there are numerous opportunities with certain brands where the inefficiency is sufficient to offer positive mathematical expectation on favourites, and was something I exploited as part of a tipping service from August 2003 to June 2008 that showed a profit over turnover of 3.04% from 1,294 tips. Sadly, these are the same brands that restrict customers for exploiting these opportunities on a regular basis, either through stake limitation to the amount a customer can wager or an outright ban on them wagering altogether. In my last book, I told the story of a 20-month campaign attempting to exploit the favourite–longshot bias at the bookmaker Sportingbet, a brand that would appear to regularly (and probably intentionally) offer arithmetic value on numerous occasions. My record closed with a +0.9% profit over turnover from 5,751 selections before their stake limitation became so extreme (£1) that it effectively made the account redundant. That I lasted as long as I did was probably on account of my participation in their ‘Last Man Standing’ competition.

Other brands have limited or refused my custom much sooner. Betway, a brand regularly advertising on the UK TV network, cut my stakes to £1.50 after 14 bets. Stan James, another well-known UK bookmaker, lasted a little longer before the account was terminated by the traders, describing my custom as being of an unprofitable nature to them. Intriguingly, Stan James views its business as providing an opportunity for a customer to “*pit their wits against its traders.*” Seemingly, that is not at any price. In this, they appear to be radically different to a brand like Pinnacle Sports with a well-established reputation for passive odds management. Rather than going to war with its customers, Pinnacle Sports simply allows them to fight amongst themselves as to what they think a particular outcome will be, quietly and efficiently adjusting the odds to reflect the money backing those opinions. Stan James, and the other brands like it, by contrast, is in the habit of offering outlier prices containing loss-leading arithmetic value, and holding those prices for longer, even if they attract liquidity. Presumably, it does this to attract new customers and to provide a favourable impression that it offers competitive pricing. Of course, if some of these customers take advantage of such generosity they will be dealt with accordingly. It’s horses for courses really: one brand manages risk purely through the dynamics of the market, the other also through active manipulation and customer

interference. Sometimes, you don't even need to be profitable. If the bookmaker identifies you as someone who regularly beats the closing price, you will be marked as 'sharp.' Passive brands will use such players to help adjust their odds, thereby offering a more efficient market; the others just get rid of their 'winners,' or at least those perceived to be.

In my experience, at least, Stan James was not the worst offender in terms of betting restrictions. Blue Square, a brand since liquidated after its customer database was purchased by Betfair, saw fit to close my account after just one wager. Meridianbet, registered in and operating out of Malta, didn't even afford me that luxury. Having opened a new betting account, I attempted to back San Antonio Spurs over Cleveland Cavaliers in the NBA on 13 February 2013, for a stake of £50 at a price of 1.41 compared to the market average of 1.27. My bet was rejected and amended odds of 1.22 were offered instead. Meanwhile, the published price offered to other customers on the market page was lowered to 1.26. Not only did Meridianbet use my first attempt at a bet to adjust its price, but in the process also chose to offer me less than everyone else. Was I naive to expect anything more? Probably; but if a betting price then becomes nothing more than an advertising gimmick that can't actually be exploited, what we are left with is just a market for lemons⁵².

One brand which appears to engage in excessive artificial manipulation of prices is the Russian-backed sportsbook Marathonbet. Perhaps more than any other brand, with the possible exception of 1XBET, another Russian bookmaker, it seems to pride itself on offering the best market prices⁵³ right the way through to closing time. Indeed, a study of 6,264 tennis match betting prices collected from 15 June 2014 to 8 March 2015 reveals that it offered best market price for any player as much as a third of the time (implying it had best price in as many as two-thirds of matches). This compares to just under a quarter for Pinnacle Sports. Of course, both brands have low margins (about 2% for tennis match betting), so this is to be expected, but it is the manner in which Marathonbet appears to fiddle with its betting odds that is most intriguing. Given bettors' greater demand elasticity with respect to favourites, it seems puzzling that it spends more of its time inflating the odds on the underdogs. Does it know something that Nobel Prize-winning psychologists do not, or am I merely suffering from

halo bias in deference to Daniel Kahneman?

Of the occasions when Marathonbet was top price, 62% of the time this was for the longer-priced player, compared to just 47% of the time for Pinnacle Sports. In absolute terms, it offered roughly twice as many top market prices for underdogs compared to Pinnacle Sports. Whether these would consistently prove to offer profitable expectancy cannot be ascertained from such a relatively small sample size. Betting all 1,277 of these ones would have seen a small loss of -0.2% on total turnover to level stakes. This compares to a -7.7% loss on turnover from all 3,098 of its underdogs. Nevertheless, many of their closing prices lend credence to the suggestion that it's not operating in the same way that Pinnacle Sports would appear to be. The match played between Marsel Ilhan and Jarkko Nieminen at the Istanbul Open on 28 April 2015 is a case in point. Other bookmakers rated both players fairly equally, with average closing prices of 1.84 and 1.91 respectively, and with little price movement in the final 12 hours before match time. Pinnacle Sports' price for Nieminen, for example, fluctuated between 1.93 and 2.05, and ultimately closed at 1.96. Marathonbet and 1XBET, meanwhile, engaged themselves in a little bubble of stupidity, ramping up his price in ever increasing increments from 1.83 to 2.63 with seemingly no method to such madness other than to determine who could offer the best value. The next best price was 2.03. An arbitrage bet including Pinnacle Sports' 1.93 for Ilhan would have yielded a guaranteed return of 111%. This sort of nonsense is not untypical.

See how long you can bet Marathonbet's 'value' opportunities. I tried it and managed three stakes of £50 before a limitation was imposed. After a few further attempts, where most wagers were limited to under £10, the stake limit was further reduced to £1. From reading the popular betting forums, I am evidently not the only one to have experienced this. It's hardly surprising really. Customers who target value like this are assumed to be arbitrage hunters. I suspect most of them, like me, are not. Rather, they are looking for opportunities where one brand is out of step with the others, but to a bookmaker who discourages arbitrage it all looks like the same thing anyway. Almost every bookmaker is intolerant of this type of value hunting, and I've been restricted by most of them for trying it. Needless to say, about the only one that isn't is Pinnacle Sports, and here you'll rarely find value that's theoretically out of step with the market. On the contrary, they

positively encourage arbitrage and value hunting, knowing that, with their laissez-faire odds management and a belief in efficient market principles, it will generally not be the brand on the wrong side of the value line. It's hardly surprising, then, that Pinnacle Sports is one of the most popular online sportsbooks and is recognised for accepting the largest stake limits and seeing the biggest turnovers. In contrast, some of the brands that choose the alternative business model are probably just operating out of broom cupboards, if their levels of internet traffic and customer service are anything to go by. As I've argued previously, interfering with people's behaviour doesn't typically yield the best outcomes.

Practically at least, then, misjudgement of low and high probabilities is probably insufficient to create consistent value expectation for contrarians wishing to exploit such inefficiency. In a sense it's obvious why; the favourite-longshot bias, evident with bookmakers but apparently not exchanges, doesn't really represent a typical inefficiency at all. If it did, favourites would offer consistent opportunities for profit. If that happened, how long do you think they would last? All this is, in fact, is a market manipulation exploiting people's irrationality to help market setters pay the costs of offering the market in the first place. As McDonald has shown, it's much easier to do that with longshots. Perhaps the misperceptions of bettors may not be enough to create consistently predictable and profitable opportunities. Or are they?

Exploiting the Hot Hand Fallacy

One of the most common ways sports bettors exhibit a systematic bias is via the hot hand fallacy, sometimes also called the reverse gambler's fallacy. The error initially arises because the extent to which randomness or luck in a repetitive pattern or streak is present is underestimated. In its place, other causal explanations for such streaks are assumed to be more relevant, causes which should prolong the longevity of the streak. In expressing such a fallacy, the influence of regression to the mean, the tendency for a variable to be closer to the average on a subsequent measurement following a previous extreme one, will be ignored. 'What goes up has a tendency to come down' is substituted with 'what goes up will probably stay up for

longer.'

In a football context, for example, consider a team on a 6-match winning streak. How much of that will be due to causal factors and how much due to luck? If the (fair) odds for winning those games were respectively 1.50, 2.75, 1.72, 3.80, 1.66 and 2.50 then arguably the probability of such a hot streak occurring by chance would be less than 1%; in other words, pretty lucky. When such an extreme streak occurs, however, the probability that it could have happened will, by many, be overestimated, in much the same way that many overestimate the chances of low probability outcomes as described by the possibility effect. To put it another way, because we have assumed that the probability such an outcome arising as a result of luck is so small, given that it happened, luck cannot have had much to do with it. Arguably, such a thought process would increase the likelihood of backing the team in their 7th match. Hence, bettors expressing the hot hand fallacy and thereby ignoring regression to the mean will show a tendency to over-bet teams with better recent form than their opposition. As a consequence, we might expect the odds for such 'hot' teams to be shorter than a more objective assessment of outcome probabilities would suggest they ought to be.⁵⁴ Similarly, odds for 'cold' teams might be expected to be priced longer. That, at least, is the theory. Is there any way this can be tested in practice? Perhaps more importantly, if any systematic expression of the hot hand fallacy exists, will such a bias provide consistently profitable opportunities for contrarians looking to exploit it?

To test such a hypothesis we need some way of measuring how 'hot' or 'cold' teams are. One way is to use the betting odds themselves. If we assume that, on average, the 'fair' odds represent the 'true' probability of a team winning then over the long term our expected return betting such odds should be approximately zero. Consequently, teams on 'hot' streaks will show positive returns over the short term, whilst those on 'cold' streaks will show negative returns. We can estimate the fair odds from actual bookmaker odds using the model described earlier in the chapter.⁵⁵ If our hypothesis is correct, backing 'cold' teams should prove to be relatively more profitable (or at least less unprofitable) than backing 'hot' teams, by virtue of the fact that disproportionately fewer people are backing them. The following analysis would appear to offer considerable support to the

presence of such a systematic bias in a football match betting market.

For the 5 domestic seasons 2010/11 through 2014/15, average match betting odds (home-draw-away) were used to estimate ‘fair’ odds for 36,126 league matches in 22 European divisions. For each match, the winning team was awarded a risk adjusted score of $[1 - 1/\text{odds}]^{56}$, whilst the losing team (or both teams where they drew) was awarded a score of $[-1/\text{odds}]$. For each team, these scores are consecutively added during the course of a season, and reset to zero at the start of the next one. The process is perhaps best illustrated by means of an example, in this case for Blackpool’s first 10 league games in the 2010/11 season, as shown in the table below. We can see that, during this period, Blackpool had over performed relative to what the betting market had expected the team to achieve. After 10 games, betting risk-adjusted stakes at these theoretical fair odds would have netted the player over 2 units.

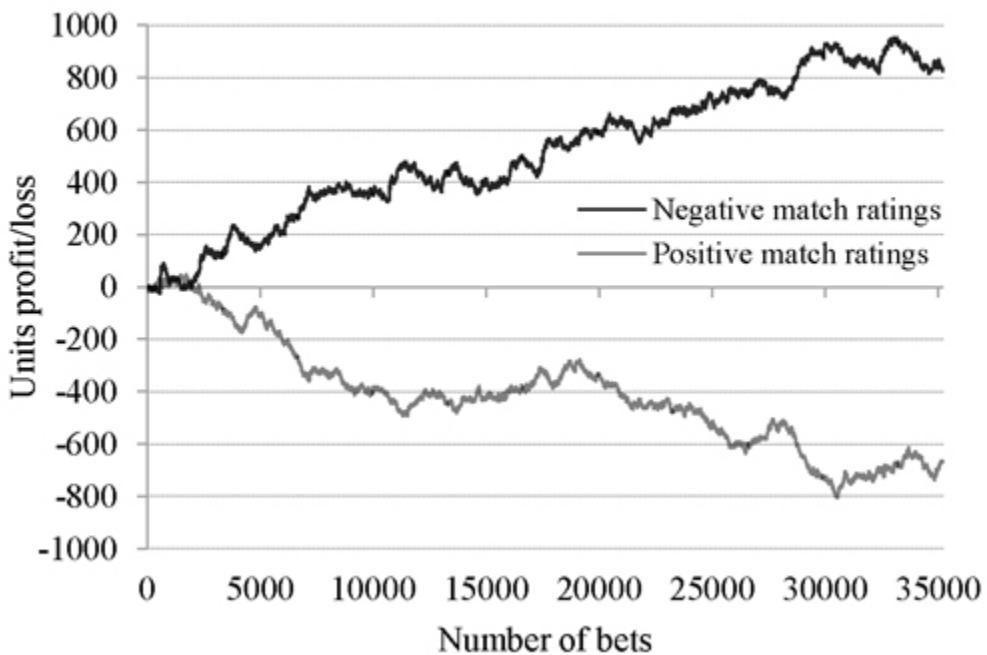
Team	Opposition	Date	Fair Odds	Result	Profit/Loss	Cumulative Score
Blackpool	Wigan	14/08/10	4.96	Won	0.798	0.798
Blackpool	Arsenal	21/08/10	27.40	Lost	-0.037	0.762
Blackpool	Fulham	28/08/10	3.29	Lost	-0.304	0.458
Blackpool	Newcastle	11/09/10	6.92	Won	0.855	1.313
Blackpool	Chelsea	19/09/10	48.20	Lost	-0.021	1.292
Blackpool	Blackburn	25/09/10	3.34	Lost	-0.299	0.993
Blackpool	Liverpool	03/10/10	12.80	Won	0.922	1.915
Blackpool	Man City	17/10/10	6.39	Lost	-0.157	1.758
Blackpool	Birmingham	23/10/10	4.72	Lost	-0.212	1.546
Blackpool	WBA	01/11/10	2.95	Won	0.662	2.208

The next step is to utilise these cumulative scores to design a predictive rating. After Blackpool’s first game, for example, their cumulative score was 0.798, on account of winning their match against Wigan. This score is therefore taken as their team rating for their next game. In other words, it is a measure of how ‘hot’ or ‘cold’ they are going into their next game. Naturally, prior to their first game of the season, their rating will be 0.

Repeating this process for every team, we can then finally produce a match rating, simply defined by subtracting the rating of one team away from the rating of their opposition. Swapping the teams around will simply provide a rating of opposite sign with equal magnitude. So, for example, Blackpool entered their game with Manchester City on 17 October 2010 with a team rating of 1.915. Manchester City, similarly, entered the game with a rating of 0.521. Hence, we can calculate the match rating by $1.915 - 0.521 = 1.394$ (or -1.394 if calculated the other way around). In other words, this is equivalent to saying that prior to their match, Blackpool had been performing relatively better than expected compared to Manchester City. Of course, Manchester City, with a positive team rating themselves, had also been over performing, but just not to the extent that Blackpool had been. In contrast, when Blackpool met Birmingham in their next game, Birmingham had a team rating of -1.509 , implying they had been doing worse than the betting market had predicted. The match rating for that game was 3.267 in favour of Blackpool.

Had someone been able to bet every home and away result for each of the 36,126 matches in this sample at the theoretical ‘fair’ odds and to level stakes (a total of 72,252 bets), their profit over turnover would have been 0.22%. The fact that it wasn’t exactly zero will be a consequence of either model inaccuracy in the way the ‘fair’ odds have been calculated, slight over performance of longer odds relative to shorter ones during this 5-season period, or a combination of the two. Nevertheless, the figure is reasonably close to what we could expect betting at ‘fair’ odds to yield. The time series of accumulated profits/losses, furthermore, shows a fairly typical random walk about the break-even line. Contrast that to the time series for betting, on the one hand, all negative match ratings (where we favour a ‘colder’ team over a ‘hotter’ team), and on the other hand, all positive ratings (where we favour a ‘hotter’ team over a ‘colder’ team). Again, bets are struck at theoretical ‘fair’ odds and to level stakes. The results are graphed below.

Theoretical bankrolls from level staking



The conclusions appear stark. Backing teams (at theoretical ‘fair’ odds and to level stakes) that had been performing relatively ‘colder’ to their opposition in this sample would have delivered a 2.36% yield. What is more, the profit-taking appears consistent across the 5 seasons. In contrast, backing relatively ‘hotter’ teams would have shown losses of -1.90% on turnover. The difference between these two yields for this sample is statistically significant.⁵⁷ The implication must be that relatively ‘colder’ teams are overpriced as a consequence of fewer people backing them and, as such, provides evidence of a systematic bias arising out of the hot hand fallacy and a disregard for regression to the mean. Being a contrarian in football match betting might very well be profitable.

We should not, however, let excitement get the better of us just yet. Firstly, statistical significance compared to 0% expectation is weaker.⁵⁸ Furthermore, to a sizeable extent, it arises out of the large sample size rather than the magnitude of the advantage the bias gives us. The profitable yield from backing relatively ‘colder’ teams, after all, is small. A further consequence of that will ensure that short term variance in the swings of profits and losses will be considerable. Finally, we need to remember that the odds used to produce this performance were theoretical. How would

things look in the real world? In fact, blindly backing best available market prices to level stakes for this sample would have seen a -0.51% yield. Backing relatively ‘colder’ and ‘hotter’ teams would have seen yields of 1.41% and -2.39% respectively. The lower figures are a consequence of the fact that, on average, even best prices are not quite ‘fair’. In this sample, the mean overround was 100.6%. Furthermore, for reasons already discussed, many of those best prices will be made available by bookmakers known for limiting or discontinuing the activity of customers looking to consistently exploit them.

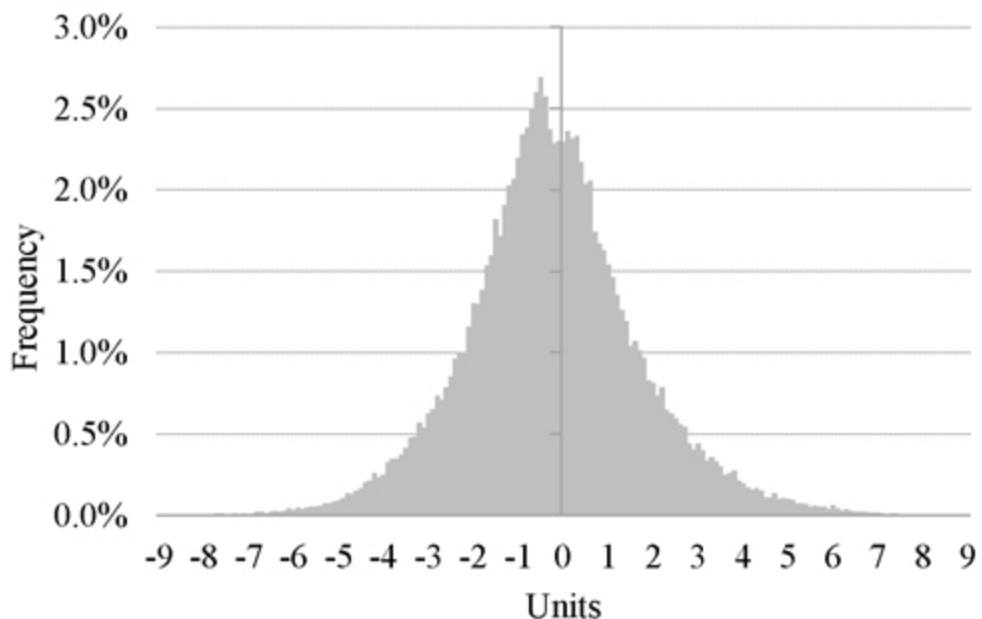
We could, of course, choose to be more selective in our betting criteria. What if our expression of betting interest was limited to match ratings that were not simply negative but below -1, -2, -3 and so on? The effects of such selective betting can be seen in the next table.

Match ratings less than	Bets	Yield from ‘fair’ odds	Yield from best market odds
0	35,200	2.36%	1.41%
-1	23,215	4.44%	3.46%
-2	14,230	4.87%	3.77%
-3	8,355	5.27%	4.17%
-4	4,742	2.85%	1.72%
-5	2,564	-0.17%	-1.25%
-6	1,315	-7.01%	-8.19%
-7	672	-7.90%	-8.94%
-8	316	-3.58%	-4.86%
-9	159	-7.48%	-9.08%
-10	61	18.56%	16.21%

As one might reasonably expect, restricting our betting to ever decreasing match ratings improves returns, but apparently only to a point. Below match ratings of -4 the improvement ceases, and indeed results in unprofitable returns below match ratings of -5. Of course, it should be observed that as match rating selectivity increases, the number of betting opportunities vastly diminishes. The apparent failure of our hypothesis to explain returns

at these low match ratings could simply be a consequence of greater variance and bad luck due to the much smaller sample sizes. Alternatively, however, it could be a result of some other factors not taken into account, where the relative difference between the ‘heat’ of two teams is large. Finally, the whole association found here could simply be the result of good fortune, arising from the mining of a large data set until something of value was discovered, but which has no basis in causality whatsoever. Lest you need reminding, humans have an almost never-ending capacity to be fooled by randomness and for demanding causal explanations where none exists. In this context, it’s worth taking a look at the distribution of team scores produced by this model for contrarian betting (below). Whilst these scores should contain the signal of the systematic bias we have been looking for, their distribution is distinctly normal. Remember, normal distributions imply random processes. Whatever bias is present in the data is hidden amongst the noise. Explicitly, we probably knew that already, since it only amounts to about 2% compared to theoretical ‘fair’ odds and lower still in the world of real betting odds. The exploitation of such a small advantage, even if valid and consistent, will require discipline to maintain, where wild ups and downs of profitability are typical, and where bookmakers’ restrictions provide further practical constraints.

Distribution of team scores



Stock Market Irrationality

Misperceptions of probability, overconfidence, confirmation, anchoring, loss aversion and other sources of irrationality aren't exclusive to betting markets. All are present in the greatest gambling show on Earth: the stock market. For a century and more, total return on equities has outpaced GDP growth over long periods in many countries. Such an apparent paradox has caused some commentators, mostly notably Bill Gross⁵⁹, American financial manager and co-founder of Pacific Investment Management (PIMCO), to question its sustainability. In the US, for example, inflation-adjusted returns from stocks since 1900 have averaged about 6.6% per year (known as the Siegel constant and named after Jeremy Siegel, the author of *Stocks for the Long Run*). Over the same period, annual wealth or GDP growth has only averaged 3.5%. Given such a discrepancy, how is it possible for stockholders to be skimming off an additional 3%?

GDP and the stock market index measure different things. The former measures annual output of goods and services, the latter the value of all future corporate earnings, that is to say, cash flow. Frequently over much shorter time scales the two measures don't even correlate particularly well. Over the longer term, however, they should, in theory at least, be intrinsically related. Think of the economy in terms of a river: economic growth determines its size whilst the amount of water flowing through it provides a measure of cash flow. Stock prices can grow faster than GDP for two underlying reasons. Firstly, the ratio of stock prices to company earnings (known as the P/E ratio) can rise. At the end of the 19th century, the S&P 500⁶⁰ P/E ratio was about 10; today it is double that. Essentially as the P/E ratio increases, the value of company rises for every unit of profit that it produces. Facebook's current P/E ratio is about 100; the higher the P/E ratio, the greater the likelihood that a stock is overvalued. Some internet stocks during the Dotcom bubble of 1997 to 2000 had essentially infinite P/E ratios since their companies weren't earning anything. Secondly, corporate earnings can rise faster than GDP where the share of national income going to capital increases relative to that going towards labour. In the US, for example, wages and salaries as a percentage of GDP have declined from above 50% in the 1960s and 1970s to 42% by 2013.

Neither, for obvious reasons, can go on indefinitely. If stocks continue to appreciate at a rate 3% higher than the underlying economy itself, stockholders would ultimately command, not just a disproportionate share of wealth, but nearly all the money in the world. Of course, the flip side to this is that price to earnings and earnings to GDP ratios can fall too, with the potential for financial markets to underperform relative to the wider economy. As a consequence, the year-to-year variability in the market tends to be much greater.

For Bill Gross at least, “[t]he Siegel constant of 6.6% real appreciation, therefore, is an historical freak, a mutation likely never to be seen again.” He is probably correct on the first suggestion: GDP-plus returns from equities are evidently not sustainable over the long run. Whether we’ll never see the like of it again is another matter. Financial bubbles come in all shapes and sizes, and the detrimental effects encountered when one bursts don’t usually preclude another one following it. People’s memories are too short and their overconfidence too enduring for that. Some are short-lived like the South Sea bubble of the 18th century or the Dotcom bubble of the last years of the 20th century; seemingly others can last decades or even longer. What they have in common, however, is the capacity to fool people that they are even in one. Soothsayers like Vince Cable, the UK government’s Secretary of State for Business from 2010 to 2015, may argue after the event that they saw them coming. In truth, almost all are just overconfident victims of hindsight bias.

Irrational optimism, herd behaviour and a disposition for looking backwards to explain the future (remember, we are built to recognise patterns to explain causality) are all psychological characteristics that will account for why today’s investors believe that the persistence of success so manifest in the last century will be inevitable in this one. Don’t try to tell them they might be wrong, that Siegel-size returns might just have been lucky, or worse still, manufactured out of Ponzi-style investment thinking⁶¹, where all you need is a greater fool willing to pay more for an asset than the last one, regardless of how much it might intrinsically be worth. “*What do you think I am; a nobody?*” When the time frame of a bubble extends beyond lifetimes, it’s understandably natural to deny its existence. In contrast to the overweighting of probabilities when evaluating choices from

description, we frequently do the opposite when considering choices from experience. “It’s never happened to me,” or so the thinking goes. For example, people frequently underestimate the chances of getting cancer, even though at least 1 in 3 will eventually do so in one form or another. Similarly, fund managers prior to the 2008 financial crash couldn’t anticipate such an event, since they had no idea what one felt like.

Daniel Kahneman and his colleague Mark Riepe have made an interesting exposition of investor overconfidence⁶². They ask the following question:

“What is your best estimate of the value of the Dow Jones one month from today? Next pick a high value, such that you are 99% sure (but not absolutely sure) that the Dow Jones a month from today will be lower than that value. Now pick a low value, such that you are 99% sure (but no more) that the Dow Jones a month from today will be higher than that value.”

In other words, an investor should be 98% confident that the market index will lie within a specified range. Correct predictions would lead to actual outcomes falling outside the predicted range only 2% of the time. In fact, in most studies asking similar questions, the actual failure rate is closer to 20%. Patently people are not well calibrated to estimate probabilities.

According to Terrance Odean⁶³ from the University of California, such overconfidence leads to overtrading which results in poorer overall investment performance. Men in particular are notoriously overconfident in this respect, a characteristic which undoubtedly is evolutionary in origin, given their propensity for greater risk taking. Using brokerage account data for over 35,000 households covering the period 1991 to 1997, Odean revealed that, whilst men trade 45% more than women, such over-activity reduces their net returns by 2.65% per year as compared to 1.72% for women⁶⁴. In a separate investigation analysing the trading records of 10,000 brokerage accounts from 1987 to 1993, Odean also demonstrated that investors sell far more winners than losers⁶⁵. Moreover, the winners they sold outperformed the losers they clung on to by 3.4% in the following year. Counterintuitively, prospect theory explains this via loss aversion. Rather than sell poorly performing stocks, investors hang on to them in the hope that they will turn around. Furthermore, denying the error of picking the wrong stock contributes to ego defence. In this way, our fear of losses

actually leads us to experience more losses. In the extreme, we even end up chasing them, in an attempt to correct past errors with future success. Barings Bank learnt that to its cost when Nick Leeson, one of its traders, lost £827 million chasing past mistakes with Martingale-style trading gambles that ultimately led to the bank's collapse.

With regards buying, Odean's investors were also seen to purchase stocks that had already risen by 26% in the previous 2 years, only to find that they declined by 3% in the 12 months thereafter. Such behaviour of buying high and selling low (assuming they've chosen to abandon their underperforming stocks at all) is classically known as the dumb money effect. At its heart is simply the inability to recognise regression to the mean. Rising stocks are typically perceived to be winners with causal explanations for why their values are increasing. In random markets, however, much of what happens is just noise, not signal. Things that go up, on average, have a tendency to come back down. If you bought after most of the gains had already taken place, regression to the mean will inevitably mean that your purchases will perform less well than you might have expected. Furthermore, from the perspective of market efficiency, once the stocks are rising, the positive expectation they might have held has probably already disappeared. By the time news is available to most investors that something is worth buying, it's probably already too late. Indeed, Odean speculated that some of the investors he studied were among the last buyers to contribute to the rise of overvalued securities. More crucially, he concluded that, whilst investors do have useful information at their disposal to make investment decisions, they are misinterpreting it just as prospect theory predicts. To be sure, financial investors really do make bold forecasts but timid choices.

If most investors are buying high and selling low, what would happen if you purposely did the opposite? Werner De Bondt and Richard Thaler, professors in behavioural finance, sought to test this hypothesis by comparing investment performance between so-called 'winner' and 'loser' portfolios on the New York Stock Exchange over the period 1926 to 1982⁶⁶. 'Winner' and 'loser' portfolios were constructed from the best and worst 35 performing stocks based on the preceding 3 years. Their performance was then tracked over the subsequent 3 years. 'Losers' outperformed the market by 19.6%; 'winners', meanwhile, fell short of the benchmark by 5%. Such

findings seem remarkably similar to the apparent systematic bias arising in football match betting as a result of over-betting ‘hot’ teams whilst under-betting ‘cold’ ones. According to De Bondt and Thaler, investors, expressing both availability bias – focusing excessively on immediately available information – and recency bias – extrapolating recent events into the future – were overreacting to previous performance, avoiding stocks that had experienced losses whilst targeting those that had seen gains. Prices for losing stocks were initially driven down disproportionately before investors acknowledged that their earlier pessimism had not been entirely justified; the losers subsequently rebounded as investors came to the conclusion that the stocks were undervalued. Indeed, most of the rebound was found to occur in the second and third years. In contrast, previously well-performing stocks were eventually shunned as investors came to the conclusion that their earlier exuberance hadn’t been totally justified.

If only stock picking were so easy. An investment strategy that seeks to exploit the lack of full immediate adjustment to stock valuations is an example of contrarian investing. More generally, a contrarian investor (like the contrarian football bettor we considered a little earlier) is one who attempts to profit by investing in a manner that differs from the conventional wisdom, when the consensus opinion appears to be wrong, as often behavioural economists seem to think it is. But if investors can make so much profit simply by backing previous losers, why isn’t everyone doing that? Moreover, what do you think would happen if everyone did? If a contrarian’s response to ‘winners’ and ‘losers’ (or ‘hot’ and ‘cold’ football teams for that matter) was inevitably and predictably successful, wouldn’t the market then evolve differently as professional investment managers developed strategies in an attempt to exploit the pattern, thereby threatening its self-destruction? According to Burton Malkiel, author of the definitive *Random Walk Down Wall Street*, the “*more potentially profitable a discoverable pattern is, the less likely it is to survive.*”

Hindsight, of course, is a wonderful thing; indeed it is consistently the best gambling system ever invented. Sadly, its profitability is only available retrospectively. Hindsight bias ensures that events that even the best-informed experts did not anticipate often appear almost inevitable after they occur. Why? Malkiel argues that such errors are sustained by having a “*selective memory of success,*” in part because investors attribute profitable

decision making to their own abilities and rationalise negative outcomes as resulting from external and uncontrollable events, including bad luck. Ask any gambler how much profit he has made over his lifetime; almost certainly you'll be given an incorrect answer, weighted heavily in favour of winning. We always remember our successful investments, bets and gambles, not so much our failures. This is not a purposeful denial of the facts, merely an innocent self-deception. Hindsight promotes overconfidence, fostering the illusion that markets are far more predictable than they really are.

Malkiel offers several other explanations for why contrarian investing in response to psychologically driven inefficiencies is probably not the golden goose that lays the golden egg. Many of these so-called predictable patterns may simply be the result of data mining. Modern computing power ensures that financial data analysis is a fairly effortless process. Look long and hard enough for associations between financial variables and investigators can pretty much find anything they want to. Sports bettors who like playing with big data, too, are notoriously guilty of making spurious correlations that have no basis in causal reality. Indeed, I had considered that possibility with the contrarian football betting model developed earlier. Remember, however, we are designed to find patterns even where none exists. As Nassim Taleb perennially reminds us, investors (and other gamblers) are endlessly fooled by randomness. Secondly, the literature on contrarian investing may very well be biased in favour of reporting positive findings. How probable is it that boring confirmation of meaningless randomness will be published in professional journals? Generally, we only ever get to read about the interesting patterns (an example of survivorship bias) because they make for much better stories. Finally, profiting from overreaction to recent form in the stock market may not be possible if the reversals represent nothing more than regression to the mean. If luck accounts for most, if not all, of the profits of investors following such momentum strategies, the law of large numbers will surely imply that predictable and consistent profitability will be beyond their reach.

Aside from these proposed weaknesses in contrarian investing, it's actually very difficult to follow the strategy. One common explanation for why is because of something called 'groupthink'. This psychological phenomenon arises out of a natural herding instinct, in which the desire for

harmony or conformity within a group of people results in irrational or dysfunctional decision making behaviour. Much of the time individuals express diversity of decision making, acting independently and without centralised control. Sometimes, however, a herd instinct arises where individuals may choose to follow the decisions of others simply to acquire what is termed ‘social proof,’ the tendency to assume that if others are doing something it must be a good idea. Psychologists Stanley Milgram, Leonard Bickman, and Lawrence Berkowitz had a great deal of fun in 1968 experimenting with social proof⁶⁷. First, they placed a single person on a street corner and had him look up at the sky. 4% of pedestrians stopped to see what he was doing, whilst 42% took a glance skyward to see what all the fuss was about. When they increased the number of accessories to 15, 40% of people stopped and 86% looked up.

In 1951, the psychologist Solomon Asch revealed the extraordinary extent to which individuals will seek to conform⁶⁸. In his experiment, a short video of which is available on YouTube⁶⁹, male college students were shown two cards: the first with a single line drawn on it; the second with three lines drawn, only one of which was the same length as the one on the first card. The participants were then asked to say aloud in front of the whole group which line they thought was the same length. Unbeknown to one of the participants in the group, who would always go last, the remainder were accomplices in the experiment who had been instructed to give correct or incorrect answers. In total there were 18 trials, the first two for which the accomplices were always instructed to give the correct answers to help the real subject feel at ease. Of the remaining 16 trials, the accomplices were instructed to make incorrect responses 12 times. The results were truly astounding. On average nearly a third (32%) of subjects chose to conform with the majority view when the accomplices gave incorrect responses. Of the 50 subjects who took part, three-quarters of them chose to conform with the wrong answer at least once. In a control experiment without any pressure for group conformity, the error rate was less than 1%.

Perhaps even more astonishingly, neuroscientists investigating which parts of the brain are involved in group conformity have confirmed that the craving for social proof arises not because people choose to lie but because

actual perceptions are shifted⁷⁰. In other words, conformity within a group actually changes what people believe they experience. In *The Real Story of Risk*, Glenn Croston reckons that our compulsion to conform stems from the “*thing we fear more than death*” – ostracism. For most of humanity’s history, being part of a group meant survival, helping each other to find food and to defend against predators. In such an environment, rejection by the group could result in death itself. Whilst the world for most of us is now a much safer place, herding and the tribalism that necessarily arises from it persist across time and cultures; just think of football fans, Facebook followers, high-school cliques and inner city gangs, all seeking solace within the confines of a group. The risk of not fitting in weighs heavily on people’s minds; the fear of looking foolish, of experiencing shame and losing face, causes anxiety that is almost unparalleled, particularly for oriental cultures schooled in the philosophy of Confucius. Those who’ve had to stand up and talk in front of a crowd will know exactly what Croston is talking about. Such self-effacement will undeniably be evolutionary in origin.

In an investment setting, the fear of not conforming and looking foolish may be accompanied by the fear of regret, of missing out on a profitable opportunity. If people are buying stock it must mean it is stock worth buying. Failure to do so could result in what Daniel Kahneman calls ‘regret of omission’, failing to do something that might make us financially better off. Of course, in choosing to avoid omission regret we might make a bad decision; we might purchase the stock and find that its value falls. In this case, we will then experience the stronger ‘regret of commission’, regretting something we did. Commission regret is more powerful than omission regret because of loss aversion: losses hurt more than gains. In this case, actual losses will hurt considerably more than missed opportunities to gain.

For James Surowiecki, author of *The Wisdom of Crowds*, overreaction as a consequence of groupthink is the textbook explanation for financial bubbles and crashes. When people stop thinking independently, when there is too much single-mindedness, when there is only limited diversity of opinion, when everyone starts piggybacking on the wisdom of the group, an ‘information cascade’ develops with people abandoning their own

information, opinions and beliefs in favour of inferences based on the behaviour of others who have acted before them. Independent thinking, Surowiecki explains, is the sort of thing you do when you try to figure out how much something is worth, and consequently whether it's worth buying, without worrying too much about what other people think. By contrast, the price of a stock (and perhaps the odds of a football team or tennis player) sometimes reflects a series of dependent decisions, because their evaluation depends partly on what everyone else is thinking. The average investor is concerned not just about what he thinks, but what he thinks other average investors are thinking too, and so on and so on. Needless to say, once everyone starts thinking the same way as everyone else irrationality can escalate dramatically, even for the price of tulips. If the herd is buying, prices rise, sometimes very significantly, creating what financial commentators call a bubble. When the herd is selling, usually after something precipitates a bubble to burst, prices crash, typically much faster than when the bubble was growing. Given our predisposition for loss aversion, this asymmetry is unsurprising.

Even the professionals suffer from groupthink. Much of the explanation for the 2008 global financial crash that began with the US subprime mortgage crisis can be found in the herding of credit ratings agencies. With low interest rates in the US in the early years of the new millennium, real estate was regarded as a good bet. As more American citizens became homeowners, prices began to inflate. They continued to spend freely, believing the rising value of their property, which doubled between 2000 and 2006, would cover the debts. These debts were financed by mortgage-backed securities and collateralised debt obligations, which started to unravel in 2006 as it became apparent that many were invested in poor credit quality – the so-called subprime borrowers. In fact, subprime mortgages had grown from \$35 billion in 1994 to a staggering \$1.3 trillion by 2007. Ultimately much of it defaulted or lost most of its value. Perhaps none of it would have happened had the credit ratings agencies not graded the debt securities as triple-A, the safest of all ratings with the lowest probability of default. Such a rating implied roughly a 1-in-1000 possibility that it would fail to pay over the next 5 years. In the event, nearly a third of collateralised debt obligations defaulted. Talk about overconfidence! Many have since argued (with hindsight?) that the credit ratings agencies were

part of an oligopoly, a market or industry dominated by a few significant players. In the US, these were the ‘Big Three’ rating agencies: Moody’s Investors Service, Standard & Poor’s and Fitch Ratings. Such a setup was arguably a perfect breeding ground for a lack of independent thinking, particularly since all of them were benefiting financially from issuing such ratings.

Financial markets, then, are not particularly rational places; David Bernoulli would be appalled. On the contrary, they are driven by the sentiments of optimism and pessimism, greed and fear. In *Mean Markets and Lizard Brains*, Terry Burnham cautions that our emotional, intuitive automatic System 1 “*makes us greedy when we ought to be fearful and fearful when we ought to be greedy.*” Optimism drives overbuying leading to overvalued stocks; pessimism drives overselling leading to undervalued stocks. It’s a kind of pricing bias but on a grand scale. Any profitable opportunities that arise, therefore, do so not through a better prediction of the future but through betting against the irrationality. Speculating on the financial markets, and the business of sports betting as well, for that matter, are really just like a game of poker, where decision making is dependent not just on what the decision maker thinks but also on what all the other decision makers think. Indeed the renowned 20th century economist John Maynard Keynes defined speculation as “*the activity of forecasting the psychology of the market.*” Game theory made great strides in recognising all of this, but it took prospect theory to reveal that players in games of uncertainty are not always acting rationally. Yet if so many people persist in biased thinking, why aren’t there very many others either willing to exploit it or profiting from doing so? Or, as Peter Bernstein, in *Against the Gods*, jokes: “*if people are so dumb, how come more of us smart people don’t get rich?*”

³⁵ $[(1/6) + (2/6) + (3/6) + (4/6) + (5/6) + (6/6)] = 3.5$

³⁶ The paradox takes its name from the Russian city, St. Petersburg, where Daniel Bernoulli was resident. He published his theory in the *Commentaries of the Imperial Academy of Science of Saint Petersburg*. Originally published in Latin it was not translated into English until 1954, when it assumed the title *Exposition of a New Theory on the Measurement of Risk*, published in *Econometrica*, 22(1), pp.23-36.

- 37 Kelly, J.L.Jr., 1956. A new interpretation of information rate. *Bell System Technical Journal*, **35(4)**, pp.917-926.
- 38 Kahneman, D. & Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, **47(2)**, pp.263-291.
- 39 Chen, K., Lakshminarayanan, V. & Santos, L., 2006. How Basic Are Behavioural Biases? Evidence from Capuchin Monkey Trading Behaviour. *Journal of Political Economy*, **114(3)**, pp.517-537.
- 40 Pope, D.G. & Schweitzer, M.E., 2011. Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes. *American Economic Review*, **101(1)**, pp.129-57.
- 41 Tom S.M., Fox C.R., Trepel C. & Poldrack R.A., 2007. The neural basis of loss aversion in decision making under risk. *Science*, **315(5811)**, pp.515-518.
- 42 <https://www.youtube.com/watch?v=aoScYO2osb0>
- 43 <https://www.youtube.com/watch?v=76p64j3H1Ng>
- 44 John is a man who wears gothic inspired clothing, has long black hair, and listens to death metal. How likely is it that he is a Christian and how likely is it that he is a Satanist? Most people underestimate the probability of him being a Christian, and overestimate the probability of him being a Satanist. This is because they would ignore that the base rate probability of being a Christian is vastly greater than that of being a Satanist.
- 45 Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which is more probable? Linda is a bank teller or Linda is a bank teller and is active in the feminist movement. Most people choose the second option. However, the probability of two events occurring together (in conjunction) is always less than or equal to the probability of either one occurring alone. ‘What you see is all there is’ ensures System 1 short cuts to the second option by means of substituting a difficult question with a simpler one. Daniel Kahneman describes the ‘Linda problem’ as his best-known and most controversial experiment.
- 46 The book margin, M, can be calculated directly from the published odds by summing the inverse of the odds for every runner in a book. In this example $M = \{(1/1.05) + (1/11)\} - 1$.
- 47 ‘Theoretical,’ in this context, means betting blindly or randomly. Whilst most bettors claim to be using ‘systems’ to identify good value in the odds, their results demonstrate that barely any of them are actually achieving that. The likelihood is, then, that if you choose to bet long prices you can expect to do worse, in percentage terms, than those betting favourites.
- 48 Players, teams and runners who have lost many times in succession tend to be longshots, and thus a gambler’s fallacious belief that they are due for a win may contribute to their over-betting.
- 49 See for example Jullien, B & Salanie, B., 2000. Estimating Preferences under Risk: The Case of Racetrack Bettors. *Journal of Political Economy*, **108(3)**, pp.503-530 and Snowberg, E. & Wolfers,

J., 2010. Explaining the Favourite-Long Shot Bias: Is it Risk-Love or Misperceptions? *Journal of Political Economy*, **118**(4), pp.723-746.

[50](#) McDonald, D. C. J, Sung, M-C., Johnson J. E. V, & Tai C. Forecasting the presence of favourite-longshot bias in alternative betting markets. Conference presentation at 33rd International Symposium on Forecasting, Seoul, 23-26 June 2013. David McDonald kindly provided the original paper via personal correspondence: *Toward an explanation of the favourite-longshot bias in competing betting markets*.

[51](#) In fact this assumption is probably not quite valid since even at best market prices a very weak favourite–longshot bias was observed, although its significance has not been tested. Betting all prices longer than 4/1 for example would have lost over 4% on turnover. Nevertheless, for the purposes of this exercise it probably represents a fair approximation.

[52](#) The phrase comes from the economist George Akerlof, in his discussion on market information asymmetry, where the seller knows more about a product than the buyer. A lemon is an American slang term for a car that is found to be defective only after it has been bought. I have devoted a later chapter in the book to the occurrence of such phenomena in gambling.

[53](#) This is based on a comparison of closing prices for about 50 well-known bookmakers via Oddsportal.com.

[54](#) To a significant degree, the bookmaker’s odds are simply a reflection of the opinions of his customers, expressed through money wagered. If more money is bet on an outcome relative to other possible outcomes, the bookmaker will likely shorten his odds for that outcome, and *vice versa*. This idea will be further examined later in the book.

[55](#) Of course, if the real odds we are using to estimate ‘fair’ odds themselves contain the systematic bias we are trying to find, this in itself presents a problem. For our purposes here, however, it is probably safe to assume that the relative rates of winning and losing by teams will provide a much bigger influence towards short term returns than small inaccuracies in odds. For example, the difference between winning and losing at odds of 2.00 is 200% of the stake. By contrast, the difference between winning at 2.00 and winning at 1.95 is just 5% of the stake.

[56](#) This is simply the profit won from a bet with stake 1/odds. For this model I have preferred to use risk-adjusted returns over level stakes returns to minimise variance in scores. A lucky win at odds of 10/1, for example, will have a much bigger (and potentially unwarranted) influence on returns calculated with level stakes. The same would be true for an unlucky loss at odds of 1/10.

[57](#) P-value = 0.004 according to a paired two-sample Student’s t-test. I’ll be talking a little more about statistical testing, the t-test and what p-values mean in the next chapter. For now it is enough to say that a p-value of 0.004 implies that such an outcome could be expected to happen by chance 4 times in 1,000.

[58](#) P-value = 0.015 for negative match ratings, p-value = 0.016 for positive match ratings (one-

sample t-test).

59 Gross, W. H. Cult Figures. *Investment Outlook* (August 2012),

<http://europe.pimco.com/EN/Insights/Pages/cult-figures.aspx>.

60 The S&P 500, or the Standard & Poor's 500, is an American stock market index based on the market capitalisations of 500 large companies having common stock listed on the New York Stock Exchange or NASDAQ.

61 A Ponzi scheme is an investment operation where the operator, an individual or organisation, pays returns to its investors from new capital paid to the operators by new investors, rather than from profit earned by the operator.

62 Kahneman, D. & Riepe, M., 1998. Aspects of investor psychology. *The Journal of Portfolio Management*, **24(4)**, pp.52-65.

63 Odean, T., 1999. Do investors trade too much? *American Economic Review*, **89(5)**, pp.1279-1298.

64 Barber, B. M. & Odean, T. Boys Will Be Boys: Gender, Overconfidence, And Common Stock Investment. *The Quarterly Journal of Economics* (2001), **116(1)**, pp.261-292.

65 Odean, T., 1998. Are investors reluctant to realize their losses? *Journal of Finance*, **53(5)**, pp.1775-1798.

66 De Bondt, W. & Thaler, R., 1985. Does the stock market overreact? *Journal of Finance*, **40(3)**, pp.793-805.

67 Milgram, S., Bickman, L. & Berkowitz, L., 1969. Note on the Drawing Power of Crowds of Different Size. *Journal of Personality and Social Psychology*, **13(2)**, pp.79-82.

68 Asch, S.E., 1951. Effects of group pressure upon the modification and distortion of judgments. In Guetzkow, H. *Groups, leadership and men; research in human relations*. Oxford: Carnegie Press.

69 *A Study of Conformity*, Solomon Asch, <https://www.youtube.com/watch?v=rPEDS-0jMgs>.

70 Berns, G.S., Chappelow, J., Zink, C.F., Pagnoni, G., Martin-Skurski, M.E. & Richards, J., 2005. Neurobiological correlates of social conformity and independence during mental rotation. *Biological Psychiatry*, **58(3)**, pp.245-53.

THE HARDER I WORK, THE LUCKIER I GET

Donald Rumsfeld, the two-time US Secretary of Defense, once famously said:

“There are known knowns. These are things we know that we know. There are known unknowns. That is to say, there are things that we know we don’t know. But there are also unknown unknowns. There are things we don’t know we don’t know.”

In making such a declaration, many of his detractors regarded this as a classic example of political obfuscation. Others who were more supportive, however, considered it to represent a brilliant distillation of quite a complex subject: the business of trying to predict the future when faced with uncertainty. Essentially, Rumsfeld was saying three things:

1. Some outcomes are known. This is the easiest way to make decisions. For example, if I hold out my hand and drop a ball, it will drop to the ground because of gravity. It is an entirely predictable outcome. These are the known knowns.
2. Other outcomes are unknown, but probabilities are known. This is risk. As Michael Mauboussin, author of *The Success Equation* articulates, “*we don’t know what is going to happen next, but we do know what the [probability] distribution looks like.*” These are the known unknowns, things concerning games like roulette, slot machines and lotteries.
3. Finally, some outcomes are unknown and probabilities are unknown. This is uncertainty, and deals with things like political and economic forecasting, sports betting, poker and financial investing, and other markets of psychology. As Mauboussin puts it: “*we don’t know what is going to happen next, and we don’t know what the possible [probability] distribution looks like [either].*”

When faced with such uncertainty, how do we know whether the forecasts we make are any good? To be more specific, how can we tell whether the successes we experience have been achieved through hard work, knowledge and talent or simply as a consequence of blind chance? Samuel Goldwyn, the film producer, once remarked: “*the harder I work, the luckier I get.*” Traditionally, this is seen as an ironic paraphrasing of a seemingly self evident truth: better outcomes go to smarter people who try harder. Goldwyn’s insight, however, was obviously in recognising that whilst hard work and talent play their part in shaping futures, luck still has a big influence to play. In statistical terms, luck is always there but skill changes the shape of the probability distribution and the position of the average. In competitions, however, we may find something paradoxical: the harder we work, the more our outcomes are dependent on luck. This time, no irony is implied. Genuinely greater skill leads to a bigger influence of luck. How can that be, you might ask. This chapter attempts to find out.

Luck and Skill

To begin, let’s first consider each of Rumsfeld’s three scenarios in terms of the elements of luck and skill. In the first instance, luck plays no part at all in making predictions. We can essentially say, with absolute certainty, what is going to happen based on the specific knowledge about specific cases. If I play a game of chess against Gary Kasparov or a game of tennis against Roger Federer it is a given that I will lose, because of their vastly superior skill levels in their specific domains. In the second instance, skill plays no part whatsoever. Outcomes are purely a consequence of luck and predictions made are based purely on the known probabilities of each outcome. In statistical jargon, the known (prior and unconditional) probability of some outcome is called the base rate. The base rate for throwing a 6 with a fair dice, for example, is $1/6$, whilst for throwing boxcars (two 6s) is $1/36$. Scoring a slots jackpot won’t change the odds of winning another, since every outcome has no memory of the preceding one. Finally, in the third instance, since the probabilities of outcomes are unknown, there exists scope for some people to make better predictions than others. Luck still plays a significant role, but skill is also a factor.

Trying to figure out how the two things interact is no easy task.

Mauboussin has attempted to define skill as “*the ability to use one’s knowledge effectively and readily in execution or performance.*” It basically says you know how to do something and can do it when called upon. Obvious examples would include musicians, artists, engineers and surgeons, as well as chess players and professional sportsmen and women. A strong relationship between cause and effect exists, where doing the same thing again will yield a similar outcome, offering a reliable learning feedback mechanism. Mauboussin classes such skilful activities as linear and stable. Luck is a lot more slippery. Of course, Laplace believed there was no such thing; for him it merely demonstrated our ignorance of perfect knowledge. We might consider it as success or failure apparently brought about by chance, without purpose or predictable causes, rather than through one’s own intentional actions. What, then, is chance? Evidently, our definition represents a tautology; chance might be defined as the occurrence of events in the absence of any obvious intention or cause, so we’re back to where we started. Clearly, luck is an elusive concept. Mauboussin likes to think of randomness or chance operating at a system level and luck at an individual level. “*If I gather 100 people and ask them to call coin tosses, randomness tells me that a handful may call five correctly in a row. If you happen to be one of those five, you’re lucky.*”

Mauboussin has constructed a ‘Skill-Luck Continuum’ suggesting how much each element contributes to various activities, including games involving chance (roulette, dice etc), games involving sports (football, tennis and so on) and games involving psychology (poker, betting, investing). At one end we have games like chess, where skill is dominant and luck is largely irrelevant. We would have little trouble in recognising that a grandmaster’s victory over me in a game of chess is a consequence of skill. Through years of practice and beneficial feedback linking specific play to positive outcomes, the grandmaster has become a much better player than I am. In such domains, even slight differences in skill levels between one player and another have a dramatic influence on the outcome of a game where potentially just a tiny handful of more skilful moves can change the result. Another way to view this is to consider how easy it would be to deliberately lose a game of skill. I’m not a particularly skilled chess player, but I could probably manage it in about 5 or 6 moves. Luck plays no

part in that; throwing a game would be all my own work.

At the other end of the spectrum in the domain of pure luck we find the traditional casino games of roulette, slots and craps, in addition to bingo and lotteries. Remember, probabilities of outcomes are known, but not the outcomes themselves, which are clearly a matter of pure chance. What can you learn from the spinning of a roulette wheel and the ball landing black? Of course, our heuristic blind spots encourage many of us to believe a large number of consecutive blacks should make a red more likely – the gambler's fallacy – but in truth there is no meaningful feedback at all. The game is completely memoryless. Try intentionally losing money in a game of craps. How you throw the dice has absolutely no influence on the numbers that land, although unsurprisingly that hasn't stopped millions of people believing that irrational superstitions can help them. Some even believe they can physically control the way the dice are thrown.

Assessing the relative contribution of luck and skill in competitive sports like football, rugby, tennis, basketball, cricket, snooker and darts is arguably quite difficult. Undeniably, a professional sportsman will be a highly skilled operator in his field of expertise. Ask Michael van Gerwen to score 180 with three darts and he'll probably manage it about once in 5 to 10 attempts. Over 30 years, by contrast, I've done it just twice. Similarly, football players like Lionel Messi seem to be blessed with talents of ball control, pass execution and vision of play that come from another planet. These skills are not lucky; they represent the accumulation of tens of thousands of hours of deliberate practice and learning through linear feedback, where cause is intimately and overtly linked to effect. Nevertheless, there will always remain some underlying luck since not everything a sportsman does is perfect. Van Gerwen can't score 180 with every three-dart throw.

Mauboussin helps us interpret the way luck and skill interact by means of his 'Two-Jar Model.' It explains why smaller differences in absolute skill lead to a greater influence of luck, and indeed how, paradoxically, greater absolute skills can often mean that luck is a bigger contributor to outcomes. Following Mauboussin, imagine that you have two jars filled with numbered balls. The numbers in the first jar represent skill, those in the second, luck. The smallest number in the skill jar is 0, equivalent to no skill at all. Numbers in the luck jar can be negative, implying bad luck, and positive numbers, implying good luck. You draw one ball from each and

add the numbers together to get a total score. Let's define good outcomes as total scores above 0 and poor outcomes below 0. Suppose each jar has just three balls: in the skill jar we have 0, +1 and +2, whilst the luck jar has -3, 0 and +3. The range of possible scores is as follows:

Skill = 0, Luck = -3, Total = -3

Skill = 1, Luck = -3, Total = -2

Skill = 2, Luck = -3, Total = -1

Skill = 0, Luck = 0, Total = 0

Skill = 1, Luck = 0, Total = 1

Skill = 2, Luck = 0, Total = 2

Skill = 0, Luck = 3, Total = 3

Skill = 1, Luck = 3, Total = 4

Skill = 2, Luck = 3, Total = 5

We can still have a positive outcome (+3) even in the absence of any skill; we just need to be lucky. Conversely, we can have a poor outcome (-1) despite being skilled (+2) in the event that we are very unlucky.

Evidently, if we repeat the draws many times, skill, if any is present, will have an ever-increasing influence on the outcome. In the short term, good and bad luck will affect your scores, but in the long run they will cancel each other out, leaving just your residual skill to come through. How quickly that will happen, will depend upon the relative strengths of luck and skill on each repetition. Where skill is the more dominant factor, better outcomes will be achieved more quickly; where luck is the principal influence, that will take longer. All of this is just another way of saying sample size matters. The larger the number of repetitions, the less luck influences the overall outcome, provided of course some skill is present. Essentially, this is a consequence of the law of large numbers. The two charts below hopefully illustrate this phenomenon.

Evolution of performance with skill (luck small)

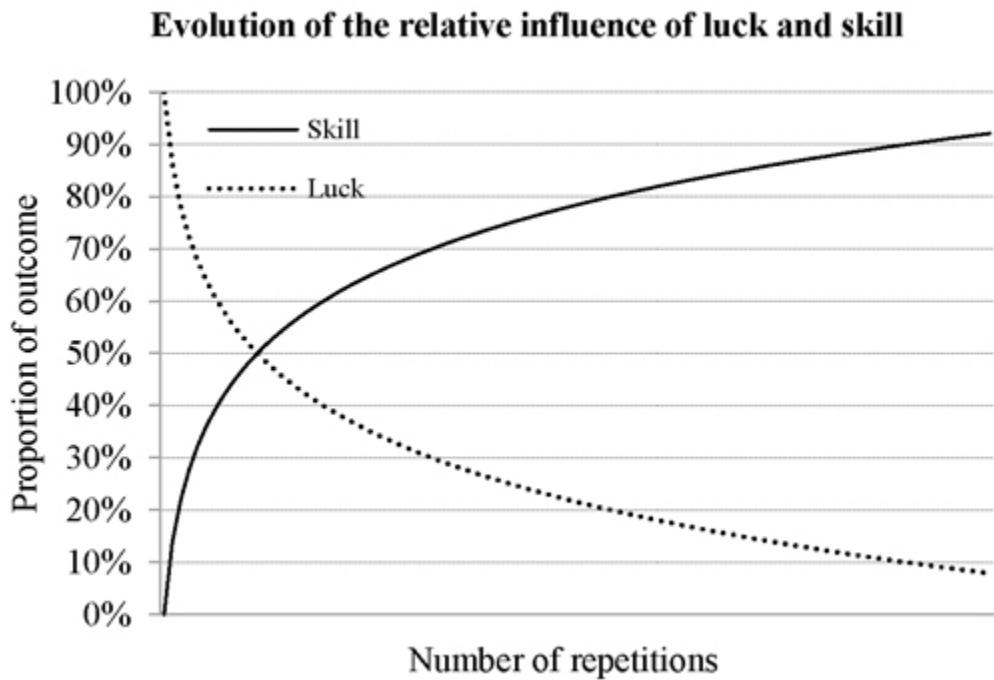


Evolution of performance with skill (luck large)



Perhaps the relevant question, then, is not so much whether something **is** a game of skill but rather **when** it becomes a game of skill. Where luck and skill are both present in games involving repetition, the relative contribution

of skill increases with the number of repetitions. I've tried to illustrate this schematically in the diagram below. Understandably, the speed at which skill starts to show its hand in a mixed (luck-skill) game will depend on the relative weighting of luck and skill for each repetition.



We can also use the two-jar model to think about games of pure skill and pure chance. For the latter, all the balls in the skill jar will be numbered 0. The evolution of outcome will then represent a random walk, switching between good and bad outcomes indefinitely, with final expectation of 0, as for the skill = 0 cases in the charts above. For the former, all the balls in the luck jar would be numbered 0, and the evolution of outcome would be illustrated by a perfect linear trend line, where the expectation is equivalent to the skill level.

Most sports, however, take place within an environment of competition. When skilled players compete against each other the outcome of a contest will then depend on the relative difference of the skills of the competitors. The smaller the difference, the less likely it is that skill will influence the outcome. Competitive sport is really just a relative skills contest. When the relative difference between competitors' skills is small, the underlying luck will become far more influential. We can illustrate this again with the two-

jar model where two equally skilled players (for example skill = 2) are competing against each other. To determine the outcome of the contest, we must subtract one total score from the other. Since balls drawn from the skill jar will always be numbered 2 for each player, skills will simply cancel out. Luck, on the other hand, will be different for each player on each repetition. Sometimes player 1 will be luckier, sometimes player 2. In such a situation, the outcome of a contest will purely be a consequence of luck, with skill playing no part. Imagine now we increase their skill levels to 20. In absolute terms, each player is far more skilled. In relative terms, however, they've stood still, and luck is the only determining factor deciding who wins the contest. When considering the comparative influences of luck and skill in single-player performance, absolute skill is the important factor. By contrast, when analysing a competition, relative skill – how much more or less skilled one player is compared to another – is all that matters.

Arguably, for two-player sports like tennis, snooker and darts, even quite small amounts of relative skill will be more influential than luck, especially over the longer term. Two-player games are conceivably more linear in terms of causality (i.e. the things players do to bring about desired outcomes) than team sports, particularly where games involve a large repetition of plays to reach a final result where a small differential in skill will be magnified exponentially. Ian Stewart, an English mathematician, has calculated that, where one tennis player has a 53% probability of winning a single point, this will translate into an 85% probability of victory in a 5-set match⁷¹. Team sports like football, rugby and cricket probably involve a greater element of luck. Differential skill levels between players will have an influence; Manchester United, for example, have better players than Cambridge United and so can be expected to win more often, but with a far greater number of player interactions which are arguably less linear and more complex, where cause and effect are less obvious, randomness will play a bigger role. How a game is scored will also be an influencing factor. Generally speaking the larger the number of scores a game has, the greater the influence of (relative) skill. Again, as for tennis, this is simply a consequence of the law of large numbers. Football, for example, will experience more lucky outcomes than rugby. Given our neurochemical

predisposition for uncertainty, perhaps that's why it is universally so appealing.

The Paradox of Skill

Counterintuitively, as absolute skills improve, performance becomes more consistent, and therefore luck becomes more important. Mauboussin calls this the 'Paradox of Skill' and has referred to an analysis of Major League Baseball batting averages by Stephen Jay Gould⁷², the late Harvard palaeontologist and baseball fan, as a prime example. In 1941, Ted Williams, a player for the Boston Red Sox, had a batting average of 0.406. Considering that typical batting averages have remained largely unchanged since the origins of professional baseball in the 1870s (around 0.25 to 0.28), this was a remarkable achievement, and something that has not been repeated since. Arguably, however, Williams would not score anything like that average in today's league, given the improvements in training, fitness, diet and general professionalism. So what's going on?

Firstly, the batting average is simply a measure of relative skill, between the pitcher on the one hand and the batter on the other. As Major League Baseball has become more professionalised, batters have individually become more skilled at hitting. At the same time, however, pitchers have become more skilled at pitching. As Mauboussin says, it's like an arms race: absolute skills improve across the board, but relative skills, on average, remain more or less the same. Secondly, whilst overall skill levels have improved, the difference between the best and worst hitters (and pitchers) has shrunk. Gould explains this by imagining there to be a 'wall' of human ability. Think of 100m times: the best sprinters have been getting progressively faster but at an ever diminishing rate. The lowering of the 100m world record occurs less frequently today than say 50 or 100 years ago. This is because the faster we become, the closer we get to the physical limits of possibility. Conceivably no person will ever be able to run 100m in 8 seconds, and probably not 9 either. More crucially, more of those running at the top level today are closer to the physical limit. In the men's Olympic marathon, for example, the difference in finishing time between the top 20 has fallen from about 30 to 40 minutes before the Second World War to

about 5 to 10 minutes in the modern era. Gould argued that, in the early years of professional baseball, a few players were already approaching the ‘wall’ but most were still quite some way away. Over time, progressively better hitters (and pitchers) were replacing the weaker ones, and as a consequence the difference between best and worst has narrowed. Arguably, this has happened across most, if not all, professional sports. The consequence is that we see fewer surprises and outlier results. Given this phenomenon, it makes the achievements of athletes like Usain Bolt, Lionel Messi, Cristiano Ronaldo and Tiger Woods seem even more remarkable.

In statistical terms, Gould’s idea should reveal itself in the amount of variation seen in batting averages, quantified by means of the standard deviation. Sure enough, during the first years of professional baseball in the US (1870s), standard deviation in batting averages was around 0.05, meaning around two-thirds of all batting averages were roughly in the range 0.2 to 0.3, with about 95% between 0.15 and 0.35. Today, the standard deviation is about half of what it was. Consequently, whilst in the 19th century we might have expected a batting average of 0.40 to appear once in every 1,000 batters, today that might be more like 1 in a million. Indeed, the almanac of historical seasonal batting averages shows that, of the 27 occasions when it was achieved, 14 were before 1900, and of course none since Ted Williams in 1941. Thinking about this from the perspective of the two-jar model, as the variance in (relative) skill diminishes the variance in luck will assume an ever-increasing importance in the calculation of outcomes. As Mauboussin says, “*if everyone gets better at something, luck plays a more important role in determining who wins.*”

So what about games of psychology – poker, financial investing and sports betting – where the interaction of players’ decisions matter? Mauboussin has positioned these much closer to the pure luck end of the spectrum. For poker, that’s pretty understandable. Indeed a paper in the *Journal of Gambling Studies*⁷³ went as far as to say that poker, under certain basic conditions, should be regarded as a game of chance. Sadly, those conditions were far too basic; it considered just 300 players, dividing them into two groups of squares and sharps, depending on their level of interest in the game, and had them play just 60 hands. On the other hand, a team of researchers writing in the journal *Science*⁷⁴ in January 2015

claimed to have solved the game of heads-up limit hold'em poker by developing an algorithm (called counterfactual regret minimisation) through the analysis of a quintillion hands (1 with 18 zeros) that was capable of perfect play. Nevertheless, it still required playing more than 1,500 hands to learn its champion-level skills, achieving this via adopting a game-theoretic approach to the very human psychological condition of bluffing. Whilst fallible in the short term, its developers claimed to have demonstrated that it was unbeatable over the long run. Evidently, as such an algorithm demonstrates, whilst being able to read your opponents' play whilst hiding your own is undeniably a talent, a much greater influence will be the cards you are dealt, at least in the short term. The role of skill may only become evident after many rounds of play, and certainly more than 60. Potentially, only a small proportion of smart players will be able to sufficiently exploit it to create positive expectation, especially when paying a rake to the poker room. According to Ingo Fiedler and Jan-Philipp Rock⁷⁵, from the University of Hamburg, the point at which skill overtakes luck in a game of poker, what they call the critical repetition frequency, will occur anywhere from several hundred to several thousand hands, depending on game design and the presence of skilled players.

For sports bettors, however, the idea that betting may well mostly be a matter of luck might seem particularly confusing. Sports, as I've previously clarified, are games involving a lot of absolute skill, and sometimes even relative skill too if played over many repetitions. This is true, but the business of sports betting is a secondary market to the sports themselves. The following conversation hopefully illustrates the point.

Question: *“Baseball is not a lottery. If it was, we wouldn’t have good and bad teams. Every team would be the same, equal chances for everyone. If that’s a lottery, it means anyone can play it. So take me and make me the pitcher, how much luck do you think I’m gonna get?”*

Reply: *“Undeniably not a lot, but if you were made a pitcher, the odds for the other team winning would shorten to 1.00000000000001”*

Prizes are not awarded to the bettor who can predict the winner. If they were, it would simply be a matter of always backing the most skilful teams and players. On the contrary, when we bet we are not merely predicting uncertain outcomes, rather we are also assigning monetary values to

opposing opinions according to the perceived probabilities of outcomes. To all intents and purposes, betting represents a derivative market of opinions about predictions. Whether those playing the primary sports on which people bet are skilful has nothing to do with whether those betting are skilful themselves. Essentially, the odds act as a kind of skills handicap, reducing their influence relative to luck. In one respect at least, the religious critics of betting were making a relevant point: by assigning odds to outcomes, bettors as much as possible are significantly increasing the influence of chance by handicapping their predictive skills. Of course, this is only to ensure that a mutually agreed wager can happen in the first place. Yes, it's easy to predict that superior teams and players will more usually beat inferior teams and players, but you'll get paid less for doing so. Whether bettors purposefully want to eliminate skill altogether is another matter. Arguably, they don't, since the attraction of betting is pitting one's wits (skills) against others, as the traders of Stan James have testified. Unknown unknowns, furthermore, provide the possibility, if not the inevitability, of positive expectation, a reward unavailable in (casino) games of risk. And as Daniel Kahneman has shown us, we all love a possibility.

Consider again Manchester United versus Cambridge United. A bet requires the consent of two parties, the backer and the layer, to agree on the acceptability of the odds. The backer thinks Manchester United will win. The layer thinks the opposite. For the actual teams, the prize is the same: entry to the 5th round of the FA Cup. Such reward equality will be completely unavailable in the betting market. Anyone with the slightest interest in football will predict that Manchester United is far more likely to win, without any recourse to sophisticated forecasting techniques. Should the backer ask the layer for an even money return on his investment it would undoubtedly be refused. Consequently, the monetary reward available for the correct prediction will need to be reduced relative to that available for a Cambridge United victory, sufficient to ensure that both backer and layer can reach an agreement whereby a bet will take place. In this case, that might be 1/10 for Manchester United (implying 10/1 for Cambridge). Essentially this process is one of bartering and compromise although in practice this will all take place swiftly and automatically. Both backer and layer will intuitively have in their minds roughly what they think a suitable price would represent for them. Presumably, overconfidence

allows for both parties to hold the perception that each of them has secured some sort of positive value expectation at the expense of the other, which of course is a logical impossibility. Of course, without this logical impossibility the bet would not take place, since both rationally self-interested parties are motivated by the expectation of making a profit based on information that is better than his opponent's, not throwing away money for the sake of it. Naturally, if one or both parties fail to agree, the bet is not struck anyway. This compromising of opinions implicitly ensures that the odds are more likely to be closer to the 'true' probability of Manchester United (or Cambridge United) progressing. Conceivably the odds probably won't perfectly represent the true probability of outcome (and we'll never know anyway since we're dealing with unknown unknowns), but by purposefully handicapping the role of predictive skill (to allow the bet to proceed) the influence of luck will be significantly greater. All of this is rather stating the obvious, but it does serve to highlight the significant role that luck plays in betting. The crucial question is whether some bettors are consistently better at perceiving what little positive expectation might be available than others. In other words, is there any skill in sports betting, and if there is will more and more people expressing it paradoxically lead to a greater influence of luck?

Much the same is true in the financial world where buyers and sellers of assets must agree on a price if a transaction can take place. No agreement; no transaction. Daniel Kahneman gets to the heart of why financial investing must logically be a relatively low-skill, high-luck activity. In *Thinking Fast and Slow* he tells the story of an encounter with an investment manager at a Wall Street firm, and specifically a question he posed. "*When you sell a stock, who buys it?*" More generally, what makes one person buy and the other person sell? What do the sellers think they know that the buyers don't? Are the sellers all possessed with superior talents of share price forecasting? And if they are, does that mean the buyers must all be very wealthy, stupid, charitable, or even all three? This thought experiment illustrates the illogicality nicely. If sellers knew consistently more than buyers, eventually the buyers would stop buying. This is not to argue that financial investors (and sports bettors for that matter) have no idea what is going to happen in the future but merely that the process of two sides – the buyer and seller (or backer and layer) –

mutually agreeing on a price pits one set of skills against the other in such a way that they are closely balanced. Whether one side will consistently perform better than the other is then simply a matter of the skill differential between them. As such, the trading of opinions might well be described as a relative skills market, much like for batters and pitchers engaged in competitive baseball. Unlike in sports, however, where good and bad sides alike get the same reward for winning, rewards in betting and financial trading are intentionally settled according to the probabilities of outcomes. Such a handicapping process arguably turns the business of betting and trading into much more of a lottery than sports.

Furthermore, as Mauboussin has explained, paradoxically as more and more forecasters make better and better predictions the differences between them get smaller, and the influence of chance increases. Indeed, referring to the performance of institutional professional investors, he says that “*the more everyone’s level of skill looks the same, the more you’d expect the range of excess returns for money managers to shrink.*” Sure enough, that’s exactly what Peter Bernstein⁷⁶ found when he analysed the variance in excess returns for mutual funds from 1960 through to 1997. Over the period, the standard deviation of Morningstar fund returns trended downwards from about 13% to 8%. Just as in baseball, the big hitters were disappearing; not because they were less skilled at forecasting returns, but because they were competing against more of the same. Similarly, in my capacity as a verifier of sports betting advisory services, I have observed a declining variance in aggregated year-to-year yields, with the standard deviation in running 5-year samples dropping from about 2.5% in the period 2002-2007 to just 1% by 2009-2014. Arguably, this period witnessed the biggest growth in online sports betting, but as more and more took up the challenge of beating the market the harder it became to do so, with sharper forecasters converging towards a ‘wall of truth’. Players may have become sharper in absolute terms, but with that profitable returns are now harder to come by. The 0.400 hitters of sports betting are disappearing.

Deciphering Luck and Skill

There are several ways of determining whether an activity is dominated by

skill or rather by luck. I've already touched on a couple of them. Firstly, can you lose on purpose? The greater the influence of luck, the harder that will be. For games of pure luck, it's impossible. For games of pure skill, it would be impossible not to if you tried. Secondly, there is sample size. Remember, the greater the number of repetitions in a mixed luck-skill contest, the more likely it is that relative skill will begin to reveal its influence. If you have an activity where the outcomes are largely a consequence of skill, you don't need a large sample size to draw meaningful conclusions. How long would it take you to realise that Roger Federer was better at tennis than I am? In contrast, how many hands could I play with Daniel Negreanu, winner of six World Series of Poker bracelets, before it was obvious I was a poker square? The sample size required for skill to show its hand will be much smaller in the former than the latter. Conceivably, it would be pretty obvious after just one or two points that I can barely serve a tennis ball, never mind return one received at 125 mph. In a game of poker, however, the much greater influence of luck in the dealing of cards will enable me to play tens, if not hundreds, of hands before my inferior ability to bluff and read opponents will prove to be my undoing. Where games involve pure luck no sample size will be sufficiently large to reveal the influence of skill, obviously because there isn't any.

Another method is determining the clarity of causality. How obvious is the relationship between cause and effect? When the relationship is clear, repeating the behaviour should deliver the same outcome most of the time. Such an activity is said to be linear, and feedback from outcomes will prove to be a very useful teacher. Clearly, Michael van Gerwen throwing darts at the Bull's eye is a fairly linear activity, in contrast to mine which, like tossing a coin, is not. How linear is determining the value expectation in the betting odds of a football team or the next movement of a share price? Based on the arguments presented above and the data following in the next chapter, we'll see probably not very much at all. Getting it right once conceivably has little bearing on whether we can get it right the next time. Of course, our cognitive biases, in particular attribution errors and our self-serving overconfidence, will encourage us to believe that causality when trading opinions with others is often much clearer than it really is, particularly when we are right. Remember, we prefer causal narratives to statistical fuzziness, because of our evolutionary craving for control.

Causality, even when spurious, is a better story than randomness. Daniel Kahneman calls this the illusion of validity. We might also describe it as the illusion of skill. Such an illusion arises because we are often blind to our own blindness. True intuitive expertise, as professional sports people will confirm, is learned from prolonged experience with good feedback in linear environments. Later in the book, I'll examine why gambling markets like betting and investing are probably not predisposed to such learning techniques (for the majority of people at least), and are in fact what Kahneman terms zero-validity environments.

Closely related to the idea of linearity is regression or reversion to the mean, a concept we've discussed previously. The greater the amount of luck (relative to any skill), the faster your score or total will revert to the (expected) mean. You can see this illustrated in the charts earlier in the chapter based on Mauboussin's two-jar model. If your activity relies entirely on skill, the balls in your luck jar will all be numbered 0. Since the score for your skill in this thought experiment doesn't change, there is no reversion to the mean, and your score expectation is simply the same as the number on your skill ball. If I play Roger Federer, it's more than likely he'll win every point – no reversion to the mean. In contrast, if your activity is purely a matter of luck, all the balls in your skill jar will be numbered 0. Since the average score of the numbered balls in the luck jar is 0, the expected value of the next outcome will be 0 too; in other words, there is complete reversion to the mean, as we see for games in a casino. A time series of scores for an activity based purely on luck looks just like a drunken man's walk; that is to say, random. Many profits series of sports bettors and financial traders look very much like that. When good (bad) luck takes you into positive (negative) territory, reversion to the mean will be expected to return you to the average. The greater the influence of luck, the sooner you can expect that to happen.

Yet another means of determining the relative contribution of luck and skill is what is known as 'True Score Theory.' This is a theory about measurement and is a very simple one, if not necessarily proven: observed outcome is true ability (skill) plus random error (luck). More specifically, it states that the variance in outcome is the sum of the variance in skill and the variance in luck. Variance is simply a measure of variability in observations (for example MLB batting averages, tennis win percentage, betting returns

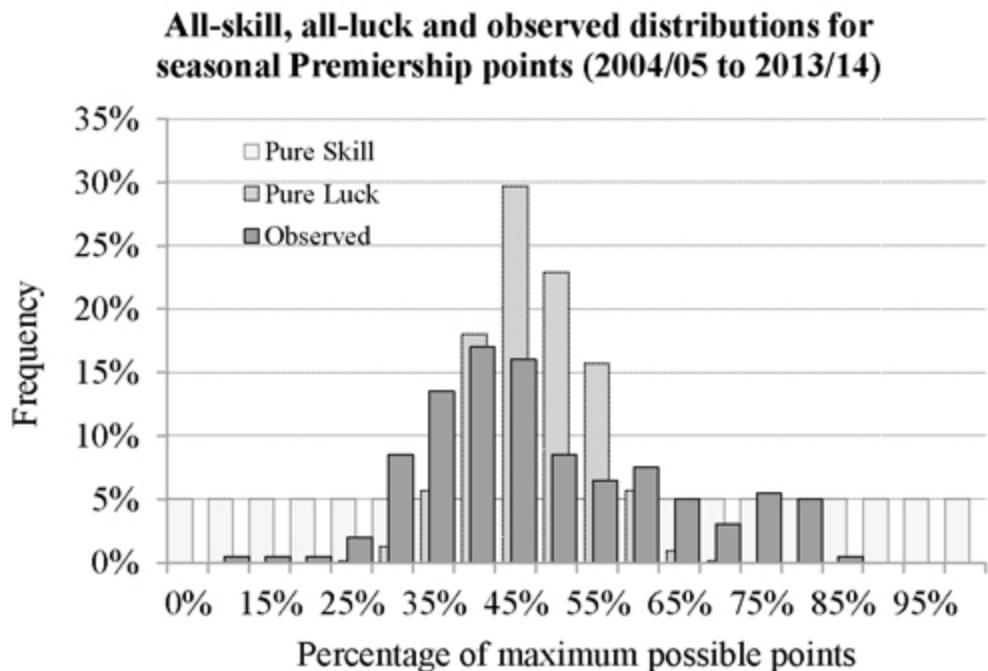
etc.), and in statistical terms is the square of the standard deviation, itself a mathematic measure used to quantify the spread of observations around their average⁷⁷. This is perhaps best illustrated by means of an example. Consider the spread of points for Premiership football teams for the 2013/14 season as shown in the table below.

Team	Games	Points	Pts per game	% of max.
Manchester City	38	86	2.26	75.44%
Liverpool	38	84	2.21	73.68%
Chelsea	38	82	2.16	71.93%
Arsenal	38	79	2.08	69.30%
Everton	38	72	1.89	63.16%
Tottenham	38	69	1.82	60.53%
Manchester United	38	64	1.68	56.14%
Southampton	38	56	1.47	49.12%
Stoke	38	50	1.32	43.86%
Newcastle United	38	49	1.29	42.98%
Crystal Palace	38	45	1.18	39.47%
Swansea	38	42	1.11	36.84%
West Ham	38	40	1.05	35.09%
Sunderland	38	38	1.00	33.33%
Aston Villa	38	38	1.00	33.33%
Hull	38	37	0.97	32.46%
WBA	38	36	0.95	31.58%
Norwich	38	33	0.87	28.95%
Fulham	38	32	0.84	28.07%
Cardiff	38	30	0.79	26.32%

The spread or standard deviation in the finishing points as a percentage of the maximum possibly achievable (shown in the final column) is 16.9% (or 0.169). Consequently, the observed variance is 0.0286. To calculate the

variance we should expect if finishing points were simply a matter of luck, we first need to calculate the expected points as a percentage of the maximum. We can do this by assuming that each of the three possible results for any game is equally likely. Consequently, a team's points expectation is 1.333 per game, or 44.44% of the maximum, although of course a third of the time a team will get 0, 1 or 3 points. A Monte Carlo simulation with 10,000 iterations to simulate finishing points according to this assumption yielded a standard deviation of 6.8% (or 0.68), and hence a variance of 0.0046. Consequently luck, according to true score theory, accounted for about 16% (i.e. 0.0046/0.0286) of the finishing points of Premiership teams in the 2013/14 season. Performing the same calculation for the last 10 seasons aggregated together revealed that skill accounted for approximately 80% of the distribution in finishing points, with luck about 20%. Intuitively, that seems to be about right since the best team in the Premiership generally wins the title with 2 to 2.5 points, whilst the bottom team gets relegated with about 0.67 points, roughly 3 to 4 times fewer.

A final method of separating luck and skill is via means of comparing the distribution of observed outcomes to those predicted by chance alone. In a sense this is a kind of visual interpretation of true score theory. It involves three steps: 1) see what happens if only skill is present; 2) see what happens if only luck is present; 3) compare them both to real observations. If all outcomes were dependent on skill alone we might reasonably assume that the superior team always beats the inferior team. From the resulting linear hierarchy of ability it will naturally follow that the distribution of points will also be linear from 114 for the top-ranked team down to 0 for the bottom team; there is no reversion to the mean. When all outcomes are dependent on luck we have full reversion to the mean which in this case is 44.44% of the maximum possible number of points. Obviously, outcomes will usually be less or more than the mean, but most will cluster close to it, forming the classic bell-shaped curve that students of statistics know as de Moivre's normal distribution. The finishing points for the last ten Premiership football seasons (2004/5 to 2013/14) have been modelled accordingly and their distributions shown in the chart below.



According to this model, there still remains an inherent amount of randomness in what happens at a Premiership football match, since stronger teams do not always beat weaker teams. Yet clearly it's not just a matter of luck; more teams perform far worse or far better than chance alone would predict. Of course, as I've already argued, skill in football is not the same thing as skill in football betting. The obvious question then follows: how much randomness exists in betting markets?

We all know what skill in sports looks like, but what about gambling? Obviously, there's no such thing in the strictest sense of the word 'gambling', and as applied to casino games, slots, bingo and lottery. When playing such games, all balls from our skill jar are numbered 0. Short term good and bad luck can be expected, in the long run, to fully regress to the mean. Sadly for players of such games, that expectation implies a financial loss (over the long term) to help the casino, bingo hall and lottery provider pay for the costs of offering them as well as make a profit. Of course, this hasn't stopped people trying to make money from pure games of chance for thousands of years. As we've learnt, a belief in our ability to control an uncertain future is quite ubiquitous, even when there is rationally no possibility for control at all. Those who win, and they are few in number, do

so purely because of good fortune. By way of example, a longitudinal study by researchers at the Harvard Medical School⁷⁸ into the effects of casino gambling behaviour at the online betting service provider bwin.com revealed that players (4,222 of them who registered in February 2005) lost an aggregated 3% on total amount wagered (in line with typical casino game expectation), with only 11% of them showing profit over the following two years. The longer they played, the greater the chance they would lose: a nice exposition of regression to the mean. 15% of gamblers who bet fewer than 100 times made some money. In contrast, just 6% of heavy players who bet over 10,000 times managed that. Financially speaking, casino play is a mug's game, played only by squares. Of course, as previously argued in the book, there's far more to gambling than just money.

What about sports betting, however, and for that matter financial investment? Here, we're dealing with unknown unknowns. As a consequence, many people who do bet consider it to be more like speculation or investment and less like gambling, according to the way we defined these terms earlier in the book. Surely there must be scope for some people to know these unknowns better than others, shouldn't there? In my opinion the answer, for almost everyone, is a resounding no. The evidence I will present in the next chapter will hopefully make it obvious why. Before I do that, however, we should take a look at some of the ways we might try to decipher the difference between luck and skill in betting. The ideas might equally be applied to other psychological markets like poker and finance.

To begin, we should define what we mean by skill in prediction markets. In investing it is typically characterised as the ability to take actions that will predictably and consistently generate a risk-adjusted return in excess of the appropriate benchmark. Such a definition is equally suitable for the world of betting. In financial investment, with the exception of shorter term trading, even after playing costs that benchmark might conceivably be positive as a consequence of the engine of capitalism, slowly transforming natural capital into individual and social capital in the form of goods (and services) that are more beneficial to our way of life. As such, it can be considered a non-zero-sum game (provided we conveniently forget about its detrimental side effects, for example environmental pollution). Betting,

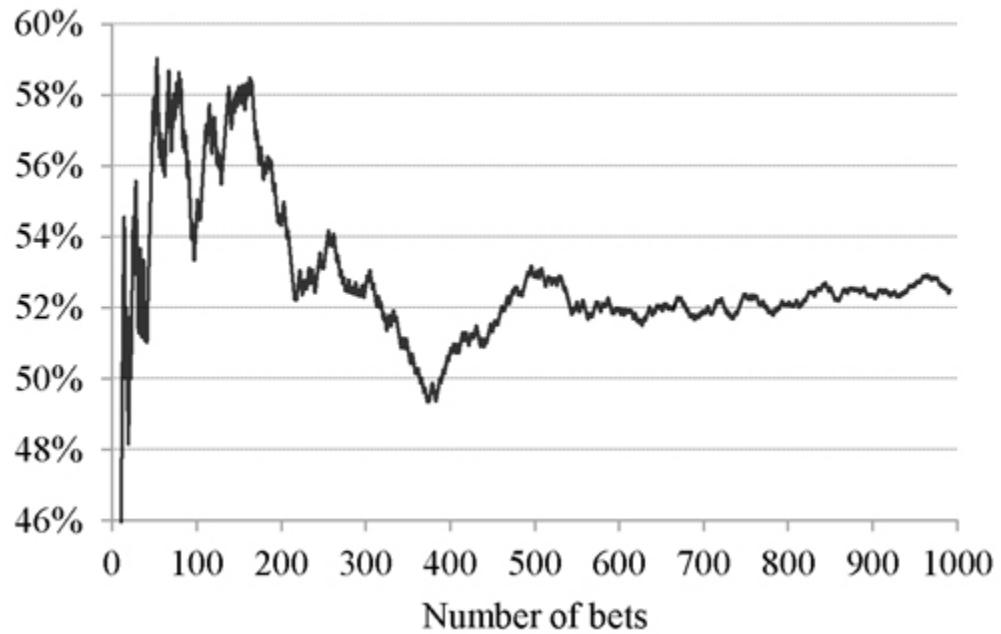
however, is zero-sum, and the benchmark will typically be negative since, like for games of chance, those offering such markets will not do so for free. Winners must wholly be paid for by losers and then some if we include the bookmaker's cut or the exchanges commission. In zero-sum games, as Mauboussin rightly points out, it is impossible for everyone in aggregate to generate returns in excess of the benchmark, as is the case in day trading. The question is whether anyone can do it consistently?

Having defined what we mean by skill, we now need to set about measuring it. Specifically, when a bettor beats his benchmark – that is to say, he makes a profit – how do we determine how much skill that involved? The previous discussion on distinguishing luck from skill dealt with groups of participants and observations, and the nature of their distributions. What about a single player in isolation? Fundamentally, the problem is one of blindness. We can never be absolutely sure what numbers we are drawing from our jars. Yes, we can expect regression to the mean to restore betting returns to the benchmark faster where they are largely the consequence of luck. But how do we know how long that should take? More importantly, if our betting returns are not regressing to the mean, how can we be sure that will continue to persist?

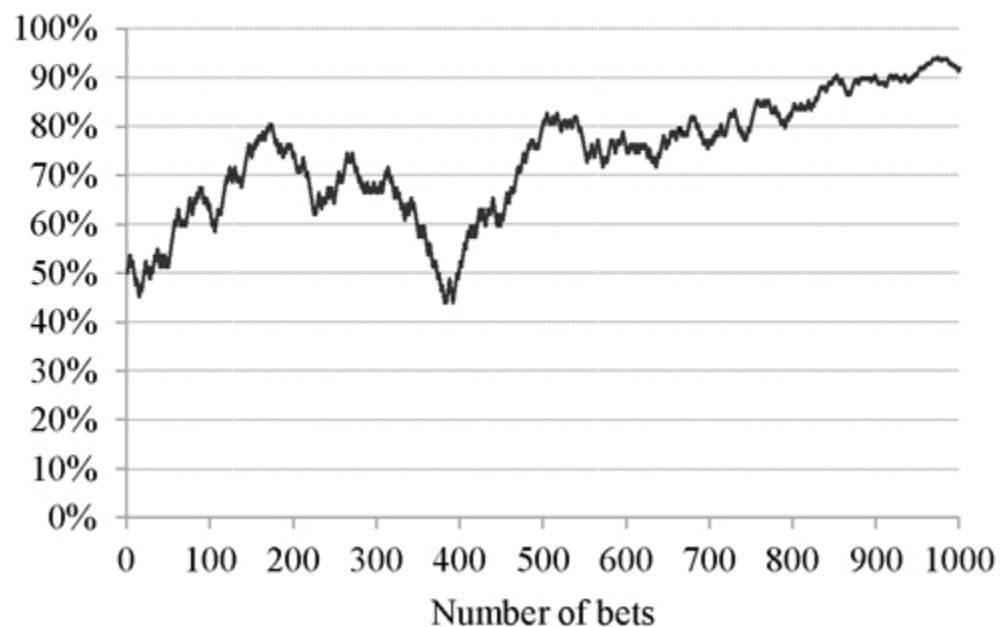
In statistics there are traditionally two approaches to this problem. The first is called the **Bayesian** approach. The data (for example, your record of wins and losses) are treated as fixed whilst the hypothesis (they arose that way because of skill) is random, and somewhere between absolutely true (1) and absolutely false (0). Starting with a probability that the hypothesis is true based on some prior knowledge or belief, each new observation allows one to update that probability. For example, before my first bet I might believe that the prior probability of me being a skilled football bettor is 50%, given that I knew nothing about it and it was a fair assumption to rate my chances of being so as 50-50. Winning my first bet allows me to update my prior probability with a posterior probability using Bayes theorem. With each and every bet I make, the posterior probability becomes the subsequent prior. Win and the probability that I'm a skilled bettor goes up, lose and it goes down. The evolution of the posterior probability that I am skilled might look something like that illustrated in the charts below. The first shows my cumulative win rate as the number of even money wagers I place increases. The second chart shows the Bayesian probability that I

might be skilled based on my previous performance up to that point.

Win rate



Probability of being skilled



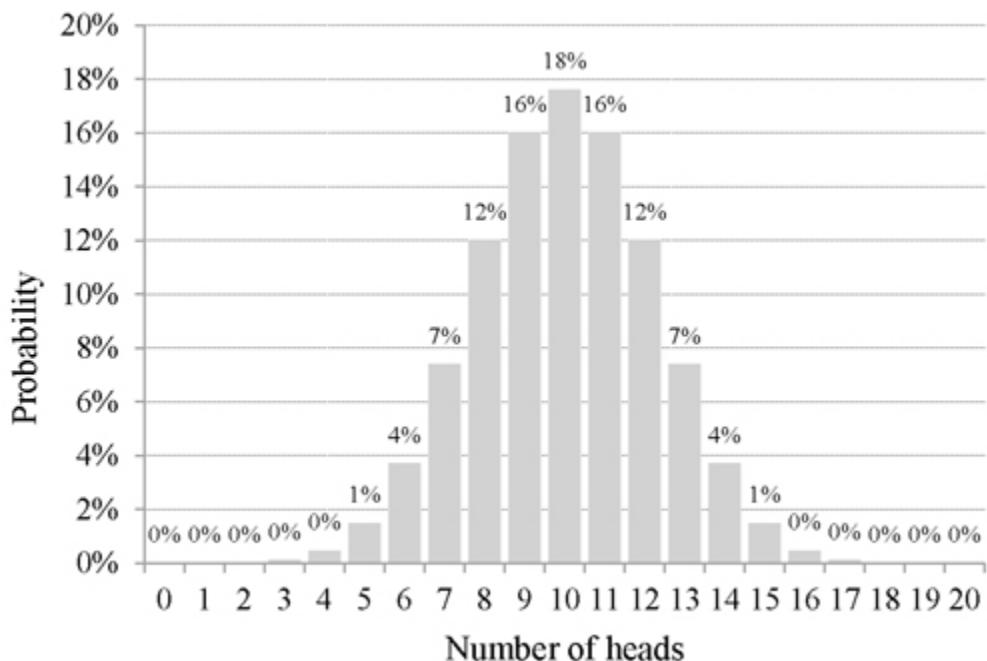
A Bayesian approach is particularly useful when predicting outcome probabilities in cases where one has strong prior knowledge of a situation.

But do we really have that in assessing the probability that someone is skilled at betting? My choice of 50% in this example was purely arbitrary and based on nothing else than guesswork. An alternative method to assessing the probability that I am a skilled bettor uses a **frequentist** approach. Whilst the Bayesian approach focuses on the probability of the hypothesis given the data, the frequentist approach focuses on the probability (or frequency) of the data given the hypothesis. This time the hypothesis is fixed – it's either true (1) or false (0) that I am skilled – whilst the data are assumed to be random. Typically, the frequentist approach starts with the null hypothesis, in this case I am not skilled and that my betting outcomes are all a consequence of luck. It then attempts to calculate the probability (usually called the p-value) by means of some statistic that the data we have observed, in this case my history of wins and losses, could have happened assuming the null hypothesis to be true. Finally, that probability is compared to an acceptable significance value (sometime called the α -value) such that, if $p < \alpha$, the null hypothesis is rejected in favour of the valid one.

Of course the frequentist approach is not without its problems either. A large one is the subjective choice of α . A typical figure in a lot of statistical hypothesis testing is 5%, with secondary and tertiary significance levels of 1% and 0.1% occasionally used as well. Its choice, however, appears to be largely a consequence of convention (groupthink?) following the pioneering work of the English statistician Ronald Fisher in the early 20th century, rather than anything else. One might reasonably argue that claiming statistical proof that a hypothesis is valid, whilst there remains a 1 in 20 chance that the observations on which it is based are just random, is not sufficiently conservative. Jacob Bernoulli was undoubtedly more cautious; remember moral certainty for him was 1 in a 1,000. Particle physicists are even more so, requiring a moral certainty of the order of 1 in millions when announcing the discovery of a new particle. Perhaps more important is the misunderstanding that results from hypothesis testing can lead to. It is easy to interpret incorrectly something having only a small probability of occurring by chance as something absolutely not occurring by chance at all. A probability of 5% that the observations are occurring by luck is not the same thing as a probability of 95% that they are occurring because of skill. It simply means that assuming the null hypothesis – that wins and losses in

betting are purely a function of chance – is true, what we have observed could be expected to occur 5% of the time. The weakness of the frequentist approach is that it treats truth as an absolute. In contrast, the Bayesian approach implicitly considers it to be probabilistic, provisional and always falsifiable. Despite these shortcomings frequentist hypothesis testing nonetheless offers us a very useful tool with which to analyse a history of betting, and to ascertain whether it is likely to have arisen through skill. In this respect my preferred statistic of choice has been the t-statistic, named after the test from which it comes, the student's t-test. To see how it is used I will first begin with a little digression into binomial probability. The material is adapted from my last book, since for the purposes of this discussion it is worth reproducing again.

Consider flipping a coin 20 times. The outcomes are purely a matter of chance provided the coin is unbiased, our null hypothesis. The chances of returning x number of heads are governed by a discrete probability distribution known as the binomial distribution. If we know that there is a 50% chance of flipping heads and a 50% chance of flipping tails on every coin toss, it's a simple enough exercise to do the maths and find those probabilities, as shown in the chart below.



Predictably, since heads and tails are equally likely to occur on each flip,

the most probable number of heads and tails after 20 coin tosses will be 10 and 10 respectively. But this does not mean that we will see 10 and 10 all the time. In fact, in this example, returning exactly 10 heads and 10 tails has less than a 1 in 5 chance of occurring. Sometimes we might see 9 heads and 11 tails, or 12 heads and 8 tails, or very occasionally 5 heads and 15 tails. What such a probability distribution does show, however, is that for the majority of occasions the number of heads will fall within a fairly narrow band concentrated around the average. For example, on nearly three-quarters of occasions, the number of heads will fall within two throws of the most typical, and in this case, average value of 10, that is to say, 8, 9, 10, 11 or 12. Such a distribution also allows us to calculate how likely certain outcomes are of occurring. For example, from this probability distribution, we know that the chance of flipping at least 14 heads from 20 coin tosses is roughly 6% (the cumulative probabilities of flipping 14, 15, 16, 17, 18, 19 or 20 heads). Suppose we win \$1 for heads and lose \$1 for tails, the binomial distribution tells us we have about a 41% probability being in profit after 20 rounds. Should we end up with a profit of \$16 (18 heads and 2 tails), we might seriously begin to doubt our null hypothesis in favour of suspecting the coin was weighted, since the probability of it happening is only 2 in 10,000.

The binomial distribution is great for simple series of win/lose betting propositions where we stake the same unit size on the same betting price every time. In the real world, however, bettors bet on all sorts of different prices with all sorts of different stakes. Furthermore, not all wagers involve straight win/lose scenarios. Full-ball and quarter-ball Asian handicaps and push results for American point spreads complicate matters. In such circumstances, we can rely on what is known as the t-distribution, and the student's t-test for statistical significance which uses it. Unlike the binomial distribution which is based on discrete values used to calculate the probability of x successes from y independent yes/no trials, the t-distribution is continuous. Of course in practice betting profits and losses will be discrete since there can only be a finite number of possibilities for any series of wagers. Nevertheless, that number is large after only a few of them and their distribution can meaningfully be treated as continuous. The t-distribution is very similar to de Moivre's normal distribution (the famous bell-shaped curve) we examined in the earlier chapter 'Cleopatra's Nose',

showing symmetry (good and bad luck should be balanced) and the presence of a single peak which coincides with the mean, μ . For numbers of trials above about 30, the t-distribution to all intents and purposes is the same as the normal distribution. Either can be used as an approximation of the binomial distribution where the probability of outcome is neither close to 0 nor 1 and where the number of trials is sufficiently large (usually only 10 to 15 will be enough). We would choose to use the normal when we know what the variance or standard deviation, σ , is of general population. When analysing the significance of a profit from a sample of betting outcomes, arguably we cannot assume that we do, so the t-distribution is preferred. Provided the means of possible samples of profits and losses selected are normally distributed we are safe to use it.

There is a variety of different t-tests used in statistical testing, but the one we are particularly interested in here is the one-sample t-test. This statistical technique tests the null hypothesis that the population and sample means are the same. If they are found to differ by a statistically significant amount, the null hypothesis is rejected. In the context of analysing a betting record, the sample is simply the series of profits and losses realised from bets settled. The population, meanwhile, represents the complete set of all theoretically possible profits and losses that could be realised from placing those bets, a very large number even for a small number of bets. The bettor's sample of profits and losses from his series of bets represents just one possible permutation. The null hypothesis, then, is that his return, or more specifically his average profit per bet, is not significantly different from the mathematical expectation as defined by the whole population of possible profit/loss permutations.

Calculating the average profit per bet from a series of wagers is easy; we just look at our record of profitability. Standardised to an average of one unit stake per bet, it is simply the same as our betting yield. If I've returned \$110 from \$100 bet (yield = 10%, return on investment = 110%) my average profit, expressed as a decimal, is 0.1. The population average, meanwhile, just represents the profit expectation from betting randomly. In the absence of any further information about true result probabilities (remember those are unknown unknowns), it is perhaps sensible to assume that the bookmakers' prices most typically represent the best possible estimate of their probabilities. (For brands like Pinnacle Sports which

respect efficient market principles this is surely a reasonable assumption; for others like Marathonbet which don't, it might not be.) All we then need to know is how much advantage is built into those prices. If a bookmaker was offering 1.95 for fair result probabilities of 50%, the population mean for such bets would be -0.025. In this example, our loss expectation would be -2.5 cents for every \$1 wagered. Of course, we now know that the favourite–longshot bias in many sports markets will complicate this picture. Bettors who prefer favourites will conceivably have a smaller *a priori* loss expectation than those who prefer longshots. Furthermore, if a bettor makes efforts to hunt best market prices, his loss expectation can be reduced even further, often to 0 and sometimes can even be turned positive. Indeed the average overround for best home-draw-away market prices in domestic European football is of the order of 101% whilst for tennis match odds it is typically below 100%. Of course, as I explained earlier, consistently taking those prices with brands that have a tendency to frown upon such activity is another problem in its own right.

The t-test, then, simply compares the bettor's observed return to a theoretical expectation defined by the market he's betting in, and analyses whether the difference is statistically significant. If it is, we may reject the null hypothesis, which assumes that observed profits and losses are purely the result of chance, in favour of a different one. The t-test does not tell us what that new hypothesis would be, but the underlying presumption is that any statistically significant return by the bettor would be the result of skill. The t-statistic or t-score is really just a measure of the departure of the observed sample mean away from expected population mean, and is defined by the following equation:

$$t = \frac{\sqrt{n} (\bar{x} - \mu)}{s}$$

where \bar{X} (pronounced x-bar) is the sample mean (the average profit per bet), μ is the population mean (the random profit expectation), s is the sample standard deviation in profits and losses of the betting history (note not σ , the standard deviation of the population which remains unknown) and n is the sample size (the number of bets). The t-score is proportional both to the square root of the number of bets and average profit per bet, and inversely proportional to the sample standard deviation. The first two are intuitively obvious: a yield of 20% from 100 bets is less likely to be the result of chance than a yield half the size from the same number of bets. Similarly, a yield of 10% from 1,000 bets will be a far more reliable indicator of consistent forecasting ability than the same performance from just 100 bets. Less intuitive is the influence of the betting odds. In fact, a 20% yield from betting odds around 1.50 will be a much better indicator of skill than an equivalent yield from betting odds around 3.00. The reason is because the standard deviation in profits and losses is nearly two and a half times larger for the latter than the former, meaning the t-score will be correspondingly that much smaller. Betting on lower probability outcomes (longer odds) is inherently riskier (assuming equivalent stakes) because it is more at the mercy of random variability. To put it another way, outcomes are more volatile.

You can think of it this way. Imagine a 20:80 betting proposition with fair odds of 5.00 for one side and 1.25 for the other side respectively. 20 winners for the former and 80 winners for the latter will both break-even. The chances of seeing one extra winner either way, i.e. 19-81 or 21-79, are roughly similar. Yet the outcomes for each side are completely different. 19 or 21 winners at odds of 5.00 will give yields of +5% or -5% respectively. In contrast, 79 or 81 winners at odds of 1.25 will show -1.25% or 1.25% profit over turnover. Betting longer odds implies taking more risk to get more reward as this little thought experiment demonstrates, although really it's intuitively obvious anyway. Achieve that extra winner and you'll get a bigger bang for your buck. On the other hand, suffer that extra loser and face the consequences. Essentially, the influence of luck has a bigger effect on outcomes the longer your odds are.

The following approximation for s shows how it is related to the size of the betting odds.

$$s = \sqrt{r(o - r)}$$

Recall that r is the decimal return on investment (for example: yield = 10%, ROI = 110%, $r = 1.1$), meaning that $\bar{x} = r - 1$, whilst o represents the average decimal betting odds. Plugging that equation for the standard deviation of profits and losses for a betting history into the one above for the t-score, we get:

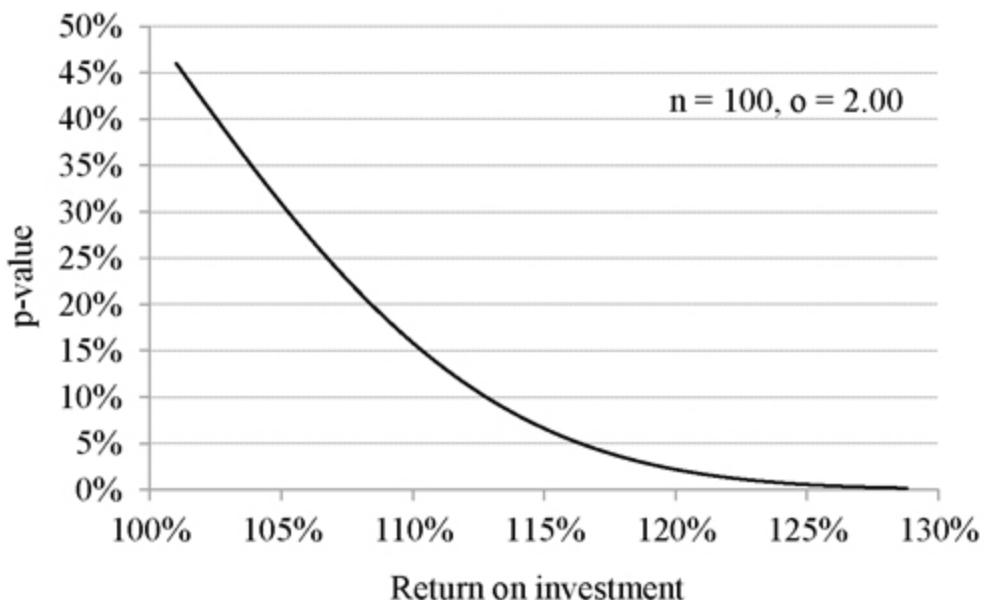
$$t = \frac{\sqrt{n} (r - 1 - \mu)}{\sqrt{r(o - r)}}$$

This equation for the t-score is robust for betting histories that are based on level staking (the same stake for every bet), even when the range of betting odds is quite large. For negative returns, the t-score will be correspondingly negative. The corollary is that superior yields achieved through betting longer odds, as is typical in markets like horse racing, are not necessarily a sign of better forecasting talent. As in the example above, the same amount of luck will deliver much bigger percentage returns. Hence, comparisons of tipping services that rank only by yield are fundamentally misleading. In effect, by involving the odds the t-score provides a measure of the quality of the risk-adjusted return in excess of the benchmark (μ), which is exactly what we want to determine.

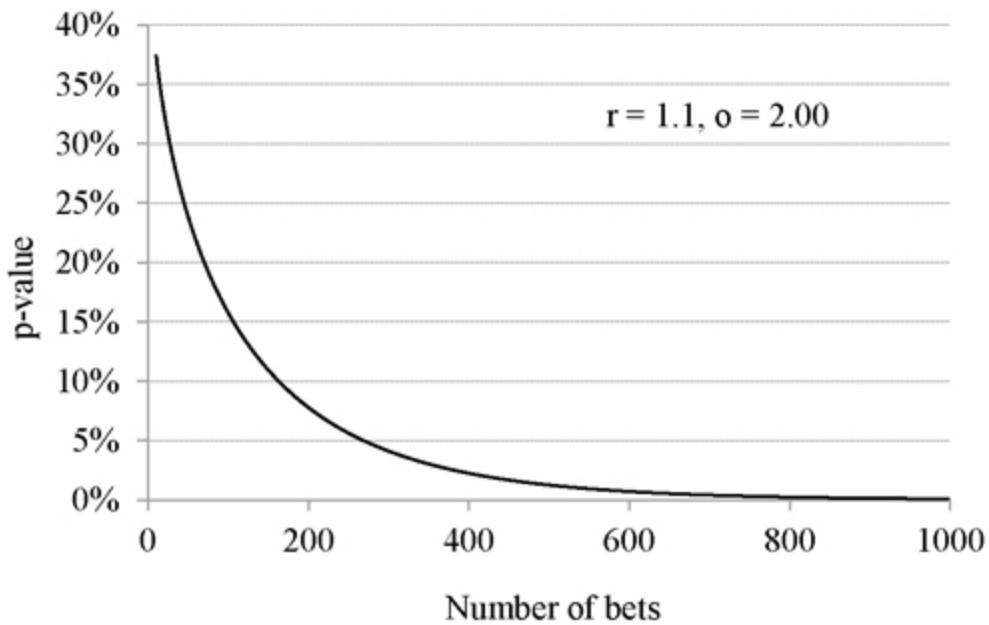
Let's try plugging some values into the t-score equation. Suppose I've placed 100 wagers with average odds of 2.00, and made a return on investment of 120% ($r = 1.2$). I've made an effort to find the best market prices so let's assume that my profit expectation, μ , is 0, that is to say, break-even if I just picked things to bet on randomly. My t-score for this betting history is 2.04, with a corresponding p-value of 0.022 (or 2.2%), the probability that such a score could arise by chance. Those with Microsoft Excel can do this using the TDIST (t , degrees of freedom, tails) function⁷⁹. The degrees of freedom (DF) merely describes the number of independent pieces of data used to make the statistical calculation, and is a measure of how certain we are that sample is representative of the entire population. It is equal to $n - 1$. For 100 wagers $DF = 99$. The tails argument can be either 1 (for the one-tailed t-test) or 2 (for the 2-tailed t-test). Since we're really

only interested in whether a profit is statistically significant, I have used the former. In this example, we calculate the p-value with =TDIST(2.04,99,1). There are also many online ‘p-value from t-score’ calculators that will do the same job⁸⁰. If our chosen α -value (the significance level) is 5%, we would be able to reject the null hypothesis that our profit had arisen purely as a result of luck. On the other hand, if it was 1% or smaller, we would still not be able to rule this out.

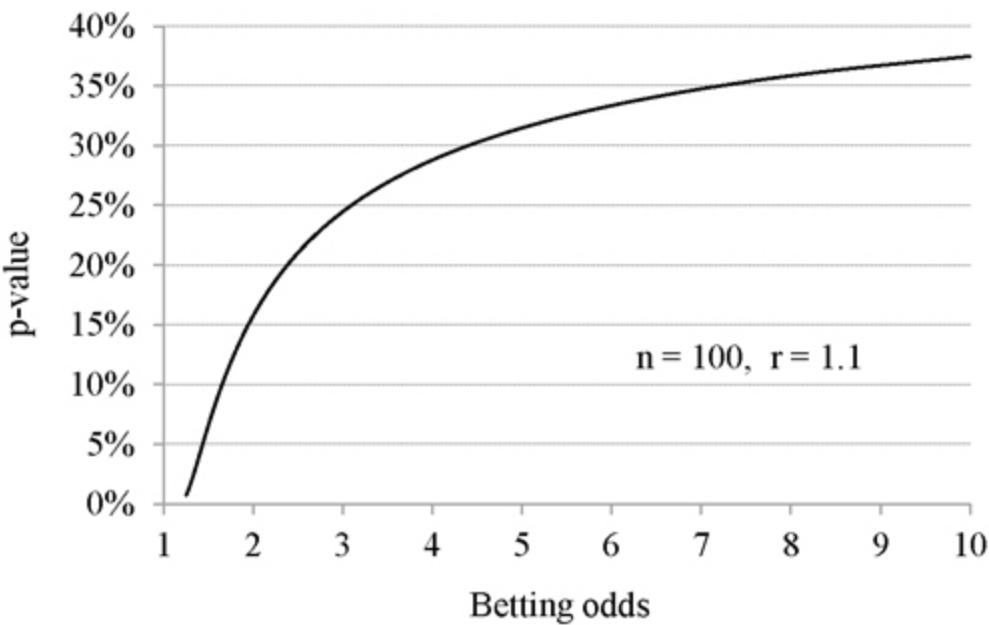
The next three charts illustrate how p-value changes with varying returns (r), bets (n) and odds (o), assuming μ is 0. In the first instance, the number of bets is fixed at 100, whilst the average odds are fixed at 2.00; p-value falls with increasing size of returns.



Second, the return is fixed at 110% with the average odds again fixed at 2.00. The p-value diminishes with increasing number of bets.



Finally, the number of bets is once again fixed at 100, with the returns also 110%. The p-value falls as the odds shorten.



One of the drawbacks of this one-sample t-test method in trying to estimate the likelihood of skill being involved is that the population mean, μ , is fixed and equivalent to the mathematical expectation. There is no reason, however, to assume that, during the time period from which the

sample of bets is taken this should be the case. The betting market won't always provide a perfectly accurate measure of mathematical expectation. Sporting outcomes are subject to random variation. Sometimes more outcomes will happen than is implied by the betting odds, sometimes fewer. In fact, leaving aside the favourite-longshot bias, over the long term the market does an incredibly good job of replicating intrinsic probabilities. In the short term, however, there will be some fluctuation about the average.

Conceivably, the best way to circumnavigate this problem is to compare a bettor's odds to the closing prices. Why? As I'll explain later in the book, the more people who express an opinion about an outcome probability, the better the chance that it will be accurate. It's called wisdom of the crowd. By definition, closing prices are representative of the largest number of opinions, and are therefore the best approximation of the intrinsic outcome probability. Odds shorten when more people back a proposition than those laying it or backing alternative propositions. Conversely, they lengthen when the opposite is true. This is Bayesian statistics at its finest. When odds shorten this implies that proportionally more people now believe the outcome has a greater probability than was previously thought. If a bettor consistently bets prices that are subsequently shorter by the time the market closes, this implies that he is beating the market, and the bookmakers are reacting to his activity. The key word is 'consistently'. Sometimes a bettor will beat the closing price, sometimes he won't, but if he does it more often than not, and if his average advantage is greater than the size of the bookmaker's margin, this will be a good sign that he has found positive expectation. The benefit of this method lies in its ability to look for skill even in bettors who have been unlucky and unprofitable in the short term. Even where losses have been made, if the bettor is consistently beating the closing odds by more than the bookmaker's margin, that's a good sign that he is sharp and will prove to be profitable in the long term. Indeed, it provides the explanation for why bookmakers will restrict or close betting accounts that aren't even winning.

By the same token, however, profitable tipsters can easily be shown to be winning through luck. One tennis tipster I verified, for example, showed a level stakes yield of 6.5% from 1,606 wagers and a p-value of 0.04, statistically significant at an α -level of 5%. Yet based on a comparison of the odds he tipped versus their closing prices, we would have to conclude

that he had simply just been lucky. In all he managed to beat the closing odds just 43% of the time (someone guessing should manage 50%), and 52% of the time for his winning bets only. According to the closing odds and the bookmaker's margin, his average expectancy was -2.8%; for his winning bets only, marginally better at -0.7%. Of course, all this could mean that the tipster was neither actually backing his own tips nor had any customers who were betting them either. Consequently, there would be nothing for a bookmaker to react to. Nevertheless, to his carefully chosen Kelly staking approach he had been regressing to the mean for many months and abandoned his project at the end of 2014. Perhaps he had concluded the very same as this analysis is telling us.

With this story in mind, it's time now to consider more generally the evidence for whether markets of psychology like betting, poker and investing – so-called speculative gambling – involve games of luck or games of skill. As I've mentioned already, readers may find what follows deeply discouraging. Indeed, it is precisely because I feel that betting (and for that matter financial investment) is almost entirely a game of luck that I have purposely used the word 'gambling' in the subtitle of this book. Most players in these games might wish to describe themselves as investors but their outcomes imply that they're doing nothing of the sort.

⁷¹ Stewart, I. *Game, Set and Math: Enigmas and Conundrums* (1989), Basil Blackwell, London.

⁷² Gould, S. J., 2004. *Triumph and Tragedy in Mudville: A Lifelong Passion for Baseball*, New York : W. W. Norton & Company.

⁷³ Meyer G., von Meduna M., Brosowski T. & Hayer T., 2013. Is poker a game of skill or chance? A quasi-experimental study. *Journal of Gambling Studies*, **29(3)**, pp.535-550.

⁷⁴ Bowling, M., Burch, N., Johanson, M. and Tammelin, O., 2015. Heads-up limit hold'em poker is solved. *Science*, **347**, pp.145-149.

⁷⁵ Fiedler, I. C. & Rock, J-P., 2009. Quantifying skill in games: theory and empirical evidence for poker. *Gaming Law Review and Economics*, **13(1)**, pp.50-57.

⁷⁶ Bernstein, P. L., 1998. Where, Oh Where Are the .400 Hitters of Yesteryear? *Financial Analysts Journal*, **54(6)**, pp.6-14.

⁷⁷ Any standard statistical textbook or website will show you the equation for the standard deviation, depending on the distribution of data, and how to calculate it. Software packages like Microsoft Excel will automatically do it for you.

[78](#) LaBrie, R. A., Kaplan, S. A., LaPlante, D. A., Nelson, S. E., & Shaffer, H. J., 2008. Inside the virtual casino: a prospective longitudinal study of actual Internet casino gambling. *European Journal of Public Health*, **23**(4), pp.410–416.

[79](#) For a wider discussion of how to use the Excel TDIST function readers can refer to <http://www.excelfunctions.net/Excel-Tdist-Function.xhtml>. I've also made available a little Excel calculator at http://www.football-data.co.uk/blog/P-value_calculator.xlsx that will do the job

[80](#) For example, see <http://www.socscistatistics.com/pvalues/tdistribution.aspx>.

MONKEYS THROWING DARTS

In addition to his ruminations on luck, Samuel Goldwyn is believed to have sardonically remarked: “*forecasts are difficult to make, particularly those about the future.*” How prophetic he was. Gambling – the toss of a coin, roll of a dice, spin of a wheel, draw of a card, call or raise of a hand, scoring of a goal, rise or fall of a share price – is fundamentally about making forecasts about the future in environments of risk or uncertainty. Thus far, we’ve talked a lot about why we’re drawn to speculating on unknowns; we’ve talked about how we value them and how we rationalise them; we’ve even talked about why we want to control them. Now let’s turn our attention to whether we are actually any good at predicting them. Is the forecasting of unknowns for financial gain a game of skill or mostly a game of chance? Evidently, for pure gambling in casinos or on the lottery, the answer is fairly obvious: it’s a loser’s game, at least in terms of mathematical expectation, although our Palaeolithic pattern interpreter and belief engine manage to fool some of us into believing the contrary. Yet, for almost all of us speculatively gambling in markets of psychology – betting, poker, trading and investing – the evidence suggests we might just as well be monkeys throwing darts here too. Trying to make predictions based on where they land, as we will see, is just as good as anything else.

The science of prospect theory has revealed that, when trying to make forecasts about uncertain outcomes, we are prone to committing systematic errors due to inherent irrational biases in our decision making. In theory, this should offer profitable expectation for those knowledgeable enough to exploit such errors. In practice, it turns out that very few of us can do it consistently. Much of the time these markets of psychology resemble little more than random walks, mostly rational and efficient and constructed from a collective wisdom of crowds, even when individuals within them are not acting fully rationally, offering few opportunities to outperform them. Market inefficiencies – errors where valuations about future outcomes fail to properly reflect the ‘true’ probabilities of those outcomes – do exist, but

usually not for very long or not in any consistently predictable pattern. An old economist's joke explains this nicely:

An economist and his friend are walking down the street when they come upon a \$100 bill lying on the ground. As the companion reaches down to pick it up, the economist says, "Don't bother, if it were a real \$100 bill, someone would have already picked it up."

Not only are we rather poor at prediction, we are also rather reluctant to accept that we are so. Both failures to a significant degree arise from our own blind spot bias, an inability to recognise that we suffer from the same cognitive distortions that plague other people, most critically our inability to differentiate chance and causality, luck and skill.

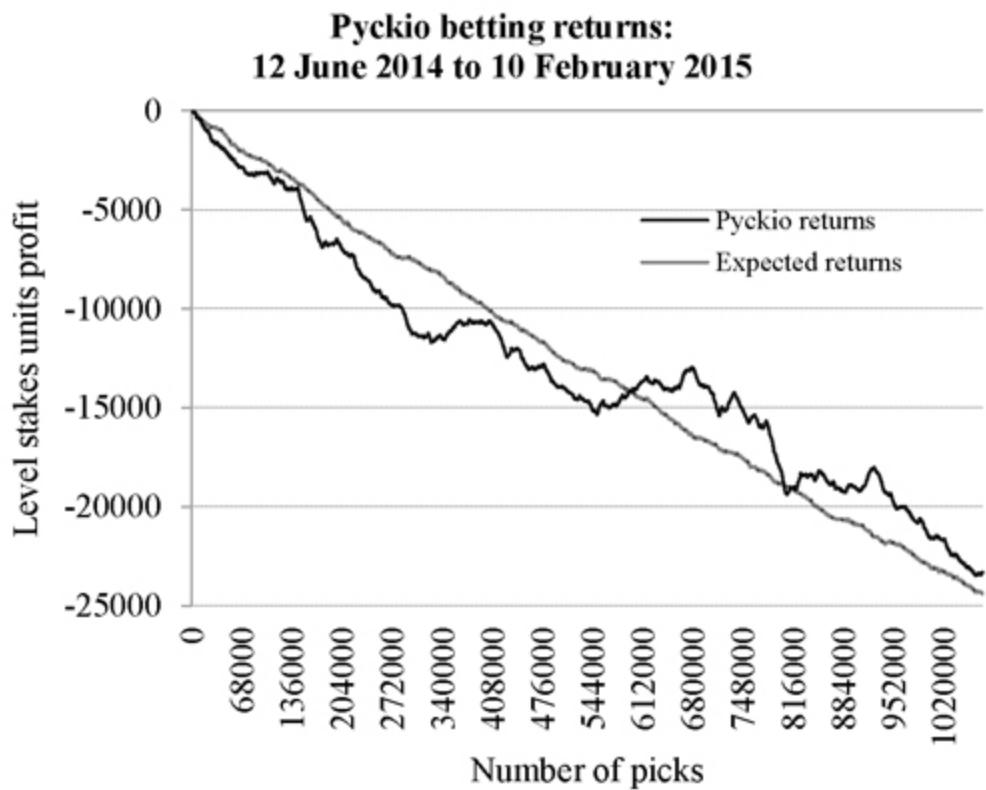
Is Betting Completely Random?

What sort of performances do bettors typically achieve? Perhaps more importantly, will the distribution of their performances lend credence to the idea that betting is more than just a game of chance, as so many who take part in it truly believe? Specifically, if there are more bettors achieving better returns than we would typically expect to occur by chance, this might very well suggest that some of them at least were skilled. Let's take a look at some real world data.

On 12 June 2014, a new international sports betting community called Pyckio.com opened, allowing bettors to both share and follow tips for a variety of sports. I should declare now that I am a 5% shareholder. By 10 February 2015, a total of 1,073,029 picks (excluding unsettled bets) from 6,044 different tipsters had been posted, with average odds of 1.99, all of them taken from the bookmaker Pinnacle Sports, which as previously mentioned has a reputation for high liquidity and an acceptance of winning punters. From the beginning, Pyckio's founders took the decision to ignore other bookmakers known for price manipulation and an intolerance of customers exploiting those outliers. Indeed, what's the point in tipping prices which will ultimately prove to be unbackable in the long run for the majority of bettors? Markets on which a player could tip were restricted to the following sports: American football, Australian football, baseball, basketball, darts, e-sports, handball, hockey, mixed martial arts, rugby,

snooker, soccer, table tennis, tennis and volleyball. Players were permitted to stake in the range 1 to 10 units, but Pyckio also analyses performance for every tipster to level stakes.

The chart below tracks the evolution of betting returns for all 1,073,029 picks analysed to level stakes. This is the most powerful and accurate method of determining what sort of value expectation a bettor, or collection of bettors, actually has. Different money management has no influence on the long term value expectation, merely the nature of risks over shorter time scales⁸¹. Furthermore, a handful of tipsters had attempted to manipulate their performance by using artificial staking methods that would be of little practical benefit to customers in the real world. Level stakes allows us to see behind this manipulation.



The aggregated yield for the full record was -2.17%. You will notice immediately that this is pretty close to the typical book margin employed by Pinnacle Sports for many of its markets. A bettor's loss expectancy can be calculated by means of the following equation:

$$EV = \frac{1}{M} - 1$$

EV is the expected value whilst M is the bookmaker's decimal margin as calculated by his overround. Aggregated losses of -2.17% on turnover imply an average margin for those wagers of 1.022 (or 2.2%).

The time series also includes the evolution of returns that we would expect to witness simply by chance. To model a series of random betting outcomes we need to be able to estimate their 'true' probabilities. For unknown unknowns like betting outcomes, arguably the best way to do this is via the odds themselves; we just need to remove the influence of the bookmaker's margin. For Pinnacle Sports I have assumed this to be 1.025 (or 2.5%). Whilst markets for minor sports have margins higher than this, the majority of picks shared through Pyckio's community have been for football, tennis and US sports, all with margins in the region of about 2 to 3%. I also considered the influence of the favourite–longshot bias using the odds setting model discussed in the previous chapter, where the differential margin weights to shorter and longer prices respectively are inversely proportional to the outcome probabilities defined by the odds. That is to say, shorter/longer odds attracted smaller/larger margins relative to the average. With 'true' probabilities estimated for every pick in the history it was then a simple matter of applying a random number generator to determine whether a bet won or lost.

The yield from randomly settled outcomes was -2.27%, almost exactly the same as for the actual yield. There was no statistically significant difference between the two profit/loss samples. Of course, the yield from random outcomes was based on just one possible sample of random outcomes; re-run the random number generator and we'll get a different answer. The full population of answers will very probably be normally distributed about an average that can be calculated from the 'true' probabilities. This turns out to be -2.36%, meaning our random sample performed a little bit better than expectation.

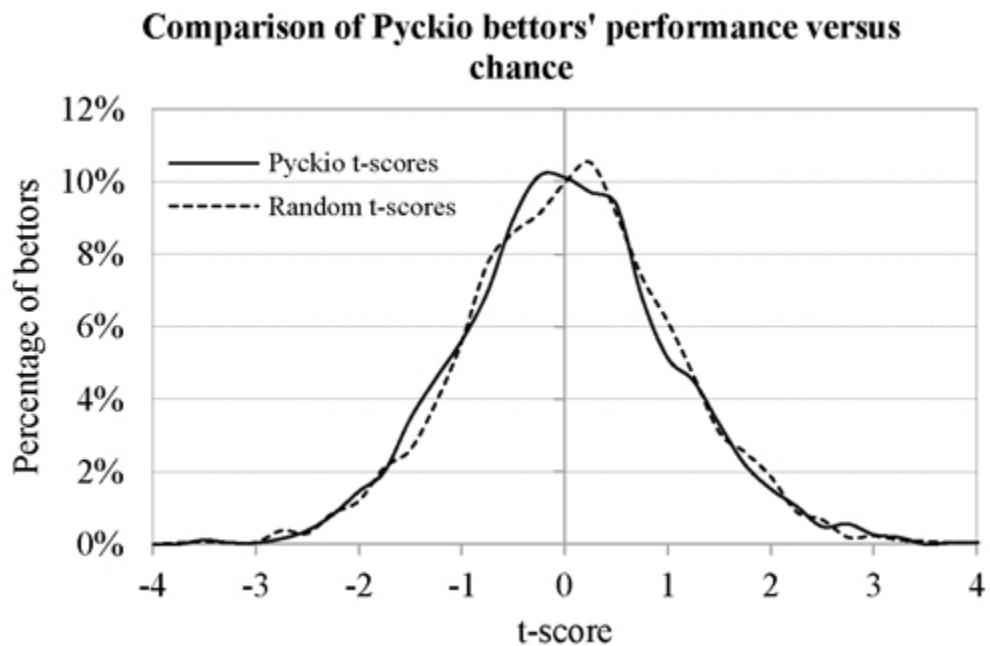
Two things are striking about the comparison between the actual and random time series: firstly, the longer term trends are almost identical; secondly there is more short term variability in real betting returns than in

the randomly generated ones. One explanation for the latter might be found in the short term variance between different betting markets. *A priori* one would expect differences between different sports and different betting markets. If a disproportionate number of community members were concentrating on certain markets compared to others and at different times, this would likely increase the variability seen in the evolution of the returns. Another explanation might be that the random sample assumes the same profit margin for every book (1.025), whereas Pinnacle Sports' actual margins will show greater variance across its markets. Finally, the greater variability might conceivably be evidence of short term inefficiencies in betting prices, induced by less than rational decision making that all of us are capable of. Evidently, however, such inefficiencies, which unmistakably yield short periods of profitability, are not consistent and ostensibly unpredictable. Economists who support the efficient market hypothesis (something I'll examine in more detail in the next chapter) frequently argue that most short term inefficiency – the technical term for ‘mistakes’ – soon disappears once it has been discovered. Remember, the more profitable a discoverable pattern is, the less likely it is to survive; just like the \$100 bill lying on the ground.

How have individual bettors performed? 2,138, or 35.4%, of the 6,044 tipsters made some sort of profit (compared to 2,513 or 41.6% for the random model). That's still a lot of people making money; aren't we always told that 95% of people playing such games are losers? In fact we're making the wrong enquiry here. We must remember that anyone can make a profit simply by luck. The key question is whether the distribution of profits that the bettors have experienced differs significantly from a distribution that could be predicted by chance. If all we have is luck, that will eventually run out and money we might have already made may then be handed back as losses, as the evolution of good and bad luck combined regresses to the mean.

In an attempt to create a risk-adjusted assessment of every bettor's performance, I have calculated their t-scores according to the methodology previously described. This allows us to compare bettors with different risk preferences (specifically what odds they prefer to bet at) and the longevity of their betting histories. One weakness in this approach is that different lengths of betting histories essentially belong to different t-distributions

each defined by the number of degrees of freedom they have (recall $DF = n - 1$). As previously noted, however, for distributions with at least 30 independent observations, the t-distribution is a reliable approximation of the normal one, and a t-distribution for 30 bets is not radically different from that for a history of 3,000 bets. Consequently, the chart below plots the distribution of t-scores for the 2,690 bettors who had histories of 30 bets or more (accounting for 44.5% of all bettors and 97.4% of all bets). Removing the 3,354 tipsters with shorter histories slightly improves the aggregated returns to -2.02%. The chart also shows the distribution of t-scores predictable *a priori* from chance, based on the odds and assumed 'true' probabilities as described above. Spot the difference. Statistically there isn't one (2-tailed paired t-test, p -value = 0.21). The average t-score of these 2,690 bettors was 0.03; that compares to 0.06 for the randomly generated equivalents. Here, a t-score of 0 implies a negative return matching that predicted by the bookmaker's margin. On average, then, bettors are simply replicating what the market says should happen.

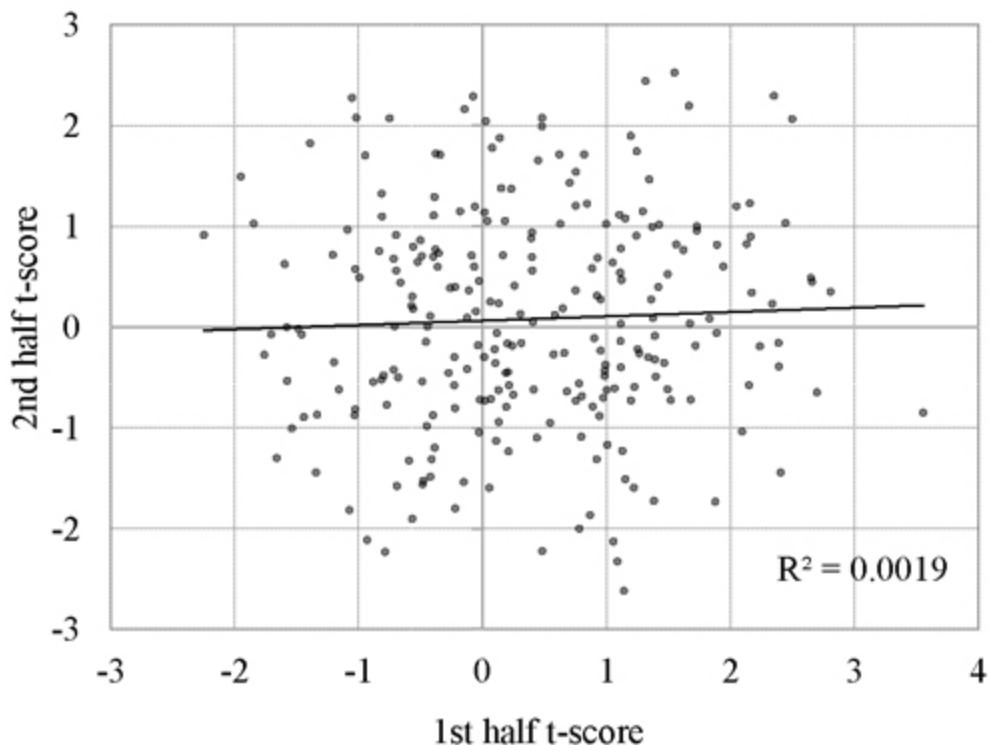


In the previous chapter, I reviewed how luck and skill can be visually distinguished by comparing the distribution of observed outcomes to those predicted by chance alone. If there is no difference, the implication is blindingly obvious: these bettors, in aggregate, have simply replicated

chance. There is barely even the slightest hint that bettors in the profitable tail of the distribution are contributing anything over and above luck. The correlation between the two distributions is almost perfect with 98.75% of the variability in the actual outcomes explained by the variability produced by chance⁸². 69.6% of t-scores in the distribution of actual performance fall within plus or minus one standard deviation (1.035) of the average, and 95.2% within two standard deviations. Recall for normal distributions these figures are 68.3% and 95.5% respectively. The standard deviation in t-scores for the random distribution – 1.027 – is almost the same. According to true-score theory this means that luck is accounting for 99.2% of the variation in observed risk-adjusted betting performances, more or less the same figure as that calculated from correlation.

One useful measure of skill is its persistence. If outcomes are predominantly a matter of skill you can expect to be able to repeat your performance. Earlier we saw that skill accounted for perhaps 80% of the rankings in Premiership football. With such a high contribution we would expect rankings to persist from season to season, and indeed that is more or less what we see. The top 4 teams in seasons 2004/05 through to 2014/15 have been occupied by just 7 teams. It's a similar story in professional tennis, arguably even more dependent on skill. By contrast, where luck is predominant, scores will more quickly revert to the mean, and there will be little correlation between one outcome and the next. Statisticians know persistence as reliability, which can be measured by means of the correlation coefficient, r . Values of r vary between 0 (no correlation) and 1 (perfect correlation). The chart below shows the correlation between first and second half performance (divided into equal-sized samples) for the 249 Pyckio bettors with histories of 1,000 picks or greater, as measured by the t-score. Together they accounted for 509,779 (or 47.5%) of all bets, with an aggregated yield of -1.91% and an average t-score of 0.31, both marginally better than the performance for the whole population. I'll be considering the reasons why this might happen a little later in the chapter. Unfortunately those readers who might be interpreting this as evidence of skill will be sorely disappointed.

Correlation between 1st and 2nd half t-scores



There is almost no correlation at all. The value of r is 0.043, whilst R^2 is 0.019, with an almost flat regression (or trend) line. R^2 is a measure of the amount of variance in the data. In this case just 1.9% of the variability in the second half t-scores can be explained by the variability in the first half t-scores. To all intents and purposes there is almost no persistence in performance whatsoever, with t-scores simply regressing to the mean. If a bettor has performed well/poorly in the first half it is more probable that he'll perform worse/better in the second half, simply by virtue of the fact that his first half t-score represented more of a positive/negative outlier. The 26 first half histories that had t-scores of -1 or poorer (average -1.41) regressed to an average t-score of 0.06 in the second half. Similarly, the 76 first half histories with t-scores better than 1 (average 1.64) regressed to an average second half t-score of 0.13.

So there we have it: bettors, at least the overwhelming majority of these bettors at any rate, would seem to be simply throwing darts. There may well be some capable of methods that yield more consistent and predictable returns, but they would appear to be relatively small in number and largely

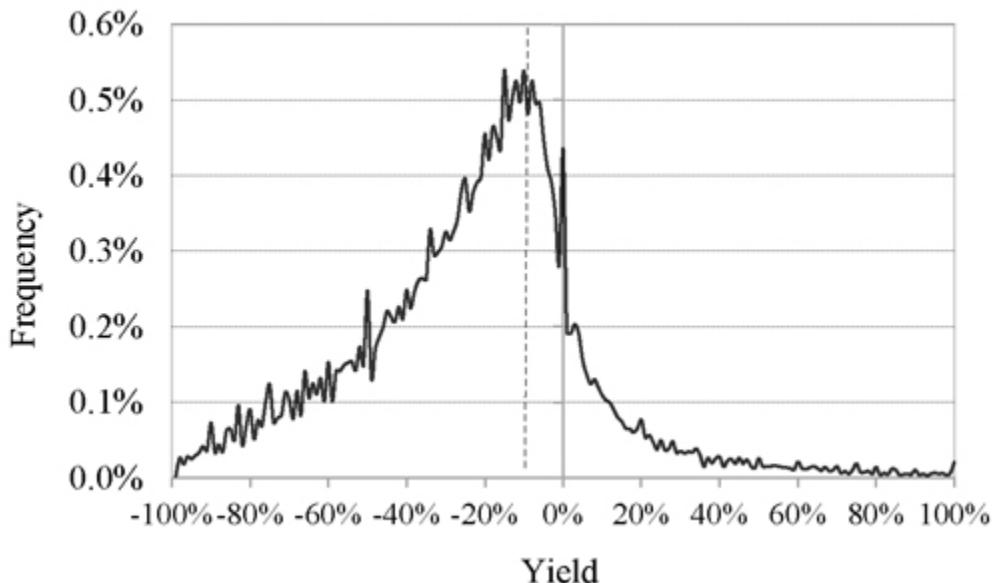
indistinguishable from the rest. Just 45 bettors with histories of at least 30 bets managed to achieve a performance level where their p-value was less than 1%. My simulation of chance produced 38 of them. The difference is not statistically significant. Theoretically, if 6,044 bettors were to continue randomly betting indefinitely, we'd expect about 60 to attain a p-value of less than 1%. The best performance from a bettor with a long record was a t-score of 4.06 (based on 701 tips, a yield of 12.6% and average odds of 1.99), equivalent to a 1 in 37,000 probability that this could have happened just by luck. But the random number generator delivered a history almost as impressive ($t = 4.05$, bets = 171, yield = 24.1%, average odds = 1.85) that had a 1 in 25,500 probability; food for thought indeed! Clearly, with enough people playing, pretty much anything is possible, good or bad, just by chance. Fortunately, it's free to follow almost all the tipsters at Pyckio. Only those who have demonstrated a significant and consistent profitability are eligible to charge a subscription, and should performance drop, that is revoked.

I'm often criticised by those running advisory networks where many tipsters offer their advice for sale that it's meaningless to look at their performances as a group like this. One in particular was Top-Tipster.com, a collection of about 300 bettors whose tips customers can register to purchase. When pressed on whether they were just throwing darts, Top-Tipster openly acknowledged that not every tipster will make a profit. Indeed, like those in the Pyckio community, only about a third of them were doing so. However, the idea behind Top-Tipster's model is to showcase a selection and let customers purchase the best ones. By treating them separately we can find the ones that are skilled. If only that was so. Unfortunately, this completely misses the point. It's precisely because in aggregate their performance almost perfectly replicates what could be predicted by chance alone that it's unlikely that more than just a small number is doing anything else. If there were substantially more we would see them. The distribution of performances would then depart markedly from normality. For Top-Tipster, the distribution of its t-scores (estimated from average betting odds, yield and number of bets issued for each tipster) did not show such departure, with 70%/96% of the scores between plus or minus one/two standard deviation of the average and the same bell-shaped curve the distribution of Pyckio t-scores exhibited. In aggregate, the tipsters

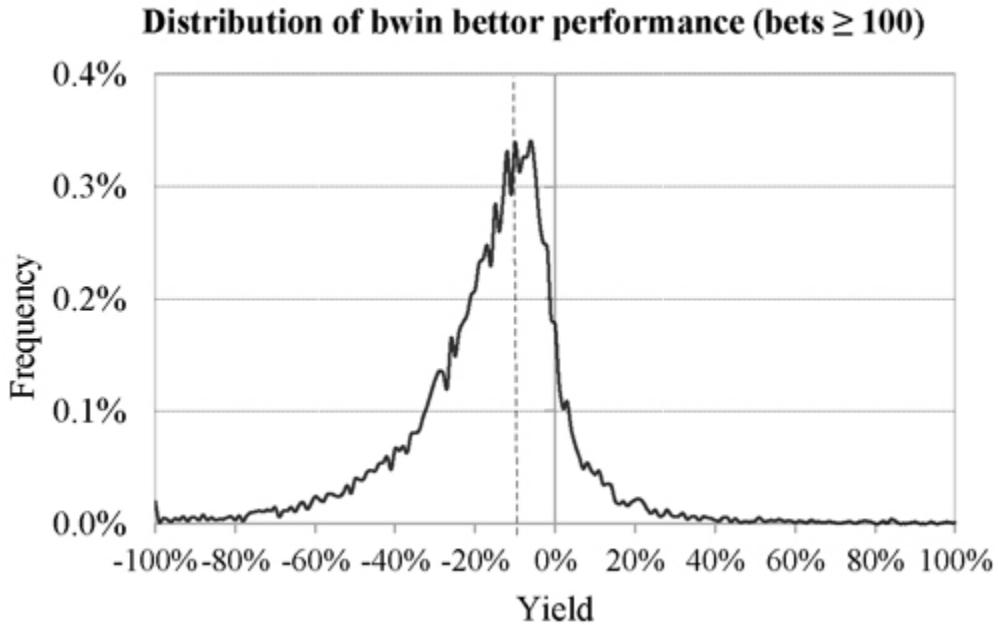
were responsible for a loss of about 2% on turnover. Top-Tipster claims to have many consistently profitable tipsters to choose from. Unfortunately, consistency was not something I have been afforded the privilege of testing since Top-Tipster wouldn't release the full historical tips data.

An even larger data set of betting activity was made available in 2005 by the bookmaker bwin.com for the Harvard Medical School team⁸³ whose research into casino gambling performance I reported a little earlier. It included 40,499 sports bettors who registered between 1 and 27 February 2005. Of those, 39,719 went on to bet on pre-match markets between registration and 30 September 2005, whilst 24,794 indulged in some live action. In total over the 8-month period, these customers made 7,815,702 bets (68.5% pre-match, 31.5% live play), wagering a total of €61,656,383 (47.0% pre-match, 53.0% live play). Unlike the Pyckio data set which contains advised tips, the bwin data is based on actual play. However, unfortunately, it does not include betting prices, making it impossible to calculate risk-adjusted performance scores for each customer. Consequently, the chart below illustrates the distribution of yield performance across the group, arguably less reliable than t-scores and evidently more prone to asymmetry. Plainly, it's not possible to have yields lower than 100%; in the right hand side of the tail, however, their size is limited only by the maximum odds that bwin makes available. For visual clarity, I have restricted the horizontal axis to a maximum of 100%. In fact, in this data set there were 454 players (just over 1% of the total) who had yields of greater than 100%, the largest being 1,527%, although obviously many of them achieved this with just a handful of high-priced wagers (nearly half had fewer than 10).

Distribution of bwin bettor performance



A number of interesting observations can be made. Firstly, the unmistakable presence of the bell-shaped curve (albeit negatively skewed), implying a significant underlying randomness, can be seen. Secondly, superimposed on it is a number of spikes, in particular at yields of 0% (break-even) and -50%, conceivably representing a disproportionate number of Asian handicap pushes (stakes returned) and half-losses for short betting histories. An additional spike at -100% is also not shown. This missing data point had to be removed from the chart for the purposes of visually clarity. Accounting for 4.21% of the observations, it would have obviously hugely distorted the remainder of distribution. Unsurprisingly, the vast majority of these customers were those who bet only a small number of times, losing all of them. Indeed 24% bet only once, whilst 83% only bet 10 times or fewer. Amazingly, however, a small handful managed to place several hundred bets (the largest being 963) despite managing not a single success. Quite what this must do for the bettor's psychology is anyone's guess. Presumably, these bets were struck at very large odds (perhaps by constructing large accumulators or permutations) that had almost no chance of success. Including only those customers with histories of at least 100 bets (34% of the total) smoothes the shape of the distribution but its general pattern remains largely the same.



Thirdly, it is apparent that the vast majority of bettors lost money over the period of analysis. Aggregated losses amounted to €5,943,808 or -9.64% on total stakes bet (shown as the dotted line on the chart), with only 5,444, or 13.4%, returning a profit. This compares to 35.4% for the Pyckio data set. Of course, there is an obvious explanation for such a difference. Whilst in 2014/15 Pinnacle Sports' margin was of the order of 2 to 3%, in 2005 bwin's typical margin was much greater, for example 10% for pre-match home-draw-away football odds. In fact, aggregated losses on pre-match bets were considerably larger (-13.25%) compared to those for live-play (-6.45%). Assuming that a 10% margin is typical across all of bwin's pre-match markets, one might wonder why such a large proportion of customers (73%) underperformed this mathematical expectation. The shape of the bell-shaped curve, and the fourth observation, provides the clue.

In contrast to the Pyckio t-scores which are almost symmetric about the mathematical expectation, the distribution of bwin yields is quite asymmetric and negatively skewed. Why should this be? One obvious suggestion is that yields, in contrast to t-scores, are neither adjusted for risk (specifically preferences for shorter or longer betting odds) nor the longevity of betting history. Shorter histories and those with more longshot betting will show a greater variance in betting yield than longer histories and those focused more on favourites. Arguably, unlucky customers who

experience earlier losses are more likely to quit than lucky customers who will continue to play. Indeed the original research team reported a negative correlation between customer yield and total money wagered, consistent with the idea that more successful betting, and specifically winning, encourages continued play, whilst conversely losing discourages ongoing play. Consequently, lucky players have more opportunity to regress towards the mean, ensuring disproportionately fewer of them will show yields significantly above (or indeed below) the arithmetic expectation. Unlucky losers, by contrast, will have placed fewer bets on average, leading to a greater variation in yields below expectation. Sure enough, the average and median number of bets placed by those customers showing yields lower than the aggregated average (-9.64%) were 167 and 39 respectively. This compares to 264 and 91 respectively for customers who performed better than the aggregated average. It seems unlikely, however, that this influence could be operating in isolation. You can see above that the asymmetry in the distribution of customers' yields persists for those who've made at least 100 wagers. Indeed, it persists even in the distribution of the 1,401 customers who'd managed to bet 1,000 times and more. Something else must be going on.

A couple of possibilities spring to mind. Firstly, we could hypothesise that a disproportionate number of losing customers exhibits a preference for longer odds and/or accumulators. The greater margin-weights accompanying such prices, on account of bwin's strong favourite-longshot bias, could conceivably contribute to some of the additional variance in yields observed in the left hand side of the distribution. Given the inferiority of longshot prices at bwin, however, one could reasonably question the rationality of such customers. Why bet long prices at bwin when they are so inferior compared to other brands? If customers do so, perhaps this merely confirms the idea we explored in an earlier chapter that bettors, given the opportunity, will happily become victims of the possibility effect, regardless of how poor the value expectation is. Without access to individual betting prices, however, we can only speculate.

A second explanation might be found in the way bwin manages its winners. At the risk of being blunt, it either restricts them or gets rid of them. Conceivably, this would help explain such a steep decline in customer frequency as we pass through profit line. Almost half the players in this

research sample were no longer active by the end of June 2005, a full 3 months before the end of the study period. Of course, some of those will have returned beyond September, whilst many others will have opted to quit, either because of losses or lack of interest. Given bwin's apparent intolerance towards winning customers, however, conceivably some of them will have been forcibly prevented from betting further. Speaking from my own experience, I found myself limited to essentially meaningless stakes after my third bet. Many others have expressed negative comments about the brand on the review website Top100Bookmakers.com⁸⁴, with several commenting on similar stake restrictions or outright account closures, typically instigated without explanation. In contrast to Pinnacle Sports, which is seen as a high rollers' paradise, bwin appears to prefer small-staking neophytes who are just looking for a bit of fun. Indeed, the average stake placed was just €7.89 and, for pre-match bets only, just €5.41. If the fun turns into something more serious and profitable, however, customers should expect to have their action curtailed, particularly if they are identified by the bookmaker, rightly or wrongly, as being sharp.

Whilst this sample did include some big winners, their numbers were small, with only 245 (less than 1%) of the 40,499 players during the 8-month analysis period showing a profit of over €1,000 and just 6 of them over €10,000. Without access to the betting odds for individual wagers it's hard to say how these players, particularly the 2% with 3-figure stakes, managed to avoid restriction. However, judging by the sizeable yields that many of them enjoyed, one can infer that their preference was for high risk and high return. Luck in such a high variance environment will deliver some pretty generous returns in the short term, but if bwin regards any of these players as squares it will hold fire on the restrictions, presumably because it expects them to regress towards the mean. Given that betting appears to be a domain almost entirely grounded in guesswork, it is surely all the more puzzling why brands like bwin would choose to restrict any players at all. Presumably, however, if such bookmakers perceive a greater advantage in the acquisition of new customers via the manipulation of betting odds (specifically the offering and holding of loss-leading value expectation) and the advertising potential that it generates, a necessary consequence is the restriction of players who choose to abuse their price

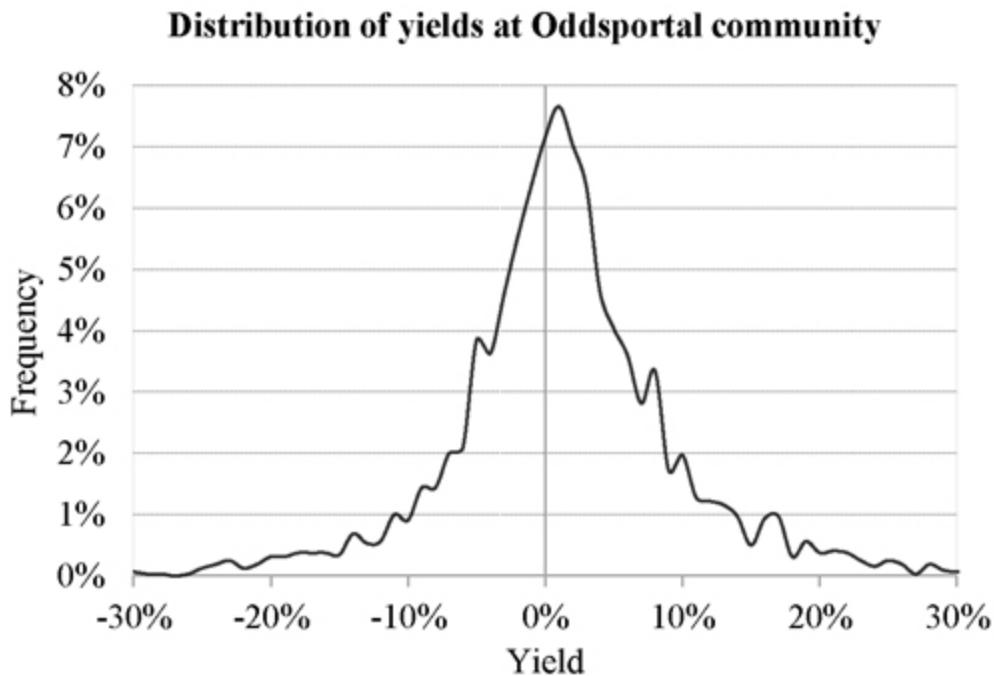
generosity.

Of course, bettors generally don't restrict themselves to just one brand, unless they happen to be particularly loyal customers or other restrictions have been imposed on them elsewhere. More usually, they will seek the best prices available, something that's easy to do by means of an odds comparison like Oddsportal.com, Oddschecker.com or Betbrain.com. In doing so, a bettor can conceivably eliminate the influence of the bookmaker's margin and sometimes even build in some arithmetic value on top. For example, by backing Interwetten's 1.35 for Arsenal, Pinnacle Sports' 12.67 for Sunderland, and Marathonbet's 6.00 for the draw (Premiership match played 20 May 2015), it was theoretically possible to lock in a sure profit of 1.39%, provided the appropriate stakes are wagered on each outcome⁸⁵. Needless to say, bettors who indulge in such practice, or even those backing just one of the outcomes, will quickly find themselves restricted by the majority of bookmakers. I've already mentioned the names of a few who use such a guillotine. Marathonbet was of course one of them. Interwetten do it, too.

However, this doesn't stop bettors tipping such prices. Like Pyckio, Oddsportal.com has its own tipping community. Unlike Pyckio, however, it doesn't restrict itself simply to brands that welcome smart bettors. Indeed, as an odds comparison service it's perfectly understandable why it would choose to make the most of its full functionality. Any registered member of the website is free to post a pick on the community. When they do so, Oddsportal automatically assigns the best market price for the pick, regardless of the bookmaker. There is now an archived database of many millions of community picks. Unfortunately, this does not include all of them. Two easily accessible archives are available. The first, containing 5,955,147 picks as of 19 May 2015, records all historical picks for those 3,201 users who have tipped at least 5 times in the previous 30 days. Presumably, there must be many more tipsters without such recent activity, including those who have since stopped posting their advice through the community. A second database, with 9,759,217 picks (again to 19 May 2015), would appear to contain a lot of those missing picks. This one records total aggregated performance by nation (114 of them) provided at least 1,000 picks have been submitted and 100 tips have been posted in the

past 30 days by tipsters from that nation. Again, however, there will still be picks unaccounted for. Furthermore, as for the bwin database, the odds for every pick are not accessible (without a hugely labour-intensive and time-consuming data scraping exercise that would be far beyond my technical capabilities). Sadly, from experience, Oddsportal is not in the habit of releasing additional data for private consumption, even for academic purposes.

Aggregated profit over turnover from the smaller and larger samples was 0.96% and 0.30% respectively. Given that Oddsportal is assigning the best market price for every pick, including prices from brands like Marathonbet and 1XBET, arguably this is yet further evidence of dart throwing. The distribution of yields for the 3,201 recently active users is shown below.



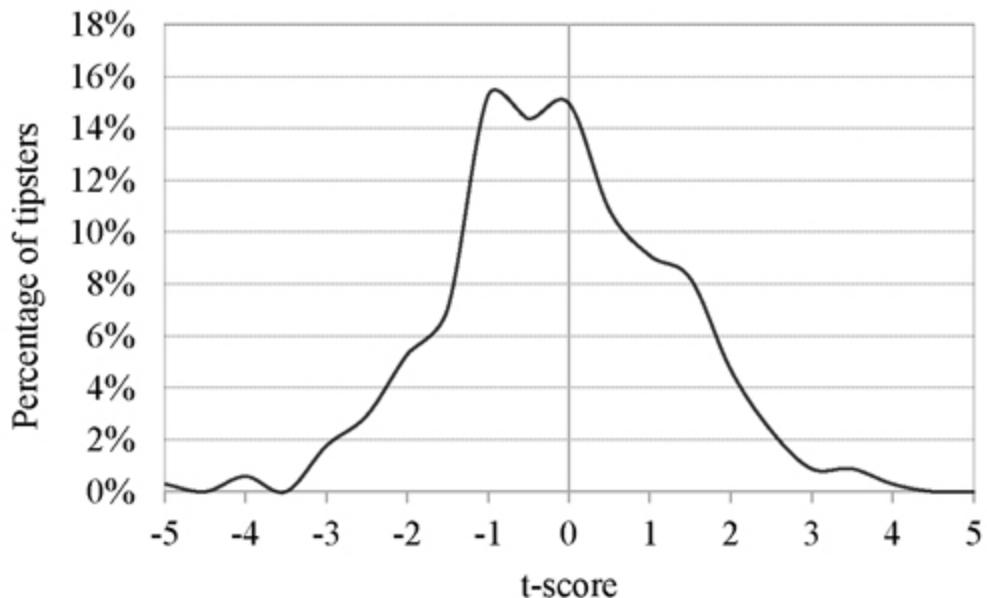
Unlike the bwin distribution, there is no asymmetry. With essentially no favourite-longshot bias, no restriction of winners (remember, these are picks rather actual recorded bets), and the majority of tipsters with at least 100 picks (89% of them in fact compared to just 34% in the bwin sample), it is logically understandable why 1,838 of the 3,201 tipsters, or 57.4%, showed positive returns. Given the shape of the distribution there is little in these figures that leads me to believe that any significant part was the result

of skill. Furthermore, the excess profitable returns are simply being achieved by the automatic selection of best market prices that have been manipulated by bookmakers aiming to attract new customers but which will prove to be consistently unbackable in the long term for those repeatedly trying to exploit them.

A final sample of betting predictions is my own. Since 2001, I have been verifying the performance, transparency and integrity of sports betting advisory services through the website Sports-Tipsters.co.uk, the majority of which choose to sell their advice for money. In 2012, the story of that work was published in my book *How to Find a Black Cat in a Coal Cellar: the Truth about Sports Tipsters*. I described a ‘black cat’ as a tipster who was capable of returning a profit through skill, over and above what can be expected to happen by chance, in other words a ‘smart’. As that book demonstrated and the analysis in this chapter has further reinforced, there aren’t very many of them; nearly all are just squares, but overconfidence and denial ensure that most refuse to accept their profits represent anything more than luck. I’ve gone back to the original data set and added 3 further years of verified picks that I have received since that time. By 31 May 2015, I had a database of 201,849 picks (excluding 2,326 void bets, usually for postponed matches or tennis retirements and walkovers). From actual stakes, the aggregated profit over turnover was 1.07%, whilst to level stakes, 1.11%. The similarity in figures should hopefully provide confirmation for those still doubting that variable money management doesn’t change the long term expectancy, merely the nature of short term risks. Such performance is more or less equivalent to the forecasters operating through the Oddsportal community. Given that the vast majority of advisory services has made a habit of advising best market prices (of course why wouldn’t they?), evidently most, if not all, of this small yield comes by way of the manipulative tactics of bookmakers that I’ve previously talked about. 152 of the 341 verified betting histories or 44.6% were profitable (to level stakes this figure was 150 or 44.0%).

The chart below shows the distribution of t-scores for the 341 services. In contrast to those t-scores calculated for the Pyckio data set earlier, where the population mean, μ , was assumed to be negative and correlated with Pinnacle Sports’ typical margin (1.025), here I have taken μ to be 0.

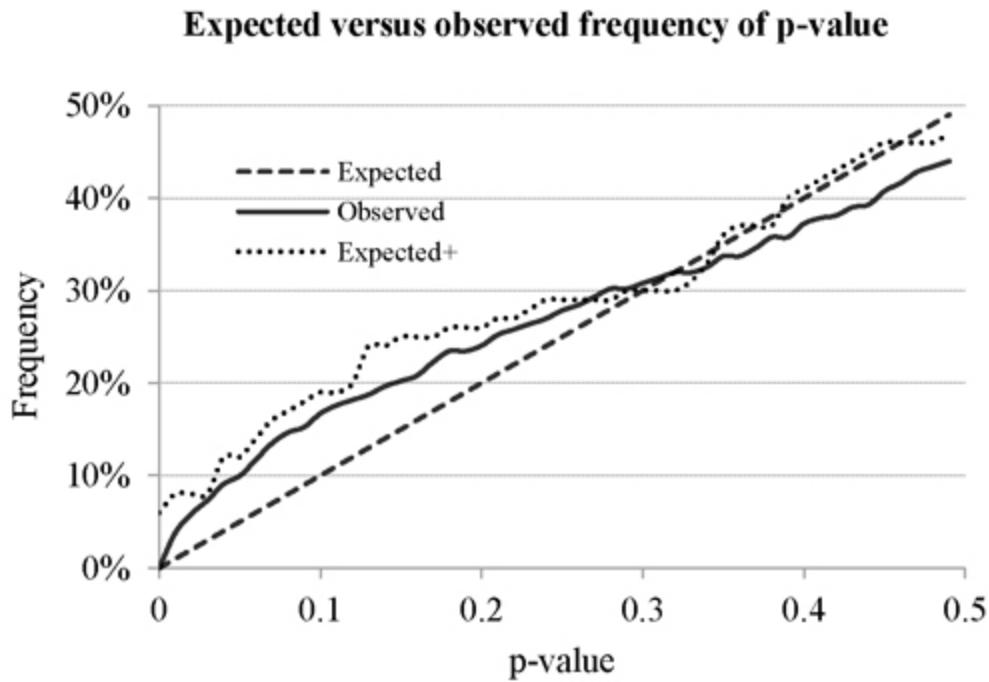
Distribution of t-scores for verified tipsters



With tipsters scouting for best market prices, arguably they have eliminated the bookmaker's margin. Indeed, one might even make the case for saying that, *a priori*, expected value should be positive with so many bookmakers now offering loss-leading betting odds in an attempt to attract new customers. The much smaller sample size compared to the previous data sets reviewed in this chapter accounts for the greater variability. Nevertheless, the underlying pattern is the same: normal.

If people selling betting advice were just guessing, how many of them would we expect to show a p-value of less than 0.01? From a sample of 341 tipsters, the answer is obviously between 3 and 4. That, after all, is precisely what the p-value means: the probability that a particular outcome could have occurred by chance. In fact, there were 13. Is this an indication that at least some of them were offering more than luck? It's possible, even probable. However, using a 1,000-run Monte Carlo simulation to model the variability in expected p-values, 2% of them produced more than 13 p-values lower than 0.01. Indeed, in one run there were 20. The chart below compares the observed cumulative frequency of tipsters achieving p-value thresholds between 0 and 0.5⁸⁶ (solid line) to that expected purely by chance (dashed line). A third dotted line shows one (extreme) outcome from the Monte Carlo simulation. Evidently, whilst more advisory services

verified by Sports-Tipsters performed better than might be expected by chance alone, their distribution of performance was not beyond the realms of possibility according to the natural variance of luck.

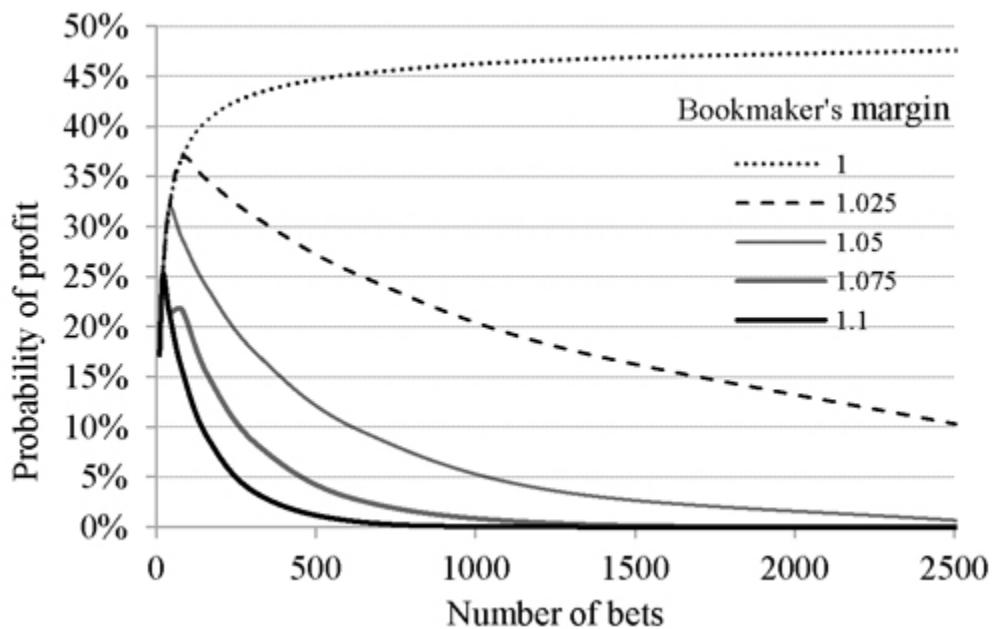


Another clue as to whether any of those 13 advisory services were genuinely offering something more than luck can be found in the evolution of their betting returns. 8 of them witnessed a slow and protracted stalling of profitability, which in most cases led to the closure of the service. Just 2 of those 8 were still active as of writing, along with 1 of the other 5. As I've previously observed, skill can be distinguished from luck by its consistency and reproducibility. Betting histories that start to tread water are clearly a sign that performance is reverting to the mean, where little more than chance was underpinning the performance that came before.

So how many people win at betting? The old adage that 95% of them are losers is not particularly helpful. As this discussion has demonstrated it appears to depend almost entirely on how lucky we've been, how long we've been betting, and with whom. We can see that lots of people make money betting on sports, and sometimes over quite long periods, particularly at brands that have small margins. Indeed, my own affiliate data of players whom I've referred to six major bookmakers reveal that 27% of

them have been profitable. As would be predicted by regression to the mean, the longer they've bet, the more likely it is that they lose. 31% who've bet between 10 and 50 times have been profitable. This falls to 26% for those with 50 to 100 wagers. It's 25% for 100 to 1,000 wagers and finally just 17% for those with over 1,000.

The chart below illustrates the hypothesised probabilities of showing profitability after a specified period of betting with variable bookmaker's margin, assuming betting to be simply a matter of chance where the 'true' probability of outcome is 50%.



If the odds are fair (2.00) unsurprisingly the probability of being a 'winner' tends towards 50% as the number of bets increases. For all other scenarios it's simply a matter of time (or rather number of bets) before the inevitable happens. Nevertheless, the chart does provide a rough picture of the proportion of bettors we can expect to find in profit. For example, we should expect about a third of lightly-betting customers (fewer than a few hundred bets) with a brand like Pinnacle Sports (having a typical margin of around 1.025) to be profitable. That's pretty much what we saw for the Pyckio community which averaged 178 bets and where 35.4% was profitable. The long term prognosis for them, however, is not good. After a few thousand bets, the probability of showing positive returns falls to 10%. At a brand like bwin with a far bigger margin (1.10), the law of large

numbers has a much faster and more dramatic impact. Even by only 50 bets (the median number of bets for the 40,499 bwin customers), the chances of being in profit are theoretically just 16% (for even money wagering). Similarly, that's close to the observed figure of 13.4%. Inevitably, in the long run squares who continue to use such a brand have about as much chance of being a 'winner' as I do of becoming Prime Minister. Regression to the mean will be, to all intents and purposes, absolute. Of course, that is provided they don't exploit any artificial arithmetic value expectation deliberately offered. If they do, the bookmaker will direct them quickly to the exit.

More important than knowing the proportion of winners, however, is the explanation for why they are winning. Is it luck or skill? Specifically, can bettors consistently keep winning if they carry on betting and can they predict that through skill, or has what has happened before simply been a matter of chance? Given the substantial size of some of these data sets, some with millions of bets, these analyses unequivocally reveal that, if any sharp bettors really exist (beyond those who can simply see a bookmaker manipulating a price), they are outnumbered by squares many hundreds of times over. Everybody else is largely just chucking darts randomly whether they choose to accept that or not. Of course, most bettors choose not to, and particularly those with profits.

"What do you think I am; a nobody?"

"Frankly, yes, I do, because you're missing the point. I never said you couldn't make a profit, I'm just saying you probably didn't make it happen."

As I've previously discussed, it's easy to mistake profits for skills because everyone likes a story, particularly a causal one, all the better if it's a self-serving narrative to massage the ego and enhance one's sense of control over destiny. Young men, in particular, are overconfident about their abilities to make money from betting, presumably for evolutionary reasons to do with risk taking. It comes as no surprise, then, to find that 92% of bettors in the 2005 bwin research study were male with an average age of 31, a statistically significant 3 years younger on average than the women. The late Nobel Prize-winning economist Armen Alchian made a related, if slightly more forgiving, observation than mine with regard to success in

business. He neither said that most successful businessmen are lucky, nor that skill doesn't matter; rather he argued that it's frequently very hard to tell the difference.

Financial Monkeys and Random Walks

In his now celebrated *Random Walk Down Wall Street*, Burton Malkiel said:

“On Wall Street, the term ‘random walk’ is an obscenity. It is an epithet coined by the academic world and hurled insultingly at the professional soothsayers. Taken to its logical extreme, it means that a blindfolded monkey throwing darts at a newspaper’s financial pages could select a portfolio that would do just as well as one carefully selected by experts.”

Such an epithet is evidently well suited to almost everyone betting on sports. I also believe it can equally well be applied to those gambling (yes, gambling, not investing) in financial markets. Moreover, as widespread as the dart throwing is the illusion of skill and the denial that it breeds. Arguably, there is even less scope for the expression of skill than in sports betting markets because the world of finance is much larger with so many more players. As Mauboussin reminds us, paradoxically this is not because investors are unskilled. On the contrary, advances in data availability and processing power, coupled with the increasing rewards of success, are attracting an ever-higher calibre of investor to play the game. As absolute skills increase, however, relative differences diminish and luck plays an ever more significant role. Perhaps more importantly, because it appears to be far easier to ‘win’ in financial markets than betting on sports, it’s far more likely that you’ll attribute that ‘winning’ self-servingly. The engine of capitalism has consistently ensured that the long term trend for returns from financial investing is positive. It would seem there is not the same cost of playing as there is in sports betting, where the value expectation is overtly negative because of the bookmaker’s margin or the exchange’s commission rate. Even that, however, is an illusion which Peter Comley, author of *Monkey with a Pin*, has thoughtfully deconstructed and thoroughly debunked. Not only are financial investors, on average, failing to have any meaningful influence on their investment outcomes, those outcomes are actually far worse than most players perceive.

John Bogle, founder and retired CEO of the Vanguard Investment Group,

has been quoted as saying that “*reversion to the mean is the iron rule of the financial markets.*” It’s a technical way of saying that what goes up in the end comes back down. The reason this happens is clear: financial markets, like their betting counterparts, are almost entirely random systems. As for bettors, investors prefer to deny such a conclusion since it necessarily implies it’s much harder to make money out of them through skill. Remember also that randomness and chance, in contrast to causality and control, are anathema to most people. I’ll come to why markets necessarily behave almost randomly in the next chapter, but to begin with let’s investigate some examples that demonstrate it.

In the chapter ‘Cleopatra’s Nose’, I first introduced the concept of regression to the mean with an example of how superior performance of mutual funds⁸⁷ during one period tends to be followed by inferior performance in the next. Specifically, the ratings firm Morningstar showed that a basket of funds that was evaluated as having performed above average during the 5-year period to 1989 performed below average in the following 5-year period to 1994 and *vice versa*. Similarly, in his bestseller *Common Sense on Mutual Funds*, John Bogle also provides an account of regression or reversion to the mean⁸⁸. Ranking mutual funds into four groups based on how they performed during the 1990s, he then compares that to their performance during the 2000s. The best performing group in the 90s suffered a 7.8% relative decline over the following decade. In contrast, the weakest performing group in the 90s saw a 7.8% relative improvement. Such powerful symmetrical regression to the mean is indicative of randomness; the balls being drawn out of the luck jar have much larger numbers than the equivalent ones coming out of the skill jar.

Daniel Kahneman offers his own typically evocative account of the illusion of stock-picking skill in *Thinking Fast and Slow*. Specifically, he writes:

“Some years ago I had an unusual opportunity to examine the illusion of financial skill up close. I had been invited to speak to a group of investment advisers in a firm that provided financial advice and other services to very wealthy clients. I asked for some data to prepare my presentation and was granted a small treasure: a spreadsheet summarizing the investment outcomes of some twenty-five anonymous wealth advisers, for each of eight consecutive years. It was a simple matter to rank the advisers by their performance in each year and to determine whether there were persistent differences in skill among them and whether the same advisers consistently achieved better returns

for their clients year after year.”

Persistence, remember, is a measure of the reliability of whether something you do can be done again for the same outcome. Where there is persistence there is skill; where there is just regression to the mean there is none, only luck. Kahneman correlated the performance of each year with every other subsequent year yielding 28 correlation coefficients for each of the 25 advisers. The average was just 0.01. Recall that the correlation coefficient, r , varies between 1 – absolute correlation – and 0 – absolute randomness. In his own words:

“The results resembled what you would expect from a dice-rolling contest, not a game of skill.”

More monkeys throwing darts. Even the Pyckio community performed better than that ($r = 0.04$). The following morning, one of the firm’s executives drove Kahneman to the airport. Defending himself against the implication that he was merely rolling dice, he exhibited the typical self-deception of a person experiencing the dissonance brought about by the information that had challenged his existing illusory beliefs: he simply rejected it.

“I have done very well for the firm and no one can take that away from me.” [“Look at my winnings, that’s all the proof I need. What do you think I am; a nobody?”]

Kahneman smiled and thought:

“Well, I took it away from you this morning. If your success was due mostly to chance, how much credit are you entitled to take for it?”

More generally, we might very well wonder how it is that so-called professionals, in both financial and sports prediction, can get away with selling what simply amounts to guesswork. Indeed, I’ll be considering that at some length in the chapter ‘A Market for Lemons’.

A key explanation for the illusion of skill in financial investment and forecasting is overconfidence. A collaboration of researchers at Duke University and Ohio State University published the findings of an investigation into the accuracy of yearly forecasts for the Standard & Poor’s index offered by chief financial officers (CFOs)⁸⁹. Analysing more than

13,300 stock market return forecasts there was no meaningful correlation between what the CFOs predicted would happen and what actually did happen after 12 months. Indeed, it was actually slightly less than zero; when they said the market would rise, it was slightly more probable that it was fall, and *vice versa*. The authors also revealed the extent of overconfidence expressed by the CFOs. When asked for a forecast, they were asked to provide an 80% confidence interval, in other words higher and lower limits which they felt had only a 10% chance of being breached. If those forecasts were valid, they should prove to be correct 80% of the time. In fact, the success rate was just 36%. To adequately reflect an 80% confidence limit would have meant predicting a range of between -10% and +30% for the 12 month growth in the index, a range 4 times wider than the typical intervals proposed.

Beginning 4 October 1988, the *Wall Street Journal* launched its 'Dartboard Contest' to test Malkiel's theory that dart-throwing monkeys could perform just as well as financial professionals in picking stocks to generate profitable returns. The monkey portfolio consisted of 4 stocks hit by the employee's darts, while 4 stock picks, each one from a professional investor, made up the competing portfolio. Whilst the rules changed over the year, stocks were limited to those listed on the major American indices with sufficient market capitalisation and daily trading volume, presumably to limit the influence of trading on much riskier stocks which could offer much bigger returns. After six months, they compared the results of the two methods. The competition ran until 2002. After the 100th contest in 1998, the professionals had won 61 of them. That seems pretty impressive, although not quite significant at the 99% confidence level (p-value = 0.0105 according to a binomial distribution). However, they only outscored the Dow Jones Industrial Average 51% of the time, as near to chance as makes no difference.

Actual returns made better reading for the professionals. After the final and 142nd contest collectively, they had generated an average 6-months investment gain of 10.2%. The monkeys meanwhile managed just a 3.5%, whilst the Dow Jones index achieved 5.6%. 1-0 to the professionals it would seem. Not so fast, cries Malkiel. From the start, he had been complaining that because the professionals' picks were published in the *Journal* they were experiencing an 'announcement effect', with readers

buying them and pushing up their value. Two days following publication, the professional picks had average abnormal returns of 4%. However, those returns partially reversed within 25 days. A bigger challenge was that, despite the competition rules, the returns were not sufficiently risk-adjusted and did not take into account the influence of dividends⁹⁰. The professionally picked stocks tended to have lower dividends, and higher earnings per share and price to earnings ratios, all a sign of greater risk. Conceivably, the professionals knowingly opted for riskier stocks and got lucky. Account for this riskier stock preference and measure returns from the day after the publication of stock picks and the advantage the professionals had effectively disappears. According to one measure, investors following the experts' recommendations would have lost 3.8% on a risk-adjusted basis over a typical 6-month holding period. Finally, after the contest ended, the monkeys' stocks continued to perform, while the professionals' picks regressed to the mean.

Beyond 2002, the *Wall Street Journal* continued to allow the public to compete against their dart-throwing monkeys. In 2013, Arthur Golden turned \$100,000 of virtual money into \$5million. By contrast, the monkeys achieved a profit just shy of \$5,000. How did he manage to beat the monkeys so comprehensively? He cheated, by exploiting a loophole in the game that made it distinct from real world trading: since the virtual stocks sold for their last traded real world price, Mr. Golden was able to buy and sell large blocks of lightly traded stocks without affecting the price of those stocks. Consequently, he was able to make large profits off minute changes in stock price. Such exploitation of the rules is not dissimilar to some of the staking activity witnessed on the Pyckio community where lucky long-odds, high-stake wins are followed by a repetitive sequence of short-price, small-stake bets. Since a tipster's rating is evaluated according to the longevity of his record, he can artificially manipulate his performance to look much better than it really is (as measured according to level staking). Of course, in the real world no rational bettor would (or conceivably could) stake 10 times as much on a 10/1 longshot as on a 1/10 favourite (given typical market staking limits). Similarly, many tipsters selling their betting advice will frequently measure their performances according to prices that no individual could realistically achieve without using numerous aliases to

cope with the endless restrictions and bans he would inevitably be subjected to. Such brazen activity serves to remind how ‘professionals’ can perform better than random chance simply by exploiting loopholes and playing in ways impossible for the rest of the market to match.

More than 40 years after Malkiel first published *Random Walk Down Wall Street* (1973), he continues to believe even more strongly in his original thesis that stock price movements essentially mimic a drunken man’s walk, despite numerous attempts by proponents of prospect theory to show that an investor’s irrationality creates outcomes that are non-random. That’s as may be, but irrationality is not necessarily the same thing as predictable and consistent. Specifically, he argues that, because of the underlying randomness, a strategy of buying and holding all the stocks in a broad stock market index would likely outperform professionally managed funds whose expense charges and trading costs detract substantially from investment returns. Indeed, in his latest (11th edition he makes the case with great simplicity.

“An investor with \$10,000 at the start of 1969 who invested in a Standard & Poor’s 500-stock index fund would have had a portfolio worth \$736,196 by 2014, assuming that all dividends were reinvested. A second investor who instead purchased shares in the average actively managed fund would have seen his investment grow to \$501,470.”

Academic and journalistic literature is littered with examples of how poor a typical professional fund manager’s performance is when compared to his or her benchmark index. Malkiel himself states that since he first wrote his book, more than two-thirds of professional portfolio managers have been outperformed by the unmanaged S&P 500 Index. The Motley Fool, a multimedia financial-services company providing financial information for investors, claims that around 80% of all actively managed funds undershoot the stock market average over the long term⁹¹. In *The Wisdom of Crowds*, James Surowiecki writes that between 1984 and 1999, almost 90% of mutual fund managers underperformed the Wiltshire 5000 Index (a market-capitalisation-weighted index of the market value of all stocks actively traded in the United States). In a study of more than 24,000 mutual funds, the consumer-finance website NerdWallet⁹² found that only 24% of professional investors beat their benchmark indices over the 10 years to 2012. Actively managed mutual funds turned in an asset-weighted

average return of 6.5% over the decade, compared to 7.3% for the passively managed index products. It should be conceded, however, that before fees active managers did actually outperform the market by 0.12%, but because the costs of dealing with them (1.07% on average) are much higher than for passive fund management (0.15% on average), active investors are left with less in comparison.

0.12%. Wow! Is that a figure investors thinking of paying experts for advice should get excited about? Hmm. Far more importantly, what does it actually mean to ‘beat the market?’ Investopedia describes it as achieving a better return than the market average or benchmark. We can regard the market average as something equivalent to value expectation in betting, although of course the former (excluding costs and the effects of inflation, as we will shortly see) is typically positive, in distinct contrast to the latter. If your returns exceed the percentage return of the chosen benchmark, you have beaten the market. Well, so what? How did it happen, because of your skill, or through blind luck? Earlier, we saw that many bettors beat their market, that is to say, their value expectation as defined by the bookmakers’ prices, with some actually managing to show a profit. Similarly, it’s easy to find fund managers with positive alpha⁹³ (α) who’ve beaten the market, even after the effects of their charges are taken into account. It’s much harder to determine if a particular manager was lucky or skilful at doing it. I’ve demonstrated that, in betting, the procurement of consistently profitable outcomes via skill is almost non-existent. Is there any similar research that’s tested the same hypothesis in finance?

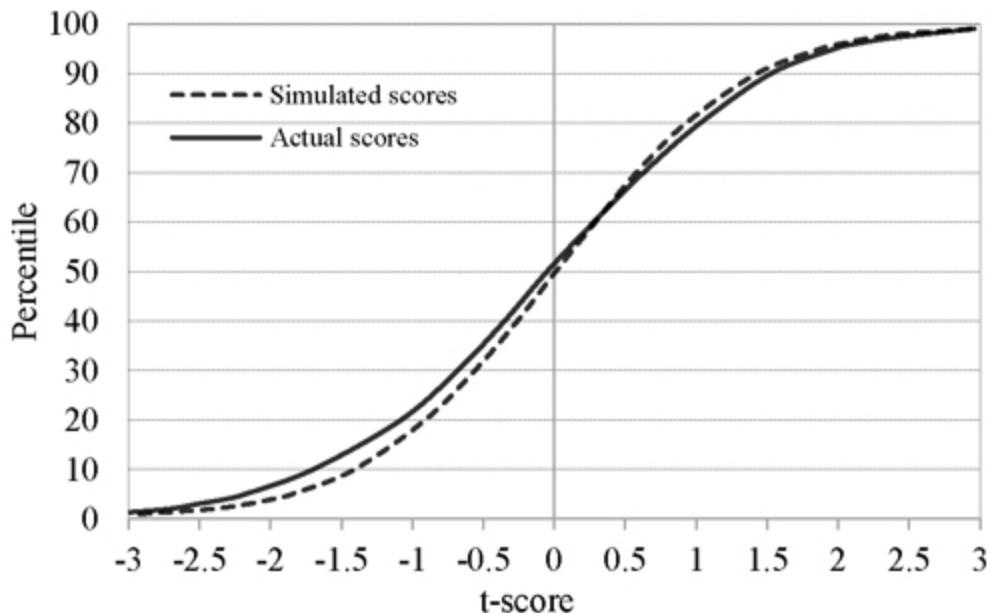
Noble Prize-winning economist Eugene Fama, with colleague Kenneth French, searched for a significant distinction between luck and skill in mutual fund returns⁹⁴. They didn’t find one. Focusing on the performance of US equity mutual fund managers from 1984 to 2006, they found that in aggregate they had performed close to the market before costs ($\alpha = 0$) and below the market after costs ($\alpha < 0$). The aggregate results imply that, if there are mutual funds with positive α , they are balanced by funds with negative α (before costs). In terms of α , this is equivalent to a zero-sum game. Indeed, relative to the market, that is how it has to be. If you have beaten the market (as measured by some benchmark index) someone else must have fallen short. Essentially, this is no different to any other form of

zero-sum gambling, whether on sports, poker, lottery, bingo or casino games. Winners have to be paid for by losers.

To test whether any of the ‘winning’ fund managers were exhibiting skill, Fama and French compared the distribution of fund returns to a distribution of 10,000 simulated portfolio returns formed with randomly selected stocks, assuming aggregate $\alpha = 0$. Essentially, this is the same approach I used to compare the performance of Pyckio’s community of bettors to that which would be predicted purely by chance. Recall the distributions of actual and simulated t-scores: they were almost identical. Similarly, for fund managers Fama and French’s simulation results were disheartening. The overlay of actual returns (even before costs are considered) with simulated ones was very close, meaning almost all were the result of random stock picking and not skill. By way of example, the chart below is a graphical adaptation of their results comparing actual and simulated t-score percentiles of α for gross returns from mutual funds that had at least \$5 million of assets under management (AUM) during the study period. Very similar outcomes exist for \$250 million and \$1 billion AUM fund groups.

It’s hard to convince oneself that anything other than chance is operating. Yes, the chart shows slightly more positive t-scores in the actual outcomes than would be predicted by chance alone, but from this marginal difference Fama and French concluded that only 2.38% of funds (equivalent to 2 standard deviations) exhibited true $\alpha > 1\%$ per year, and that furthermore, as I concluded for bettors, if there are managers with sufficient skill to cover costs, they are lost among the mass of managers with insufficient skill. Not that the fund managers would probably agree: “*What do you think I am; a nobody?*”

Percentiles of t-scores for actual and simulated gross fund returns (AUM > \$5m Jan. 1984 to Sept. 2006)



So much for professional fund managers; can private investors do any better? Sometimes it's suggested that individual investors have an advantage over large institutional investment firms because they can act more quickly and buy stocks that are too small for the latter to bother with. Do such hypothesised advantages improve the odds that individuals can beat the market? The UK Stock Challenge⁹⁵ offers one way to find out. Running since 2003, their annual and monthly competitions have attracted thousands of non commercial, non professional private investors to pit their wits against the London stock market. Whilst no real money changes hands, in Peter Comley's opinion the annual competition at least does a fairly good job of mimicking the buy-and-hold strategy that longer term investors adopt. Indeed, such a strategy is considered by Burton Malkiel to be the only genuinely credible strategy capable of delivering consistent investment returns, largely because it avoids over-trading and the excessive costs that accompany such behaviour. Costs of trading, however, are not considered within this competition, which asks participants to select a portfolio of 5 different stocks listed on the London Stock Exchange at the beginning of the year. The relatively smaller number permitted will inevitably increase the variability of observed returns compared to a benchmark like the FTSE

100 index which, as its name implies, is made up of the leading 100 stocks on the London Exchange. In contrast to the Wall Street Dartboard Contest, dividend payments do count towards final returns. The winner at the end of the trading year is the investor with the highest portfolio value. Let's see how the investors have performed.

Based on 11 full years of competition data (2004 to 2014), there have been a total of 3,286 portfolios submitted, although many of them have come from the same individual investor playing over several years. Conveniently, the period of the competition has covered at least one major boom and bust stock market cycle, allowing for a large range in annual returns. Over the 11-year period, the FTSE 100 index has grown by 47% with an average annual growth of 3.17%. Only 3 of those years (2014, 2011 and obviously 2008, the year of the global financial crash) experienced declines. By contrast, the average competition portfolio lost 0.41% per year (an underperformance compared to the FTSE 100 of 3.58%), with 59% of them failing to show any profit over a 12-month period and only 35% of annual portfolios beating the annual FTSE 100 performance. Chance alone would predict 50% before the deduction of costs.

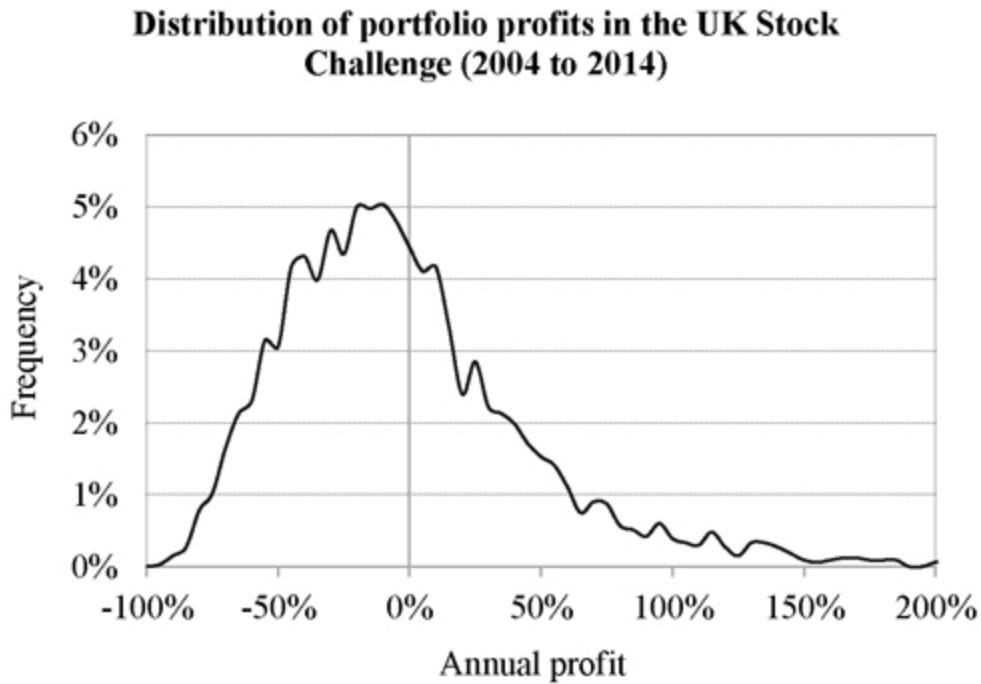
Year	Players	Average player profit	FTSE 100 growth	% best FTSE 100	'Monkey with a Pin' profit	'Monkey with a Pin' rank	'Monkey with a Pin' % rank
2004	133	6.10%	4.21%	43%	4.95%	57	43%
2005	274	6.91%	16.10%	36%	-4.54%	176	64%
2006	292	2.03%	9.68%	36%	40.40%	26	9%
2007	349	-12.11%	2.51%	26%	-13.43%	176	50%
2008	327	-58.41%	-31.91%	7%	-46.63%	78	24%
2009	309	76.08%	18.00%	83%	65.60%	144	47%
2010	351	41.12%	8.85%	67%	13.68%	215	61%
2011	418	-35.60%	-7.40%	10%	-5.79%	40	10%
2012	327	-4.51%	4.68%	32%	26.28%	59	18%
2013	232	15.51%	12.69%	45%			
2014	274	-23.75%	-2.57%	15%	-17.37%	98	36%
All	3286	-0.41%	3.17%	35%			36%

Evidently, some of this underperformance might be explained by the impact of the greater variance in 5-stock portfolios (in contrast to the 100-stock benchmark). Remember, in higher-risk environments good/bad luck will tend to deliver proportionally better/worse returns. Understandably, investing in 100 stocks would spread risks much more thinly than investing in just 5. Perhaps investors just experienced more bad luck than good over the 11 years. Overall, the standard deviation in annual player profit was 54%, compared to just 14% in the FTSE 100. Clearly, in years where the market grew, participants did much better. In 2009, for example, where the market rebounded from the financial crash with an 18% growth, the average competition portfolio experienced a 76% profit, with 83% of them outperforming the FTSE 100. The previous year, however, when the crash wiped out almost 32% of the value of the market, competition investors saw average losses of 58%, with just 7% outperforming the FTSE 100 (and only 1 player out of 327 actually showing a profit). Across the 11 years, players in aggregate outperformed the FTSE 100 only four times (2004, 2009, 2010 and 2013, all years with positive growth in the market).

In 10 of the 11 years (with the exception of 2013) the competition included a randomly picked portfolio, fittingly named ‘Monkey with a Pin.’ The table above shows its annual profit, competition ranking and percentage rank. In 2011, for example, ‘Monkey with a Pin’ managed to beat 90% of the other players. On average over the 11-year period (excluding 2013), it typically managed to beat almost two-thirds of them. Adjusting for the missing data in 2013, ‘Monkey with a Pin’ actually outperformed the FTSE 100 by an average of 3.53%, and the other players by 7.97%.

How were the annual returns for competitors distributed? The worst performance saw over 93% of the value of its portfolio wiped out. The best saw a profit of 469%. Understandably, positive growth figures are less constrained than their negative counterparts. The worst you can do is -100%. The best, by contrast, is potentially limitless. Consequently, the distribution of player returns is negatively skewed (although presumably their associated t-scores, if we could compute them, would be more symmetric). In fact, 150 (or 4.6%) of the annual portfolio profits achieved were greater than 100%. Yet, evidently, the underlying shape of the overall distribution of performance is depressingly familiar. The implication, once

again, is that stock pickers are exhibiting little, if anything, more than guesswork. The fact that a ‘monkey with a pin’ has tended to do better really says it all.



It would appear, then, that not only does stock market investment represent a zero-sum game (when compared to benchmark), but that almost all its players (whether amateur or professional) are engaged in a kind of random Brownian dance. Skill is almost non-existent and regression to the mean is rampant. On the plus side, and in contrast to betting, players have the market on their side. In sports, the benchmark is zero or worse. In finance, the engine of capitalism ensures that it’s positive. All would be well if dancing in the market was free of charge, but it’s not. In betting, market facilitators (the bookmakers and exchanges) charge us for the pleasure of letting us play. In finance, similarly, there are all sorts of costs that will eat into any returns we may have generated. In his ebook *Monkey with a Pin*, Pete Comley carefully deconstructs these charges, many of them hidden or at least embraced with an air of self denial about the true size of the impact they have. Let’s take a look at what he found.

To begin with, there are trading commissions when you buy and sell shares. In recent years these have come down as online trading firms have

competed for business. The percentage you pay will significantly depend on how much you trade, with larger trading volumes typically eligible for discounts. A typical small trader might be paying in the region of 1 to 2% per transaction. Secondly, you lose in the spread, the difference between the buy and the sell price. This is essentially the commission the market maker takes to facilitate the buying and selling. The more buyers and sellers, the easier it is for transactions to take place, and the smaller the spread will be. It's exactly the same in sports spread betting and not dissimilar to traditional fixed odds or handicap betting, where generally the bigger the sports market the smaller the bookmaker's margin since it's much easier for him to balance action on both sides. Spreads as low 0.05 to 0.1% will exist for heavily traded shares as well as full indices. In contrast, unpopular and thinly traded penny stocks might see spreads as high as 20%. Evidently, the more you trade the more these transaction fees will cost you.

Investors who prefer to let 'professionals' manage their funds for them will alternatively pay fees to the firm doing so. The obvious downside to letting someone else worry about how your investments are performing is that these costs will typically end up being higher than if you did it all yourself. Many funds are sold with an initial charge, often around 5% of the value of your investment. If you hold it for 5 years this equates to a 1% additional annual charge. Some funds might also apply an exit charge, too. All funds, however, charge an annual management charge (or AMC), often reported, together with other administration, legal and audit costs the investment firm incurs, as the total expense ratio (or TER). In the UK the average AMC and TER are around 1.1% and 1.3% respectively. Sadly, however, the term 'Total' is a misnomer. It doesn't actually include all the costs, most specifically the hidden transaction costs of trading, the price spreads and stamp duty, in the UK a 0.5% tax charged on pretty much all share purchases. The more the fund turns over its portfolio of stocks, the greater these hidden costs will be. One way to evaluate how significant they are can be provided by the portfolio turnover rate (PTR), a measure of how frequently assets within a fund are bought and sold by the managers over a 12-month period. Unfortunately, since 2011 UK regulated funds are no longer required to publish a fund's PTR, the thought being that it was too hard for investors to understand. The irony was that it was introduced by the EU to increase transparency in the industry, allowing investors to make

more informed choices. Evidently, investors in the US and Australia are cleverer than those in the UK, since the PTR is still reported there. Pete Comley has estimated that a PTR of 100% will translate into about an additional 1% in extra hidden costs. In 2011, the average PTR of a managed fund was 89%.

One way of reducing the TER (and presumably PTR as well) is to invest in passively-managed funds, those which simply track a variety of market indices, for example the FTSE 100, that require little effort on the part of the fund manager. Typical TERs for passive funds in the UK are of the order of 0.1 to 0.3%. Of course, if the fund is merely replicating the benchmark, any additional charge you pay for its management, no matter how small, will mean you are underperforming the market. The argument for a more active approach is that you are paying for the skill, knowledge and research of a fund manager. Yet, as we've seen, true α (that is to say, not reliant on luck only) amongst the professionals appears to be almost non-existent. According to Vanguard⁹⁶, the investment management group, high-cost managers, furthermore, have underperformed relative to their low-cost rivals over the past decade or so. Only through embracing higher risk have some active managers been able to show superior returns relative to those passive tracking indices. Presumably, others following similar strategies fared far worse.

In addition to trading costs and management fees, you also have to consider the tax implications of your investing. I've already mentioned the 0.5% stamp duty. Capital gains tax is also charged on the realisation of profits over a certain annual threshold, in 2015 £11,100. Any gain made over that figure will be taxed at 18% or 28%, depending on whether you are a standard or higher rate tax payer. Tax is also applied to dividends at 10% for basic rate payers and 25% for higher rate payers once an additional 10% tax credit has been taken into account. One way to avoid paying tax on financial investments is to put them into vehicles called Individual Saving Accounts (or ISAs), but the amount you can put in is limited (in 2015 the annual limit was £15,240). Plainly this is one area where bettors have an advantage, at least in the UK. Since the Government scrapped the 9% betting tax punters had to pay every time they placed a bet, betting in the UK has essentially been tax free. The Inland Revenue prefers not to insist

bettors declare winnings on their annual tax returns for the simple reason that it avoids the headache of dealing with repayments the vast majority of losing bettors would be eligible to claim. Whether these same rules would apply to the few professional gamblers who earn a significant living from their betting is a moot point, although I'm not aware of any case where the Inland Revenue has tested it. Players, however, should not be complacent that this will remain the status quo indefinitely. In Germany, for example, wins from sports betting have been taxable at 5% since July 2012, although poker still remains untaxed.

So where do all these costs leave us? According to Pete Comely, in aggregate they can amount to as much as 6%. In the three years since he wrote *Monkey with a Pin*, TERs have certainly fallen (in 2012 they were averaging around 1.7%). Furthermore, I think he may overestimate the impact so-called skilled (or positive- α) professionals have on the rest of the market. His thesis is that, since there is a group (mostly city professionals) who do consistently manage to beat the market, their existence inevitably means that the average investor must have an effective α of less than 0, given the zero-sum nature of this game. Whilst the work of Fama and French would imply that the size of any true-positive- α group is small, presumably their ability to play with much larger sums of capital impacts disproportionately on the rest of the market. I'll be taking a closer look at this 'winner takes all' phenomenon later in the book. Nevertheless, even with today's more competitive TERs and an assumed value of 0 for α for the typical investor, we are still left with aggregated costs in the region of 3 to 5% depending on the type of investing (direct investment or active/passive fund management).

Seen in this light, it makes Betfair's commission rate of 5% on profits (effectively 2.5% for even money wagering), Pinnacle Sports' 1.5 to 2.5% margin, and indeed even the casino edge in roulette (as low as 1.4% for some European tables) seem very competitive. Of course, the benchmark profitability in betting and casino gambling is 0 (the fair expectation) from which our playing costs are deducted. In financial investment, however, it has historically been positive. The FTSE 100 index, for example, grew at an annualised rate of about 3.5% between 2004 and 2014. Nevertheless, where will that leave us if our costs of investing have been 4 or 5% per year? Furthermore, we haven't even yet considered the effects of inflation.

Mutual and pension funds conveniently forget to talk about this little fly in the ointment. In fact, it's not so much a fly but a fire-breathing dragon. Pete Comley illustrates influence of inflation with his own pension fund.

"I have paid in £17,250 and my policies' current value after 23 years is £27,500. At first sight this may seem a good return – a growth of 4.2% pa. However, inflation averaged 3.5% during that time – i.e., the policy beat inflation by a mere 0.7% pa – so it is effectively worth hardly more than I paid in. This is despite the investment period including a large chunk of the biggest stock market returns in recorded history and an average return of 5.7% pa above inflation. In reality the fund lost 5% pa against the market and also did not get anywhere near paying out the absolute minimum projected at inception of £48,000."

Observe that his 5% annualised shortfall is roughly in line with the size of his charges. Unfortunately, inflation erodes the true value of wealth only slowly, whilst the costs of investing are perceived by most to be small enough to be conveniently forgotten about. The reason they matter, however, is because of compounding. Consider a stock investment initially worth £1,000. Suppose it grows at an annualised rate of 10% over the next 25 years. How much will it be worth? In actual terms, almost £11,000; in real terms, by contrast, just £1,640 if annualised costs and inflation are 5% and 3% respectively. Of course, as an investor you'll be doing very well to maintain a 10% average annual return. The legendary investor and CEO of investment firm Berkshire Hathaway, Warren Buffett, argues that the typical zero- α investor should expect to see about 6 to 7% (driven by an average of 3% GDP, 2% inflation and dividend payments on top). In other words, discounting for dividend income, you should expect the growth of your investments to keep pace with the background economy. During periods of slower economic growth and higher inflation, you'll be doing well simply to stand still. Inevitably, however, our capacity for irrationality, loss aversion, systematic bias, self-delusional overconfidence and illusory feelings of control will ensure that many investors in the world of finance do considerably worse. You just have to look at the asymmetric distribution of annual returns from 11 years of the UK Stock Challenge to remind yourself of that.

⁸¹ For a more detailed examination of money management and its associated risks I refer the reader to my first book, *Fixed Odds Sports Betting: Statistical Forecasting and Risk Management*.

82 $R^2 = 0.9875$. See below for a discussion about R^2 .

83 LaBrie, R. A., LaPlante, D. A., Nelson, S. E., Schumann, A. & Shaffer, H. J., 2007. Assessing the Playing Field: A Prospective Longitudinal Study of Internet Sports Gambling Behavior. *Journal of Gambling Studies*, **23**, pp.347–362.

84 <http://www.top100bookmakers.com/comments/bwin.php>

85 Such a bet is called an arbitrage. To calculate the correct stakes simply calculate the inverse of the decimal odds for each outcome and scale according to staking preferences.

86 Remember, a p-value of 0.5 from a 1-tailed t-test is associated with a t-score of 0. Negative t-scores imply unprofitable betting histories and are of no interest here. Excel returns a void p-value for negative t-scores using a 1-tailed t-test.

87 A mutual fund is an investment programme that is made up of a pool of diversified holdings including stocks, bonds, money market instruments, property and similar assets, bought by multiple investors and managed professionally.

88 Bogle, J., 2009. *Common Sense on Mutual Funds*. 10th anniversary edition. Hoboken, New Jersey: Wiley.

89 Originally released as a Duke University Working paper in 2010, it was revised and published in 2013 in the Quarterly Journal of Economics (Ben-David, I. & Graham, J. R., 2013. Managerial Miscalibration, *The Quarterly Journal of Economics*, **128**(4), pp.1547-1584.)

90 A dividend is a sum of money paid regularly (typically annually or 6-monthly) by a company to its shareholders out of its profits, after all expenses and taxes have been paid. Investors can typically choose to receive dividends as income or reinvest them back into the company in the form of additional shares.

91 Can You Beat The Market? <https://www.fool.co.uk/investing-basics/how-to-invest-in-shares/can-you-beat-the-market/>

92 Only 24% of Active Mutual Fund Managers Outperform the Market Index. <http://www.nerdwallet.com/blog/investing/2013/active-mutual-fund-managers-beat-market-index/>

93 Alpha is one of five technical risk ratios used to assess the risk-reward profile of a fund; the others are beta (a measure of risk), standard deviation, R-squared, and the Sharpe ratio. Alpha is a risk-adjusted measure of the so-called active return on an investment. A positive alpha of 1.0 means the fund has outperformed its benchmark index by 1%. Correspondingly, an alpha of -1.0 implies an underperformance of 1%.

94 Fama, E. F. & French, K. R., 2010. Luck versus Skill in the Cross-Section of Mutual Fund Returns. *The Journal of Finance*, **65**(5), pp.1915-1947.

95 <http://www.stockchallenge.co.uk/>

96 <http://www.thisismoney.co.uk/money/diyinvesting/article-2616861/Is-fund-manager-worth-active-funds-failed-beat-benchmarks.xhtml>

GINSBERG'S THEOREM

High school students of physics will well remember learning all about conservation of energy and the impossibility of perpetual motion machines. These form part of the laws of thermodynamics which describe how the quantities of temperature, energy and entropy behave under various circumstances, and forbid certain phenomena. In 1959, the British scientist and novelist Charles Percy Snow gave the Rede Lecture (at Cambridge University) entitled *The Two Cultures and the Scientific Revolution*, in which developed a gaming metaphor as an excellent way to remember them.

1. **You cannot win** (a restatement of the first law of thermodynamics; that is, you cannot get something for nothing, because matter and energy are conserved).
2. **You cannot break-even** (a restatement of the second law; that is, you cannot return to the same energy state, because there is always an increase in disorder; entropy always increases).
3. **You cannot get out of the game** (a restatement of the third law; that is, because absolute zero is unattainable).

More popularly, this parody of the laws of thermodynamics in terms of someone playing a game is attributed to the American poet Irwin Allen Ginsberg, from whom the theorem gets its name. Put simply, in gambling, there is no such thing as a free lunch. Presumably, Snow and Ginsberg had casinos in mind when formulating the metaphor. The data and discussion from the last chapter, however, arguably mean we should extend this to betting and investing as well. Regardless of the theoretical possibility of positive expectation, few consistently manage to achieve it. Almost all bettors and investors don't win (at least relative to their respective benchmarks, over the long term and as a consequence of skill), they don't break-even (after accounting for costs and commissions), and the only way to get out of that game, or to be more precise to avoid the gambler's ruin,

would be to have infinite wealth.

Ginsberg's theorem is a light-hearted take on gambling, betting and investing, but it has a serious and relevant significance. The fact is that for the vast majority of players these activities do represent zero-sum games where the dynamics of the system make it highly unfeasible, if not quite outright impossible, to beat it. Two questions then arise: why is it so hard to have a free lunch and what makes so many people believe it is possible? This chapter will attempt some answers.

History is Written by the Winners

I've already explored the many cognitive biases, and in particular overconfidence, that help explain why many bettors and investors alike are sure they can beat the 'system.' One that arguably deserves its own special attention here is what is known as survivorship bias. Put simply, bettors and investors are fooled into believing that it's much easier to 'win' through skill than objectively is the case because all they see around them are the success stories. History, it seems, is always written by the winners. The losers are ignored, either because we choose not to see them – they don't tell an interesting story – or because they've disappeared from view. It's easy to be impressed with success if that's all there is to see (WYSIATI). As Nassim Taleb, in *Fooled by Randomness*, narrates on the fantasy of monkeys attempting to recreate the poetry of Homer on a typewriter:

"If there are five monkeys in the game, I would be rather impressed with the Iliad writer, to the point of suspecting him to be a reincarnation of the ancient poet. If there are a billion to the power one billion monkeys I would be less impressed..."

As Taleb points out, not many people bother to count all the monkeys. Well, I've been counting them.

Survivorship bias is the logical error of concentrating on the people or things that 'survived' some process whilst inadvertently overlooking those that did not because of their lack of visibility. Survivorship bias can lead to overestimating the chances of success because failures are ignored. Its name was first coined during the Second World War when a free-thinking mathematician named Abraham Wald solved the problem of where to put

additional armour plating on the Allied bombers that were experiencing heavy losses. Initially, engineers assumed that, through an examination of the bullet holes of returning aircraft, those areas that showed the highest concentration of bullet holes – along the wings, around the tail gunner and down the centre of the fuselage – needed the extra reinforcement. Of course, it didn’t work. The mistake, which Wald saw instantly, was that the holes revealed where the planes were strongest, since these were ones actually making it back, that is to say, surviving. By contrast, no one had previously given any thought to the planes that were lost. Counterintuitively, Wald suggested putting extra armour plating where the bullet holes weren’t. The engineers’ original error was so significant, statisticians decided to give it a name: survivorship bias, or the tendency to include only successes in statistical analysis.

Unmistakably, survivorship bias will be a major source of overconfidence. The hot hand fallacy is an obvious example here, where people believe streaks of success inevitably increase the chances of further success. It can be found in all walks of life: business, health, education, sport, the arts, entertainment and, of course, gambling, betting and investing. For example, perhaps you might be thinking about opening a restaurant because you can see so many successful restaurants in your area. But what happened to all the other ones that didn’t make it? They’re gone, disappeared from view. Depending on the nature of the business and the time frame, a high proportion of start-ups will fail. As Nassim Taleb writes in his book *The Black Swan*, “*The cemetery of failed restaurants is very silent.*” Michael Mauboussin, following Jerker Denrell, Professor of Behavioural Science at Warwick Business School, describes the problem as an “*undersampling of failure.*” For example, managers will typically try to copy strategies adopted by the most successful businesses. The problem, as for Wald’s aircraft, is that they’re not looking to see what strategies the failures were using. Attributing success to a particular strategy that a winning business adopted may simply be wrong. The more important question is: how many of the businesses that tried that strategy actually succeeded? As Mauboussin explains, the problem here is that inference is drawn from outcomes, not processes. Where luck is dominant, there is very little connection between the process and the outcome. If all you care about is outcomes, you’re liable to draw erroneous conclusions. On the contrary,

don't study winners to see what caused them; study the process to see whether it consistently led to success.

One notorious expression of survivorship bias comes from the popularisation, by two journalists in particular, Malcolm Gladwell, author of *Outliers: the Story of Success* and Matthew Syed, author of *Bounce: the Myth of Talent and the Power of Practice*, of a concept called the 10,000-hour rule. This idea, specifically the belief that talent is not inherited but learned provided enough time is devoted to the task, is largely based on the research by Swedish psychologist K. Anders Ericsson investigating the influence of deliberate (high-level) practice for violinists at Berlin's Academy of Music⁹⁷. Students had begun playing at around five years of age, all putting in similar practice times, but by age eight the practice times began to diverge, some practising more than others. By age twenty, the elite performers totalled 10,000 hours of practice each, while the merely good students had totalled 8,000 hours, and the lessable performers had just over 4,000 hours of practice. Extrapolating this idea to explain performance in other domains including business and sport (and, in Syed's case, as an attempt to refute the considerable evidence pointing to some genetic component of sporting talent and the unsavoury social message that is wrongly perceived to accompany it) both Gladwell and Syed have misinterpreted the statistics in Ericsson's original work. Indeed, Ericsson himself posted a letter on his university website, titled *The Danger of Delegating Education to Journalists*, calling the 10,000-hour rule invented and restating the point that there was nothing magical about 10,000 hours, but rather it was merely an average about which there was considerable variance.

For the purposes of our discussion, however, I feel that Ericsson actually missed a trick here. In defending himself against the criticism from Syed, David Epstein, author of *The Sports Gene: Talent, Practice and the Truth about Success*, offers us a hint of what it is. Ericsson's original violin study sampled performers from a world-class academy; presumably, they had already been subjected to an intense selection process simply to get there. As Epstein says, “*in sports, this would be akin to restricting a study to only NBA centres, noticing that they had all practised a lot, and therefore concluding that practice is the only reason they reached the NBA – not*

practice plus being seven feet tall." Essentially, what Epstein is drawing attention to is the survivorship bias present in the formulation of the 10,000-hour rule. In developing the concept, neither Gladwell nor Syed, nor even Ericsson for that matter, considered the countless other participants in activities who had managed to clock up so many hours of play but who were missing from the field of view for no other reason than they had failed to make it to the elite level where they would be noticed. Understandably, that would be impossible, but to assume that all players who practice long and hard enough will make the grade simply from a study of those who have made the grade is arguably a far bigger misinterpretation of the statistics.

Survivorship bias is everywhere in gambling. My first introduction to it, before I was even aware there was a name for it, was when comparing the histories of betting advisory services I had been asked to verify with their corresponding records of performance before their owners had asked me for verification. Until 13 June 2011, my verification service Sports-Tipsters had agreed to publish the full record of a sports betting advisory service, including both the tips that I had verified and earlier ones that I had not. In all there were 120 such advisory services. Aggregating their performance together, 24,725 pre-verified tips made a profit over turnover of 17.4%. By contrast, the 90,451 tips that I saw made just 1.1%. How so? The explanation by now should be obvious: survivorship bias and regression to the mean. Sports-Tipsters employed a strict rule of passive acceptance of a service; it did not actively seek to recruit any for verification. As such, tipsters could simply self-select themselves for verification. Which tipsters would choose to do so: those who had already failed; or those who were laying golden eggs from golden geese? It's self evident, isn't it? Tipsters who fail generally don't carry on, and they certainly don't go asking for verification of their work. Sports-Tipsters was simply being treated to the 'winners'; the 'losers' had either already disappeared or simply opted not to announce themselves until luck turned in their favour.

It is probably reasonable to assume that tipsters asking to be verified are not in the business of cheating. Granted, there were a few high profile examples of services which I caught doing so (and for them presumably verification was seen as purely a short term form of advertising). In the main, however, there was rarely a reason to doubt the honesty of the records

these services claimed prior to verification. If the earlier records were not faked, what were they? Again, hopefully, the answer is self evident: lucky. I've demonstrated, fairly convincingly I trust, that most things that happen in sports betting, including for those who believe otherwise, are simply a matter of chance. The implication, therefore, must be that tipsters who choose to be verified after already building an earlier unverified history represent nothing more than the industry's lucky survivors. We can test this idea with regression to the mean. Already, I've reported that in aggregate the 120 services I was introduced to regressed from 17% 'before' to 1% 'after'. Individually, just 4 of them showed improved performance between 'before' and 'after' where the former had more than 100 tips. Indeed, including all the others with fewer, there were still only 9. Perhaps the most impressive regressor was a football and US sports advisory service called *Over Under Capper*; its 'before' (315 picks for a +10.9% yield) and 'after' (360 picks for a -7.0% yield) performance is illustrated below.



Remember, where there is skill, there is no reversion to the mean; where there is only luck, there is complete reversion to the mean, and how long is simply a matter of time. These 120 services did a pretty good job of the

latter. Since consistency is virtually absent, we are forced to draw the conclusion that an unverified history of performance from a tipster does not offer any reliable measure of the success of his sports advisory service. That's not because it's not genuine, rather it's because it has almost certainly arisen through luck rather than skill. The unlucky monkeys just disappear, and their records are no longer available for scrutiny. The luckier survivors who turn up soon discover that they were monkeys throwing darts, too. It was precisely for this reason that, on 13 June 2011, I removed all the unverified histories from my website, being as they were completely devoid of any informative value whatsoever.

Pinnacle Sports provides a nice summary of how survivorship bias affects the sports tipping industry⁹⁸.

“Let’s say we run a simulation which sees 10,000 tipsters (or monkeys, it really doesn’t matter) each with a 50% chance of either making \$10,000 a year or losing \$10,000 a year. If any tipster has a losing year, they are eliminated. The tipsters/monkeys make their predictions by simply pushing one of two buttons. If we run the test for one year 5,000 of our tipsters would be \$10,000 in profit and the same number \$10,000 in the red and binned. In year two we would have 2,500 monkeys with perfect records and if we keep going by Year 5 we would have 313 monkeys from that original cohort that would statistically be able through pure luck to make successive accurate predictions and \$50,000.”

Of course, monkeys (sorry tipsters) pushing buttons aren’t exhibiting much skill, are they, but if there are enough of them to start with, some will make it through Galton’s Quincunx machine.

“If you just focus on the winners in this process, ignoring all the other billions of monkeys producing gibberish, you’re being fooled by randomness. The simple fact is that, by starting from a large enough sample, some of the participants will end up looking like a savant by pure luck.”

As Taleb suggests (in *Fooled by Randomness*), when attempting to measure the probability of success, it’s no good just studying the sample that has succeeded. If we do, we risk turning causality on its head. Rather than suppose that success is caused by skill, survivorship bias ensures that we perceive the winners to be skilful because they have been successful.

Plainly, survivorship bias in the online tipping industry is bad news for anyone thinking of investing in it. With only the ‘winners’ visible, it’s easy to encourage the idea that winning is easy. I suppose it would be if the losers didn’t count. This is not a remark made flippantly. On the contrary,

much of the industry behaves as if they didn't. I've lost count of the number of tipsters who've tried, failed and disappeared, only to try again with all record of their past performances lost from the view of public and statistical scrutiny. Additionally, there is now an increasingly popular format for advisory services, following the original concept developed by the website Betadvisor.com, that sells the advice of a large collection of tipsters under one roof. It won't come as a surprise to learn that such services habitually 'lose' the performances of discontinued tipsters. Check any ranking table for any of these websites; the vast majority of ranked tipsters are 'winners.' As I write today (27 May 2015), for example, I have checked the all-time tipsters ranking on Betadvisor.com. 74 of the 85 tipsters listed there are in profit. Wow! This must be the place to be for buying tips. Indeed, the aggregated yield (weighted by the number of bets for each tipster) is a stunning 7%. Go back a couple of years⁹⁹, however, and you'll find a completely different set of tipsters listed there. Indeed, of the 41 active in March 2013, just 12 of them are continuing today. What happened to the other 29? No prizes for guessing the correct answer. Not that Betadvisor was doing anything much different then; the listed tipsters at that time were showing a 9% aggregated profit over turnover, with a staggering 39 of them profitable. Again, it's obvious why. The records of discontinued losing tipsters conveniently disappear. Fortunately, I was offered the opportunity to verify Betadvisor's activity. From 15 May 2011 to 31 October 2013, there were in fact 118 tipsters selling advice at some point during that 30-month period, and a full 70 of them lost money. Their 27,653 tips accounted for an aggregated yield of just 0.27%, much more like the dart-throwing we've seen elsewhere.

Significantly, such services make use of what are termed 'academies' where would-be tipsters can be trialled for their forecasting abilities prior to being released to the general public. Those who make the grade will be chosen, those who fail will be discarded. If you've understood the problem of survivorship bias thus far, you will see the inherent flaw in such a strategy. It's exactly the same as Pinnacle's monkeys. These academies don't appear to have the slightest understanding of regression to the mean (if they do they don't care), nor how to properly assess the difference between luck and skill, typically opting for attainment thresholds that

anyone could achieve afforded enough opportunity to do so (if you fail, just try again under another name). Keeping just the lucky trialists and discarding the unluckier ones creates what is technically termed creation bias. To be sure, however, this is not really about tipping services deliberately abusing data for their own financial gain, although undoubtedly that will happen, too. On the contrary, creation bias arises because of overconfidence and the illusion of skill. Companies like Betadvisor probably don't imagine they are discarding unlucky losers and keeping lucky winners; instead they genuinely believe that those tipsters who pass through their tipster academy have a true ability to return a profit from betting. A tiny handful may do, but the vast majority will inevitably just regress to the mean before being discarded and lost.

Yet even where the performance of 'losers' is preserved, arguably there still remains a residual bias in the data that could influence conclusions drawn from its analysis. Specifically, performance histories are of unequal length; 'losers' tend to stop whilst 'winners' tend to carry on. We can see the potential influence this can have by revisiting some of the betting history data I examined in the previous chapter. Consider the Pyckio data set. The 2,690 community members with histories of at least 100 picks had an average t-score of 0.03 with 50% of them scoring over 0, essentially indistinguishable from expectation given the normal (and symmetric) distribution of the scores. For the 249 members with at least 1,000 picks, by contrast, average t-score was 0.31 with 62% above 0. At first glance we might deduce that this is evidence of skill, with better tipsters lasting longer. Nevertheless, I feel that such a conclusion would ignore the real possibility that all we are witnessing is an example of survivorship bias. The longer records may be present simply by virtue of the fact they have survived by being luckier for longer; those which weren't so lucky never managed to achieve such long histories. Imagine that the remaining 2,441 bettors with records between 100 and 1,000 picks had continued to bet, such that all records in the sample were standardised with the same number of picks. How would the average t-score and percentage outperforming expectancy look then? For those still clinging to the hope that at least some of the difference will be explained by higher skill level in the long-history group, I should remind you of the close-to-zero correlation I found in those 249 histories between 1st and 2nd half performance. Such lack of consistency

really only points to one conclusion: aggregate dart throwing.

Sadly, we can't make all bettors place the same number of bets; we can't ask lucky bettors to stop at a predefined moment; and we can't demand unlucky ones continue when they've lost the will to do so. Trying to identify what would have happened if we could is exceptionally difficult. As Terry Burnham illustrates in *Mean Markets and Lizard Brains*, this would mean assigning probabilities to alternative worlds that never happened. We can only work with the world and the data we have. Yet theoretically, perhaps, we can attempt a good guess at what might happen if we could standardise the betting histories. Taking the full record of 1,073,029 tips, I've randomly created 1,073 histories of 1,000 tips each and computed their t-scores. Their average is just 0.07 with 54% above 0, much closer to expected values. It's also interesting to note that 25% of histories were profitable. That's not far off what I'd earlier hypothesised would be possible (about 20%) with a bookmaker's margin of 2.5%. Clearly, this methodology has its shortcomings but evidently it reinforces the view that data sets containing variable-length performances may very well contain inherent and unavoidable biases within them. These will not necessarily be limited to survivorship. Conceivably, an 'extinction' bias will be present as well, where particularly unlucky performance will show a tendency to self-terminate early before regression to the mean has properly got to work. You may also recall that the sample of bwin sports bettors revealed those with below average performances had bet considerably less than those with above average performances. If those with shorter records had bet longer, how would the data look then? Such musings are hypothetical, but this examination hopefully illustrates the dangers of working with unstandardised data.

Survivorship bias is rampant in the financial world, too. Indeed Pete Comley suggests it might reduce expected investment returns by as much as 1% compared to the market benchmark. To see why, we need to understand how a financial index is created. Consider, for example, the FTSE 100, a share index of the 100 companies listed on the London Stock Exchange with the highest market capitalisation, currently representing about 80% of the entire market capitalisation of the Exchange. Since beginning in January 1984, every 3 months a reshuffle of companies takes place with a few companies at the bottom being replaced by a few others whose market

capitalisation is higher, whilst ensuring that the index always starts the new quarter at the same value as it finished the previous one. Plainly, companies that get relegated will have below average performance whilst those which are promoted will be exhibiting above average growth. That ensures the index has an automatic tendency to rise since it places a disproportionate weighting on ‘winners’ rather than ‘losers.’ As such, the index suffers from survivorship bias. Sadly, investors do not have the luxury of being able to exclude their losers from their portfolio returns, unless of course they are invested directly in an index tracker. Funnily enough, that’s an observation I’ve repeatedly put to owners of tipster portfolio services who seem to think it’s acceptable to drop the ‘losers’ from public scrutiny. Typically, the response is that not all tipsters are profitable but their model allows customers to follow the best tipsters. Of course, that rather begs the question: how to tell which the best are beforehand? Presumably, if it was that simple, there wouldn’t ever be any losers for sale. Any monkey can be the best if they’re judged retrospectively on their successes.

How much can this really influence your investment returns? Comley asks us to consider the top 20 constituents on the FTSE 100 during the ten-year period from January 2001 to January 2011. Over this period, the index didn’t do much overall (-3%) but 16 of the largest shares declined in price with Marconi, the electronics giant, even going bust. Had you invested in them as a buy and holder, you would have lost 1.8% a year, or 23% over the whole period according to his calculations. Evidently this is not an objective measure of true underperformance, but presumably you get the picture. To properly calculate the impact of index survivorship bias one would need to obtain the price histories of all the stocks that had ever been part of it. That includes around 250 companies that have since dropped out. Only around 30 in today’s index were there from the start, with just 19 that have always been in it. Clearly, the makeup of the index is subject to considerable fluctuation, but it is geared towards favouring the ‘winners.’ Basing expectations of stock market returns on growth in benchmark indices will unquestionably lead to disappointment.

A wider problem with survivorship bias in financial investing concerns the problem of data availability. In much the same way as sellers of betting advice conveniently ‘lose’ their unprofitable histories, it can actually prove quite difficult to obtain historical records for companies that no longer exist,

either because they've been acquired by larger ones or because they've become bankrupt. Generally, such data is simply wiped out, meaning analysts must rely on samples of existing companies – the successful ones – to study past trends and make future predictions. Mutual fund companies are particular adept at messing about with their products, dropping poor performers because of weak results or low asset accumulation, resulting in an overestimation of the past returns. Rohleder and colleagues¹⁰⁰ at the University of Augsburg, Germany have shown from a sample of nearly 11,000 US mutual funds operating between January 1993 and December 2006 that nearly a third of them closed before the end of the sample period and only 6% of them survived throughout the whole period (the remainder being added along the way). From such a turnover rate, Rohleder calculated a survivorship bias of -0.95% per annum.

Like betting tipster academies, the mutual fund industry also indulges in creation bias. New managers are given some seed money to test their stock-picking ability via 'incubated' funds. 'Winners' are subsequently made available to the public and marketed aggressively; 'losers' meanwhile are silently discontinued. If outcomes are almost entirely and so palpably a consequence of luck, we might wonder whether the fund managers intentionally adopt such a strategy to suit their own ends or whether, like most of the tipster academies, they simply live in a sea of ignorance and denial about their own abilities at forecasting investment returns.

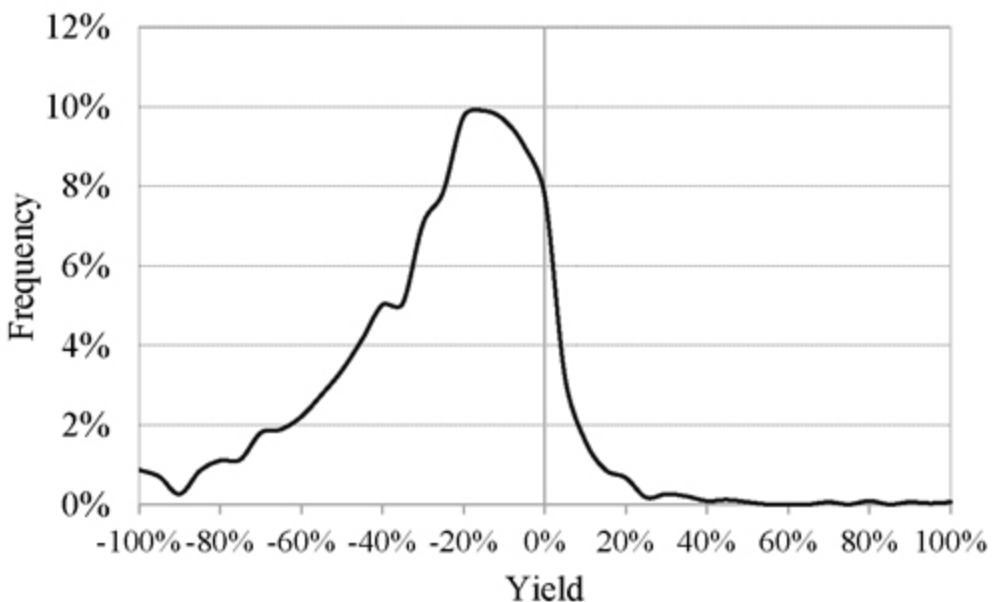
Nate Silver, author of *The Signal and the Noise* and possibly the most well-known face of statistical forecasting through his website FiveThirtyEight.com, argues that almost all professional poker players begin careers with winning streaks. He should know: he made a living playing online poker in the mid-2000s, and played the World Series of Poker Main Event in 2011. Whilst he continues to play semi-professionally, he gave up playing full time in 2007 after a series of heavy losses. Interestingly, Neil Isaacs in *Your Bet Your Life: the Burdens of Gambling*, remarks that many pathological gamblers begin their careers by winning their first bet. Are the two related? Possibly; there are many more gamblers who describe themselves as professional than their outcomes would objectively warrant. It is one thing to play for a living, quite another to consistently win and earn for your living. The commonality is surely the

illusory perception of skill and control that arises from the bias of overconfidence. This bias has also been described as the Dunning–Kruger effect wherein unskilled individuals suffer from illusory superiority, mistakenly assessing their ability to be much higher than is accurate. Winning and winning early provides the only ingredients needed to ensure we will falsely attribute our success to things we did as opposed to things that just happened by chance. If you are a successful poker player think back to when you first started playing. Did you have an early winning streak and enjoyed the feeling of defeating your opponents every time you played? New players often feel invincible when they first start winning at poker.

Evidently, what Silver describes may just be another example of survivorship bias. Silver, however, considers poker to be both a high-luck and high-skill game; it takes a long time for a player's advantage to shine through, but if it's there, inevitably it will do so in the end. Earlier, I suggested that skill would overtake luck anywhere from several hundred to several thousand played hands. Silver suggests it could take even longer, with the standard deviation in profitability as much as 16 times the value of a skilled player's expectancy. Modelling the potential profits and losses for such a player, Silver concludes that he could still be showing losses after as many as 100,000 hands, equivalent to almost two years playing 40 hours per week. If Silver is right, arguably it could be quite difficult to tell the difference between survivorship bias and evidence of skill. The winners could be genuinely skilled, or they could simply be lucky survivors. Let's take a look at some data: a sample of 3,445 poker players who registered with bwin.com during February 2005 and whose play was tracked until February 2007¹⁰¹. As for the other bwin data samples for betting and casino gambling we looked at earlier, relatively few of the players – just 11% in fact – proved to be profitable. Aggregated loss across all players was -6.76%, a figure entirely predictable given the typical poker rake on cash games is 5% whilst that for tournaments is 10%. The rake is the scaled commission fee taken by a poker room operating the game. It acts like the overround in sports betting and the house edge for casino games. Similarly, players' returns are fairly normally, if again asymmetrically distributed, with negative skew (presumably because yields are not risk-adjusted as

previously described), as the chart below illustrates.

Distribution of bwin poker player performance



Once again, we are hard pressed to pretend that anything other than luck is really operating for the vast majority of players. A closer look at those players who played for longest, however, reveals something interesting. Whilst only 8% of customers who played fewer than 100 hands were profitable, 14% had something to show for their efforts if they played between 100 and 1,000 hands. Most impressively, over a quarter of the 134 customers who played over 1,000 hands (27% in fact) were ‘winners.’ Indeed 9 of the top 20 most active players over the two-year period were profitable, as well as 5 of the top 6. As for other examples of survivorship bias, we are unable to determine whether this fat-tail sample represents customers who are winners because they’ve played a lot (lucky), or customers who’ve played a lot because they are winners (skilful). Given that the longer bwin casino customers played the less likely it was that they would be profitable, we should be encouraged to hope that at least a few of those players were demonstrating skill, as Silver has maintained should be possible. Arguably, however, we are not talking about larger numbers. Even if we were to give a sizeable number credit for being able to play poker better than the average monkey, evidently it’s still not enough for the vast

majority to be able to overcome the rake percentage. Remember, poker, like all other forms of gambling, is a relative skills contest. The elephant in the room is the commission the market facilitator demands for staging the game. It acts like an über-skilled player dominating the play of his opponents. Evidently, few players have a relatively superior skill capable of beating him, just like in betting (and investing).

Nevertheless, it is worth exploring for a minute why poker might offer the potential for a greater expression of (relative) skill than betting or financial investing. The explanation is to be found in the number of players competing, the length of play and the potential for feedback. Where outcomes are settled over short time frames, as in sports and poker, the potential for feedback and learning is greater. Poker players and sports bettors can potentially use results of play to provide confirmatory analysis of how accurate their initial predictions were, updating forecasting methodology in a Bayesian manner. Swift market closure keeps player tethered to reality. Whether such confirmatory analysis is of much benefit, given the limited relationship between cause and effect in such markets, is a moot point. For many investments in financial markets, by contrast, often there is no market closure at all, just a continuous (random?) dance of price evaluations, where nothing is ever permanently settled, until perhaps a company goes bust and the value falls to zero. We can think of this distinction in quantum mechanical terms. A bet, for example, consists of two possible outcomes (just as for Schrödinger's cat, alive or dead) whose probability is described by a wave function with a potentially infinite number of path histories that can be followed to arrive at either outcome. Prior to bet settlement, a bet, in quantum mechanical terms, is both won and lost at the same time. How we get to either is essentially unknown, according to Heisenberg's Uncertainty Principle, until we get there. When we do, the wave function collapses and one outcome or other is known with certainty. In investing, however, such certainty is never known (unless the asset you've invested in goes bust). The wave function never collapses and the value of something remains purely a probabilistic concept. Arguably, this lack of market closure (or wave function collapse if you like) makes stock market investing a tougher proposition.

Conceivably, however, the most significant reason why investing and sports betting are so hard to crack is because so many people are doing it.

Mauboussin has articulated eloquently how more and more skills paradoxically make it harder and harder to beat expectancy, since more and more players are cancelling each other out. Poker is often a contest with relatively few players, indeed where people can even choose whom to play against. Such an environment increases the probability of seeing a larger differential in skill. Betting, by comparison, is a contest typically taking place amongst hundred and sometimes thousands of players. Many see themselves as taking on the bookmaker; indeed some bookmakers rather archaically see themselves as taking on their customers. Theoretically at least, if not entirely in reality, a bookmaker should exist to facilitate the action, with players competing against themselves. The more players there are, the smaller the probability that two who meet in the market (if only in a theoretical sense) will see any significant difference in skill. Investment, of course, is a game played by millions of competitors, where prices are pushed one way and then another by buyers and sellers, all expressing opinions about the future value of an asset. The greater the number of players, the greater the likelihood (groupthink and other irrationality aside) that the perception of value will match the commodity's 'true' value (although remember in a quantum mechanical sense, there really is no such thing). Where it does so, value expectation by definition will be zero less the commission for playing in the market. A market where prices accurately reflect 'true' value is said to be efficient, and the theory describing it is known as the efficient market hypothesis. A little later, we'll take a closer look at this frequently maligned but perennially useful representation of the behaviour of betting and investment markets. In the meantime, however, I'll begin by exploring an idea that helps to explain why it frequently happens: the wisdom of crowds.

The Wisdom of the Crowd

At a local fair recently, there was a competition to guess the number of flower petals stuffed into a box. Instead of inspecting it, I simply asked the stallholder for the list of previous guesses. So far there had been 40 from which I calculated the average: 245. I submitted that figure as my own. It turned out the correct answer was 295, meaning I was out by 17%, not too

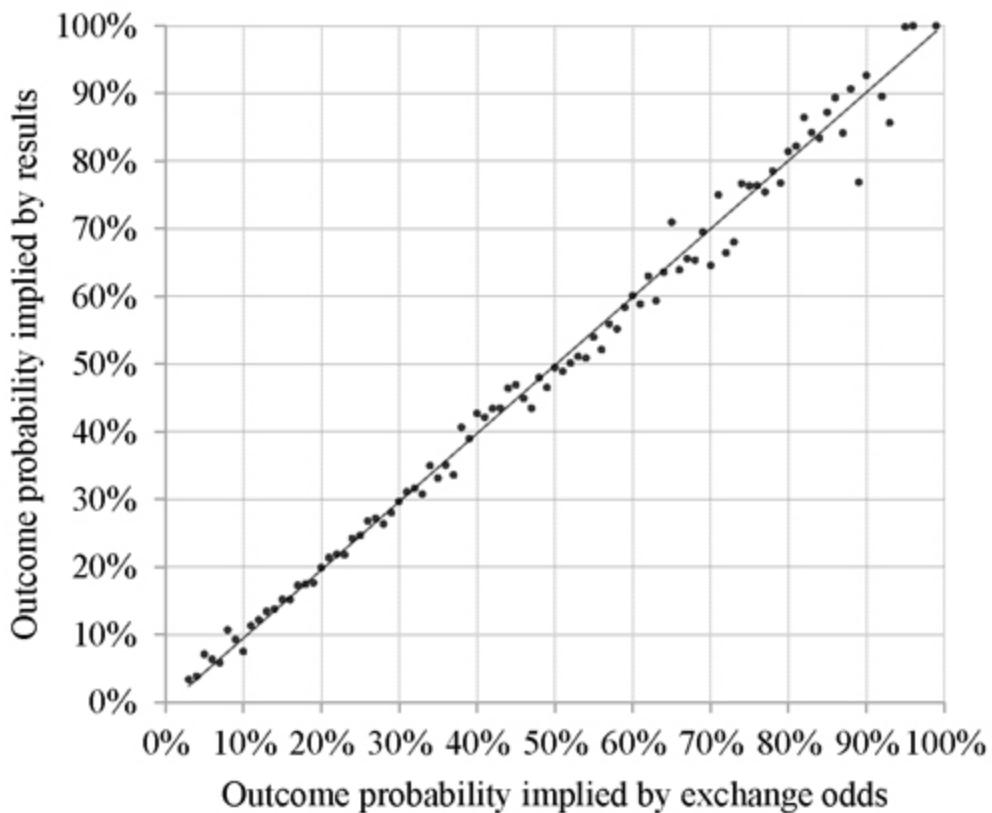
bad. Additionally, following my attempt there were a further 13 guesses, a total of 54, with an average of 272, just 8% lower than the correct answer. Seemingly, 54 completely unconnected people, acting independently of one another, collectively produced an estimate fairly close to the true figure. When you consider that the range of guesses was from 53 to 866, perhaps you'll agree that something magical is going on here. Not only was the final average pretty close to the actual number, it was also more accurate than the vast majority of individual guesses. That magic is called the 'Wisdom of the Crowd.'

The wisdom of the crowd phenomenon was first observed in the early 20th century by the eminent anthropologist, Sir Francis Galton (the man behind the quincunx machine). At a 1906 country fair in Plymouth, 787 people participated in a contest to estimate the weight of butchered ox. Galton calculated the median guess to be 1207 pounds, a figure accurate to within 1% of the true weight of 1198 pounds and again more accurate than the majority of individual estimates. In his book *The Wisdom of Crowds: Why the Many are Smarter than the Few*, James Surowiecki tells the story of the finding of the US submarine *Scorpion* that had disappeared in May 1968 whilst on a tour of duty in the North Atlantic. Whilst the *Scorpion*'s last reported location was known, no one had the slightest idea what had happened to it afterwards, with only a vague idea of how far it might have travelled. To solve the problem, naval officer Dr. John Craven assembled a team of specialists. Instead of organising a group consultation exercise, Craven asked each of them to independently submit their best guess as to what had happened structuring them as wagers with prizes for the best guesses. Everyone bet on why the submarine ran into trouble, its speed, heading, steepness of descent and so on. Individually, not one piece of information offered Craven any insight into where the *Scorpion* would be found, but collectively piecing it all together he located an area that turned out to be just 220 yards from where the wreck was finally discovered. Of course, we would do well to remember that these stories only became stories because they are winning ones. We're not party to all the examples where wisdom of the crowd fails, meaning we are potentially at the mercy of survivorship error again. Nevertheless, the wisdom of the crowd is a very real and repeatedly observable phenomenon in life, not least in the world of betting and investing, gambling markets that are dominated by player

psychology.

Remove the influence of the favourite–longshot bias and we find that betting markets actually do a phenomenal job of replicating the ‘true’ probabilities of outcomes, despite not knowing *a priori* what the results of sporting events will be. This is perhaps best observed at betting exchanges where the favourite–longshot bias is absent anyway. The chart below, based on 52,411 Betfair odds from worldwide football league matches during the period 29 October 2004 to 31 October 2005, compares the probabilities implied *a priori* by volume weighted average betting prices with the probabilities implied *a posteriori* by the actual results. There is an almost perfect correlation ($r = 0.995$).

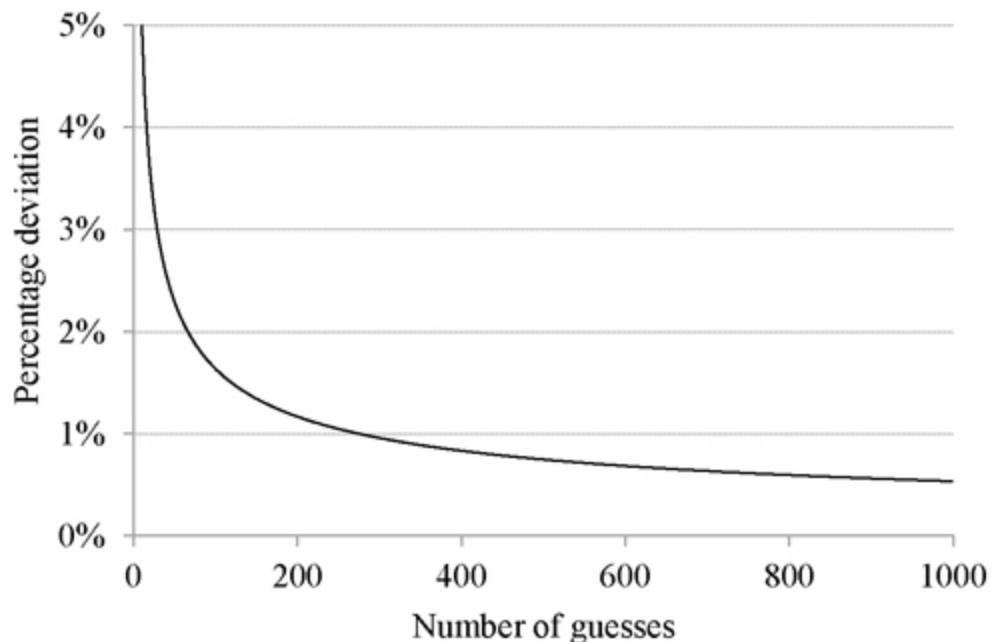
Comparison of outcome probabilities implied by Betfair odds versus actual results



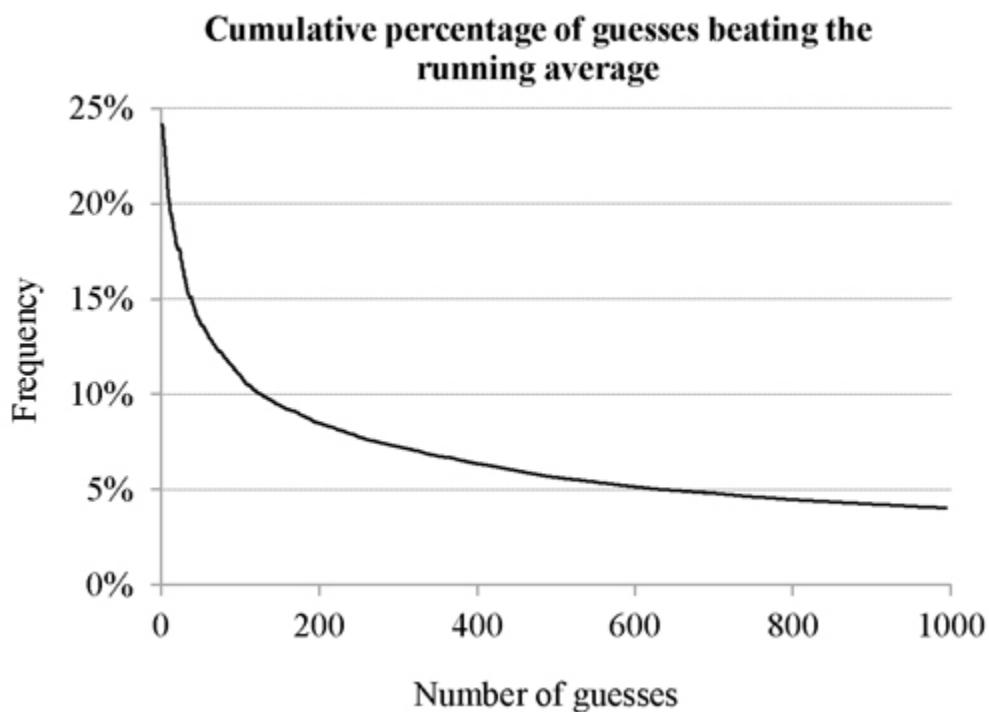
Let me illustrate, via means of a simple thought experiment, why this phenomenon might be explained by the wisdom of the crowd and why conceivably it presents a challenge for those who think they can beat the collective opinion of a betting market. Suppose I ask 1,000 people in turn to

estimate the weight of an ox, which I know to be 1,198 pounds. To simulate this little thought experiment, I've assumed people's guesses are normally distributed about an average that will match the true weight, with a standard deviation of 250. This value is just arbitrary. It means that about 68% of guesses will be within the range 948 and 1,448 whilst 95% will be within the range 698 and 1,698. Evidently, those present at Plymouth country fair when Galton made his observations were much better at estimating than my imaginary sample, since about 95% of estimates were within about plus or minus 100 of the actual weight. Nevertheless, these details don't influence what I want to show. After each estimate is submitted I recalculate the running average of all the estimates submitted to that point, creating a 1000-point time series. Unsurprisingly, as more and more people submit their estimates, the smaller is the difference between the current running average and the actual weight. Slowly, the collective average converges towards the actual value. This is the effect of the wisdom of the crowd. Repeating this scenario using a randomly generated 100-run Monte Carlo simulation, it is possible to smooth the data to estimate expected deviations between running averages and the actual value. The results from this exercise are illustrated in the first chart below.

Deviation of crowd average from true ox weight



After just 3 guesses, the average deviation from the true weight is about 10%. By 12 guesses that has fallen to just 5%. After 250 guesses, the collective guess is typically within 1% of the actual figure. For each individual, I also record whether his or her guess is closer to the actual weight than the current running average. Finally, I keep a running total of the number of people who manage to beat the average estimate at the time they submit their own. The second chart below shows how the percentage of people beating the running average evolves as more and more people submit their estimates.



Typically, after about 20 guesses, there are about 4 people with guesses closer to the actual weight than the current running average, or about 20% of the total. By 200 guesses, however, that number has only increased about 4-fold. And by 1,000 guesses it's barely doubled again. Self-evidently, it becomes increasingly harder for the next individual's estimate to be closer to the true weight than the current collective average. Of course, this is simply a consequence of the collective average becoming more and more accurate, that is to say, wiser. By the end, only 4% of people in total have managed to beat the running average at the time they submitted their estimate. Indeed, only 2.5% of the last 500 individuals were more accurate

and only 2% of the final 100. If I paid my imaginary respondents on the basis of whether they could beat the average, evidently not many of them would get paid, and proportionally fewer would do so the later they made their estimate.

This scenario is ostensibly very similar to what takes place in an evolving betting market. Imagine Barcelona is playing Real Madrid in the Champions League final. A first bettor thinks Barcelona will win the cup; another thinks Real Madrid will do it. Such a perfect balance of opinion implies fair betting odds of 2.00 for each side, in the absence of any other information about the match and any other individual expressing an opinion in the market. Now suppose a third bettor opts to back Barcelona. With 2 to 1 in favour of the Catalan side our Bayesian posterior probability will be 66.67% or 0.667 with associated odds of 1.50. Similarly, the Bayesian posterior probability for Real Madrid will be 0.333, implying odds of 3.00. These posterior probabilities now become the next prior probabilities in anticipation of new opinions expressed about the game. As each new bettor expresses an opinion, the betting prices for each team will fluctuate back and forth. Evidently, however, as more and more opinions are expressed, the posterior probabilities and their associated betting prices change less and less with each newly expressed opinion, with the betting odds gradually converging towards a collective (or average) opinion in accordance with the law of large numbers, in much the same way as for guessing the weight of the ox.

Obviously, when the match is played, only one team can win. In a classical sense all those who backed the winner were right and all the rest were wrong. In a quantum mechanical sense, however, the result doesn't change the underlying probabilities of outcomes, any more than opening Schrödinger's box to find his cat either dead or alive changes the probabilities of its death or survival. Both possibilities coexist at the same time, expressed in terms of a probability or odds. Play the match again, and again, and again; we won't always see the same outcome. Essentially the odds for both Barcelona and Madrid describe their probability of winning given an infinite number of possible path histories, that is to say, assuming the match could be played an infinite number of times. Of course, we can't play the match an infinite number of times, but we can count past outcomes and compare them retrospectively to their estimated odds. What we find, as

evidenced by the Betfair data, is that the collective opinion of bettors is, on average, a pretty good judge of the ‘true’ probability of outcome.

In a real betting context, a market is first opened by a bookmaker publishing his opening odds that represent an estimate of what he thinks the probability of outcome should be. Once his customers begin expressing their opinions via money, his odds will have to change to reflect the evolving collective opinion about the probability of the outcome. More people (and/or money) favouring one outcome will ensure the odds for it will shorten. Fewer people (and/or less money), and conversely the odds will lengthen. Clearly, early bettors have the potential to change the odds quite significantly, in the same way that early guesses of the ox’s weight shifted the running average more dramatically. But as an increasing number of bettors express their opinions, that average will fluctuate less and less. For guessing the weight of an ox, it becomes less and less likely one will be more accurate than the average. For estimating the ‘true’ probability of a betting outcome, similarly, if this model is right, it becomes less and less likely that later bettors will be wiser than the crowd. Betting, of course, is all about finding value in the odds. Consciously or otherwise, most bettors have an intrinsic acceptance that outcomes are in some way probabilistic and that we won’t win every time (even though individually many are not very good judges of probability). As such, they will have some idea of what constitutes a fair price (despite many maintaining that they simply bet on teams they think will win). This little exercise reveals that arguably the longer you wait to express your opinion, the harder it will be for you to find value in the betting odds. Of course, the bookmaker’s margin just magnifies the difficulty of that task. If you simply follow a crowd that’s collectively wise, you won’t have much chance of beating the odds.

Of course, one outstanding question remains. What is it about a crowd that often makes their collective opinion so accurate? Earlier in the book I examined the expression of systematic biases, and how they could lead to irrational outcomes, for example financial bubbles and crashes due to greed, fear and groupthink, or over-betting as a consequence of the possibility effect and the hot hand fallacy. Undoubtedly, humans have a capacity for irrationality and bias. Nevertheless, most people, most of the time and in most environments tend to behave mostly rationally, or at least not irrationally enough to allow sharper players to exploit them profitably,

where the costs of playing must be considered. An obvious example is favourite–longshot bias, very real but not sufficiently strong to practicably assist the contrarian bettor. Imagine what would happen if squares could be consistently exploited. Remember the \$100 bill? Furthermore, when errors are not systematic they will tend to cancel each other out, just as good and bad luck does. Each individual guess has two components: signal (information) and noise (error). Remove the random noise and what's left behind is the signal, that is to say, the collective wisdom. James Surowiecki proposes that four conditions are necessary for a crowd to be wise in this manner: diversity, independence, decentralisation and aggregation. I'll briefly look at each in turn. Arguably they are all present in many betting markets, and most of the time in financial markets as well.

Having a diverse set of opinions is a prerequisite for collective wisdom. Where everyone is thinking or doing exactly the same thing, the probability of systematic error or bias increases. Diversity is the basis for any competitive market; let different ideas or products compete against one another and that's usually a recipe for the best ones succeeding. Google inherently understand the significance of diversity; it invests a lot of time, money and effort into lots of little start-up ideas like Google Earth, Google Glass and a driverless car. Not all of them will succeed but by having eggs in lots of baskets, you increase the chances that at least some of them hatch. In prediction markets like betting and investing, diversity can arise because of the environment of uncertainty and the typically large number of people acting in them with different opinions, risk preferences and approaches to forecasting. When attempting to forecast the outcome of a game or the future price direction of a stock, for example, there are potentially limitless ways to skin that cat. Some prediction methods try to determine the intrinsic probability of outcome (for example value betting) or fundamental value of an asset (firm foundation theory). Others adopt a more psychological approach, believing a market to be more a reflection of opinions (and more importantly opinions about opinions) with all their biases, making use of methods such as technical analysis and charting to study trends and directions in prices and odds. Then there are different types of prediction models: linear or non-linear, static or dynamic, deterministic, probabilistic, or dynamic. And that's just for starters; the diversity in information that can be inputted into a prediction model is similarly immeasurable. The

mathematician George Edward Pelham Box once quipped that “*Essentially, all models are wrong, but some are useful.*” In doing so he perfectly captured the essence of collective wisdom and the importance of diversity. Indeed, most prediction models will be wrong, but the pooling of diverse ideas encourages collective accuracy. Paradoxically, a larger crowd with a greater percentage of squares will often prove to be collectively wiser than a smaller group with a larger percentage of smarts. Given that prediction markets would appear to be largely devoid of expertise, it’s perhaps no small wonder that they are often so collectively accurate.

Diversity arises out of independence of thought. Arguably, this is the most important ingredient for crowd wisdom. If everyone thinks the same way and does the same thing we frequently end up with poor outcomes. Everyone was betting on the 2015 UK General Election to return a hung parliament, because that’s what all the polls and pundits were saying was going to happen. Evidently, they were all doing the same thing and failing to take account of a couple of very important influences: the shy Tory effect and the lazy Labour voter. (Of course, it’s easy to say this with hindsight.) Perhaps if more polls and more pundits had demonstrated a greater independence of thought using what economists call ‘private information’, such ideas (and others) could have been used to collectively improve their predictions. Of course, this isn’t always the case. Teaching people to serve a tennis ball or to do differential calculus requires a narrow field of learning through repetition. But such activities are sufficiently predictable with clear relationships between cause and effect. In prediction markets under uncertainty, by contrast, learning through pattern recognition is limited because the patterns are largely random. What signal exists is deafened by noise, with good and bad luck dominating outcomes. In such environments having people acting independently helps to eliminate that noise, because it offers the best chance for keeping people’s errors from becoming correlated. When mistakes are random they will cancel out. As Surowiecki explains, however, independence does not imply rationality or impartiality. “*You can be biased and irrational, but as long as you’re independent, you won’t make the group any dumber.*” We’ll see shortly that this point is absolutely fundamental to defending the efficient market hypothesis against accusations that it is inherently flawed because investors and gamblers act irrationally.

The final two pieces of the jigsaw that make the wisdom of the crowd such a powerful mechanism are decentralisation and aggregation. A system is said to be decentralised if it's not acting under the influence of a top-down central authority. By definition, independence and diversity of thought and decision making will be encouraged where central regulation is not restricting outputs. One just has to look at the disappointments of top-down communism and the achievements of bottom-up market capitalism to understand the significance of decentralisation. Indeed, capitalism as a global movement inspired by the writings of the 17th century economist Adam Smith and other liberal-thinking philosophers of that era is *prime facie* an expression of the wisdom of crowds. Decentralisation ensures that a crowd of self-interested, independent people working without top-down interference will collectively find a better solution than anything else you could come up with. The process happens as if by magic. It's the mechanism behind bird flocking, fish shoaling and insect swarming, the emergence of complex and seemingly coordinated behaviour out of a few simple rules followed by the self-interested individuals. In the case of birds there are just four: stay close to the middle; keep sufficient distance between neighbours; avoid collisions; and flee predatory attack. For human interactions, Smith labelled this magic the '*invisible hand*,' describing the unintended social benefits resulting from individual actions. Indeed, arguably the whole business of complex moral behaviour emerges out of a few simple rules of self-interest, most importantly the Golden Rule: treat others as one would like others to treat oneself. Those who wish to burden markets with excessive top-down regulation would do well to remember this point. Leaving people to their own devices usually (although, of course, not always) delivers better outcomes.

Decentralisation, however, will only be of benefit if there exists a way of coordinating or aggregating all the information. For gambling, betting and investing, that aggregation process is explicit: the conversion of private information and expression of opinions into a piece of public property – the price. The odds for a football team or the share price of a company publically aggregate all the private information that exists. For a betting market, the odds represent the current balance of opinions about the likelihood of a team winning as expressed by the amounts of money wagered for and against it. In a financial market, the share price represents

all there is to know (or feel) about the future prospects of a company's cash flow, as expressed by the number of shares bought and sold. The price reflects the actions of all buyers and sellers, backers and layers, throughout the market.

The magic lies in the emergence of wisdom without individuals having a complete understanding of what the market is doing and without anyone knowing what the 'true' answer, if there is one, will be. Ask a thousand people how many runs the England cricket team will score in their 4th innings run chase against New Zealand in their 2nd test at Headingley. Few, if any, will be perfectly correct, but the collective average will be more accurate than nearly all of them. As Surowiecki says, people with only partial knowledge and limited calculating abilities actually arrive collectively at the right answer. A wonderful demonstration of this magic was accomplished by Nobel Prize-winning economist Vernon Lomax Smith¹⁰². In 1956, he set out to determine whether people with limited information would conform to the hypothesis of market clearing, where prices of traded assets adjust up or down such that quantity supplied at the market-clearing price equals the quantity demanded at the market-clearing price. Such a price is also called the equilibrium price. Giving his 22 students cards with a dollar price tag, he made half of them buyers and half of them sellers. The sellers were instructed not to sell at less than this price, whilst the buyers were instructed not to buy at more than their card value. A difference achieved between card value and actual contract price could be regarded as profit for the player. Strict anonymity was applied such that no one knew the value of anyone else's card. The students were then asked to start trading, calling out bids and offers which may, or may not, be accepted. In Smith's own words, the buyer holding the \$2.50 card might call "buy at \$1.00" whilst a seller with the \$1.50 card might shout "sell at \$3.60." Sellers and buyers were free to accept a bid or offer. If they were refused, further price compromise or bartering would be required until they were accepted. The successful trades were recorded publically on the classroom blackboard.

Economic theory was matched by reality. Traded prices quickly converged on one price, the equilibrium price or what we might also call the expectation price, despite players being completely unaware of their

competitors' demands and despite none of them preferring this outcome (self-interested traders after all want more profit). Smith also showed that collectively the convergence on the market-clearing price yielded the best possible outcome, even if some of the players had been blessed with additional knowledge telling them how they should trade. The brilliance of Smith's experiment was that it demonstrated that, for markets under uncertainty, imperfect people could collectively produce near-perfect outcomes. What allowed it to happen was a decentralised independence of action and the aggregation of privately anonymous information via the publication of a price.

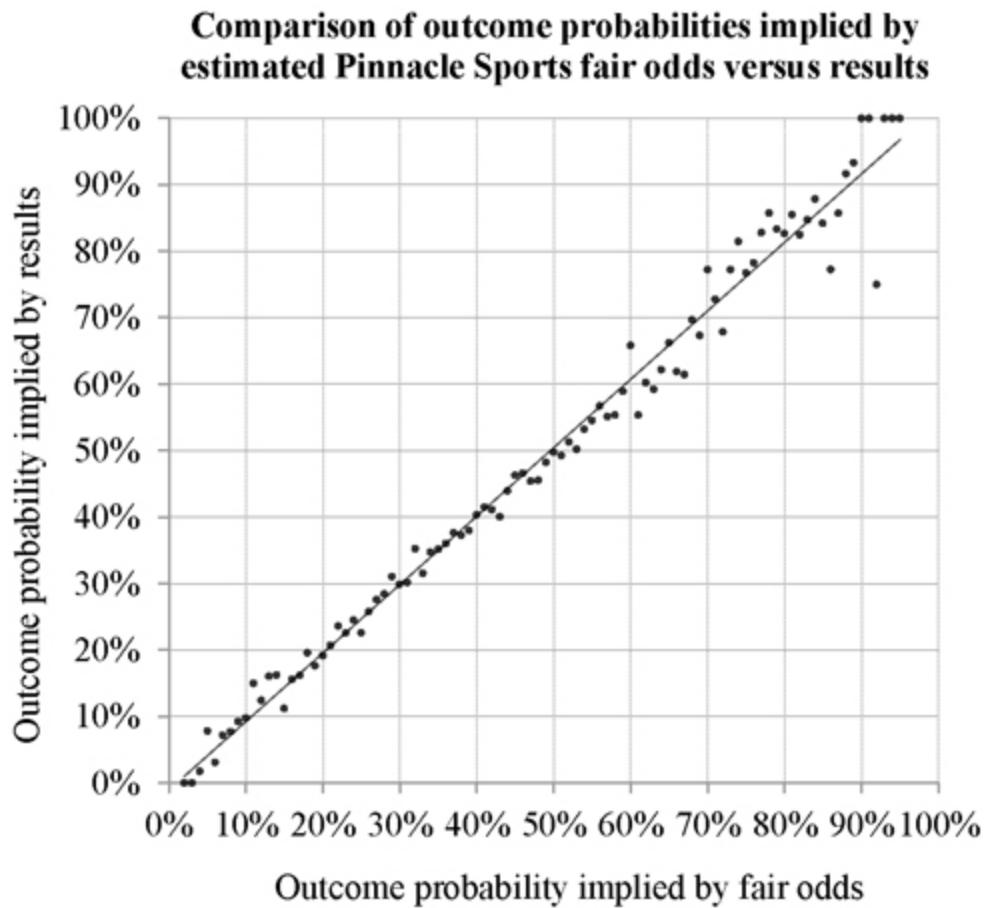
Essentially, price convergence is exactly what happens at a betting exchange. Nate Silver (in *The Signal and the Noise*) sees the invisible hand as a kind of Bayesian process in which prices are continually and dynamically updated to reflect changes in supply and demand. As such, it represents a consensus-seeking process taking advantage of the wisdom of crowds. At a betting exchange, odds move simply in response to supply and demand, just like in a financial market. The market maker sits completely outside the contest, skimming his commission percentage from the action. This process is otherwise known as 'price discovery', a mechanism for determining the price of an asset in the marketplace through the interactions of buyers and sellers, or in this case backers and layers. Remarkable as it may seem, the betting public collectively 'knows' the 'true' probability of outcome of a sporting event through its betting actions. Odds shorten on the fancied competitors and lengthen on the least fancied in a kind of Bayesian dance, settling at values that reflect all the private information that has been consumed by the players. That dance is dynamic with the equilibrium price never completely stationary because there will always be new information arriving on to the market.

For a bookmaker, things are a little different but only because they are part of the action; the fundamental process remains the same. Odds shorten because too much money has been bet on one outcome, giving the bookmaker a large liability in the event that it happens. Bookmakers are always looking to reduce their liabilities; in this case they can achieve this by shortening the odds to discourage further interest from customers. At the same time they lengthen the odds on the opposition to attract money. Through this Bayesian price clearing process they can try to balance their

books. If they get it right they won't care which team or player wins, and in effect they become more like an exchange. Some traditional bookmakers nevertheless still prefer to take some sort of position on an event, and they do this by offering attractive prices that possess positive value expectation relative to the collective market and by refusing to drop those prices when others around them are doing so. Frequently, they are then exposed to some risk on the side of the book that has attracted a disproportionate level of action. For them there are other methods of managing liability. I've already discussed the use of customer restrictions in this context. Another option is to lay off the risk at a betting exchange or another bookmaker with a smaller margin. One bookmaker that looks to take fewer positions on games than most other brands is Pinnacle Sports, which instead relies on professional odds management algorithms, allowing the market to make up its own mind. With its small margins and laissez-faire approach echoing Adam Smith's invisible hand, the brand has become synonymous with high-volume action. Such is the significance of money talking at Pinnacle Sports that the phenomenon has been given its own name in betting circles – the Pinnacle Lean. It's often the case that, when Pinnacle's odds start to lean one way, other bookmakers will soon follow (assuming, of course, they aren't too engaged in their own market interference). Of course, there is one significant consequence of Pinnacle's market being wiser than all the others: it makes it much harder to beat. It might have the best odds, but that's not the same thing as having the most profitable opportunities.

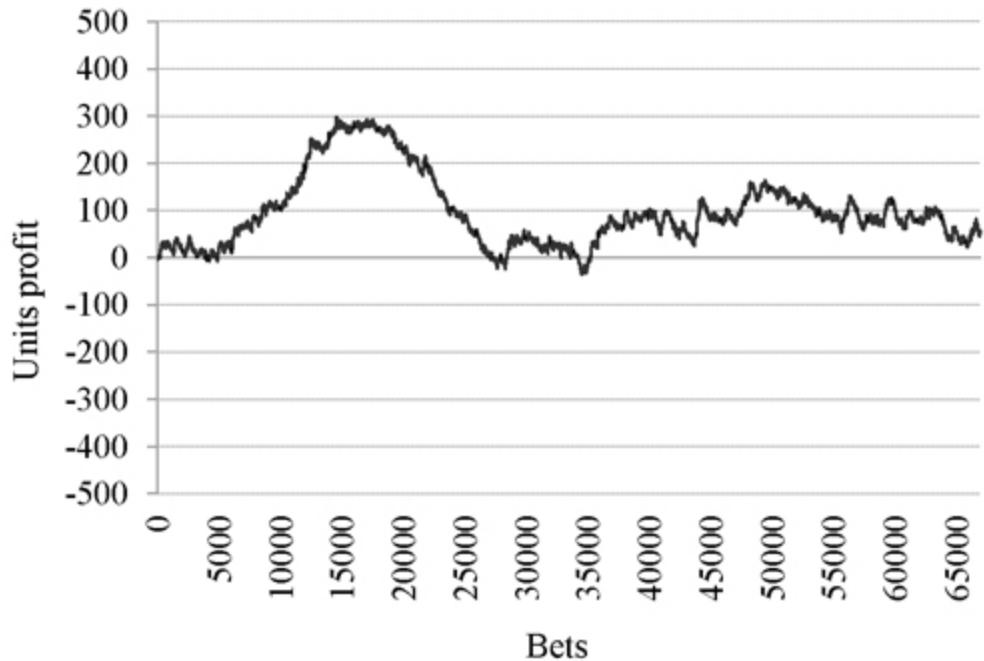
We can, however, use Pinnacle's market wisdom to estimate what the 'true' chances of a result might be. To do this we simply need to remove the influence of the profit margin and the favourite–longshot bias that Pinnacle applies to its odds, using the model I presented earlier for estimating fair odds from published ones. To test the accuracy of these fair odds estimates, and by implication how wise the Pinnacle Sports betting market really is, we can then retrospectively compare them to the actual results. I have done this for the home-draw-away football match betting market using three seasons (2012/13 to 2014/15) of European domestic league football – a total of 22,318 games. As for the Betfair exchange data I showed previously, my model estimates for Pinnacle's fair prices do a pretty good job of predicting actual outcome frequencies. Another way to test how wise these fair prices have been is to see whether they would have broken even if all of them had

been bet blindly. Fair prices, by definition, should break-even over the long term, allowing for shorter term periods of good and bad luck to even out.



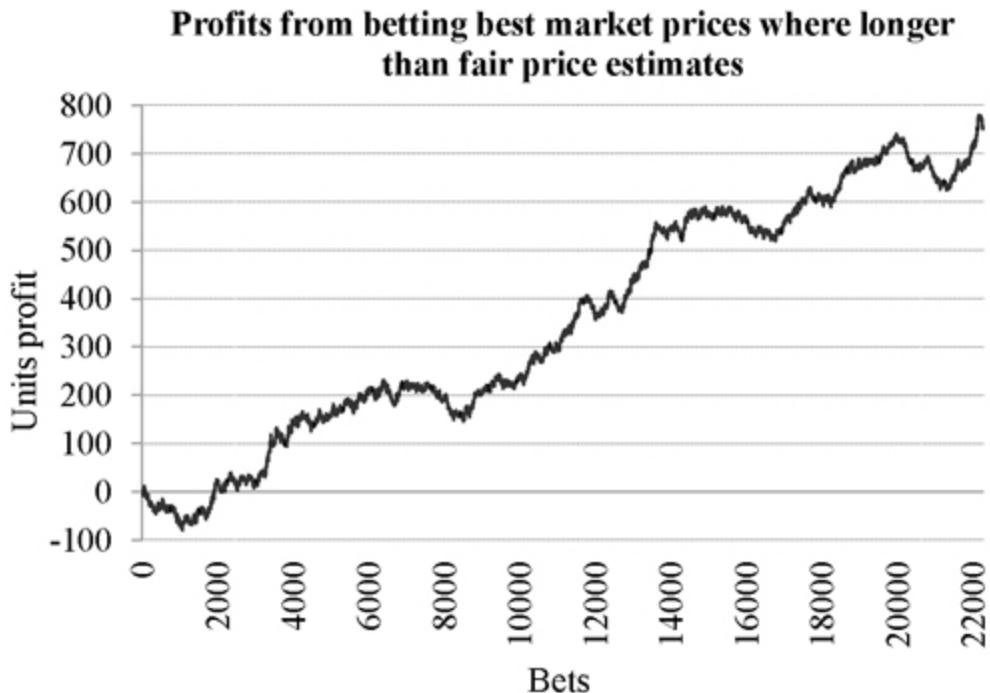
The next chart shows the evolution of profits from level stakes betting on all matches to these hypothetical fair prices.

Profits from betting Pinnacle's fair price estimates



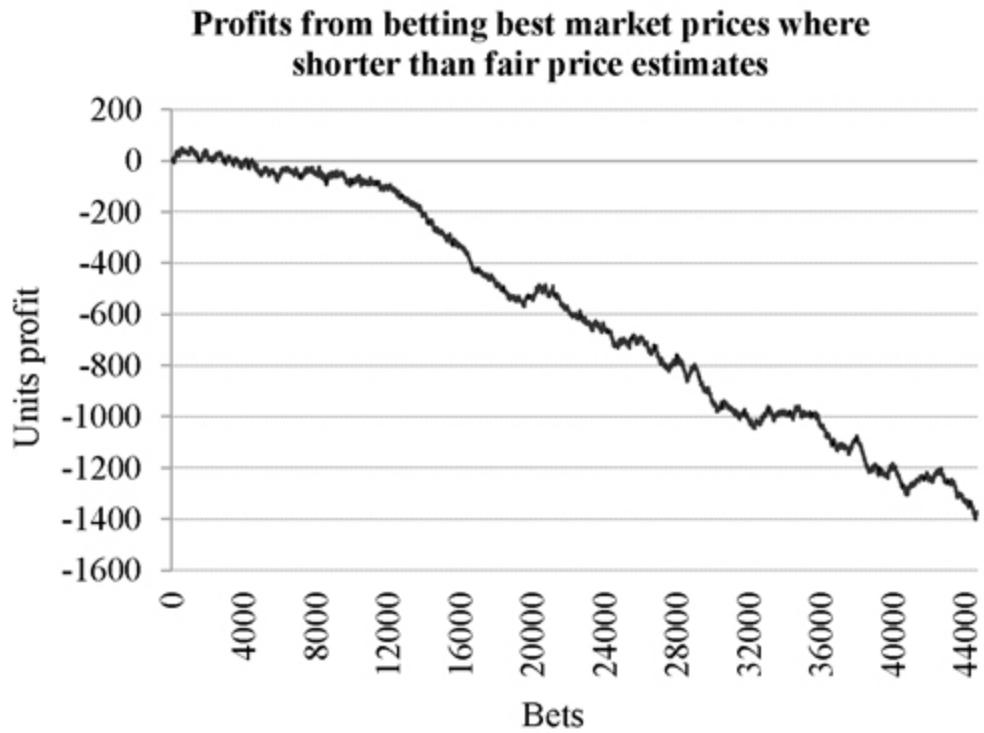
In fact, the closing yield from these 22,318 matches (66,954 home, draw and away bets in total) is 0.08%, as close to break-even as we might reasonably expect.

An obvious question now arises: if Pinnacle Sports' football home-draw-away betting market is so accurate can we use its wisdom to identify mistakes elsewhere, with a view to potentially making a profit? The answer, it would appear, is yes. Alongside the betting odds for Pinnacle Sports, I have also recorded the best market prices (as published by the odds comparison Betbrain.com). Betting every home, draw or away price where the best market price was longer than the fair Pinnacle price (as estimated by my model) gave a level-stakes yield of 3.4% from 22,281 bets with the following profit trend.



In fact, this was a little better than one might expect given the prices that were available. The average advantage over fair prices for this sample was 2.2% from which *a priori* one would expect to see a similarly-sized yield. Those best market prices came from theoretical best books with an average overround of 100.4%. In other words, if we had blindly bet all possible outcomes with appropriate staking we would have lost about 0.4% on turnover. Indeed, over two-thirds of these value opportunities were found in books that were still overround at best prices. Evidently, Pinnacle's wise betting market can help pinpoint when something somewhere else is overpriced, without requiring the availability of arbitrage opportunities.

In contrast, how would we have performed had we decided to back best market prices that were shorter than the model-estimated fair prices? This time our level-stakes yield from the remaining 44,673 bets would have been -3.08% (with an average disadvantage against fair odds of 2.4%). The profit evolution is again shown below.



Of course, we could choose to be more selective with our betting criteria. We might, for example, just decide to bet when our advantage over the fair odds is greater than 1%, 2%, 3% or higher. Naturally, this will reduce the number of betting opportunities available, but in theory it should increase the yield we will achieve. Would this have happened for this 3-season sample? Yes; the table below shows how.

Advantage over fair odds greater than...	Bets	Yield	Average advantage over fair odds
0%	22,281	3.40%	2.23%
1%	14,837	4.60%	3.07%
2%	9,196	6.63%	4.04%
3%	5,474	8.83%	5.11%
4%	3,243	12.70%	6.26%
5%	1,927	13.22%	7.51%

Clearly, for each sample, performance was better than would be predicted from the theoretical advantage gained over the fair odds. Presumably this is simply a result of good fortune. The right hand column figures are probably

more representative of what we should expect to achieve by way of yields. Nevertheless, it would appear that the wisdom of Pinnacle Sports' home-draw-away betting market, coupled with this rudimentary model at estimating fair prices from it, can provide profitable betting opportunities at bookmakers more prone to offering mistakes.

Naturally, there are a couple of caveats with this approach. Firstly, given the relatively small yields involved, one should reasonably expect to suffer fairly long periods of treading water, or worse still, losing, lasting hundreds and perhaps thousands of bets. Secondly, it is to be expected that the sort of bookmaker who will offer betting prices in excess of Pinnacle Sports' fair price estimates will also be the sort of bookmaker who won't like a customer consistently exploiting such generosity, for reasons I discussed earlier. However, this examination has at least identified that a 'wisdom-of-the-crowd' approach can identify where some bookmakers have 'made' mistakes, and that technically, if not necessarily always practicably, it should be possible to exploit them. I have repeated the same analysis for the tennis match betting market and found much the same crowd wisdom exhibited by Pinnacle Sports' betting customers.

Much of the preceding discussion has focused on the fixed odds betting market. What about point spreads? Surowiecki is clear that, with about half the favourites and underdogs alike covering the spread, this necessarily implies that it "*ends up representing bettors' collective judgment of what the final outcome... will be,*" with the bookmaker purposely equalising the quantity of money wagered on each side. Essentially, the point spread is a kind of market-clearing price. In 2004, the American economist Steven Levitt¹⁰³ challenged this viewpoint. Point spread bookmakers, Levitt argued, do not, in fact, play the traditional role of market makers, matching buyers and sellers but, rather, take positions with respect to the outcome of games, systematically exploiting bettor biases by setting spread prices that deviate from the market clearing price. The primary bias in question is an over-betting of favourites, that is to say, teams to which a negative point handicap has been applied. Analysing the actual number of wagers placed by bettors as part of a handicapping contest offered at an online sports book during the 2001 NFL season (285 entrants making 19,770 bets on 242 games), Levitt found that over 60% of bets were placed on the favourite.

Such a systematic bias towards favouring a point spread favourite argues against the hypothesis that bookmakers are doing the best they can to even out the bets on each game. On the contrary, with only 48% of favourites covering the published spreads in the preceding 21 NFL seasons, Levitt speculated that, despite the additional liability of unbalanced markets, bookmakers are intentionally biasing spreads very slightly against the favourite, thereby taking a position with respect to the outcome of a game, in order to exert their superior talent with a view to yielding greater profitability over the long term. The extent to which they artificially move the spread, however, is constrained to prevent smarter bettors from exploiting such a price distortion. The implication from Levitt's research is that point spreads do not represent the crowd's prediction of game outcomes and, therefore, that bettors in such a market are collectively unwise.

It is nevertheless worth noting that Levitt's data concerned the number of wagers rather than actual volumes of money staked; his data, after all, came from a competition to pick the largest number of winners. A bias in the number of bettors backing favourites to cover a spread may not necessarily be equivalently expressed in monetary terms. Bettors, of course, don't all bet the same stake. Levitt, furthermore, failed to offer any explanation for why bettors might be predisposed towards favourites on a point spread market, merely observing that such a systematic bias exists. In contrast, Joseph Simmons, associate professor at the University of Pennsylvania's Wharton School, and Leif Nelson, associate professor at the Haas School of Business at the University of California, Berkeley, have provided a psychological account for where the bias originates from: intuitive confidence¹⁰⁴. Tracking data from thousands of predictions of 850 professional and college football games on Yahoo.com for the 2003 and 2004 seasons, Simmons and Nelson found, as Levitt had previously, that bettors prefer to back the favourite. 65% backed an NFL favourite to cover the point spread. This rises to an aggregate of 70% when college games are included as well. Crucially, the more people believed a certain team would win, the more likely they were to also choose that team to cover the spread. In other words, the intuitive confidence bettors felt in picking the winner translated into an unrelated belief that the winner would cover the spread. By contrast, the weaker the intuition (for example, where two teams are

considered to be more evenly matched), the weaker the intuitive bias. The researchers proposed that such intuitive bias arises because intuitions often spring to mind with cognitive ease (particularly when one team is considered to be much better than the other), leading people to hold their intuitions with high confidence. As such, this would appear to represent an obvious example of Kahneman's attribute substitution, where an individual has to make a judgment of an attribute that is computationally more complex (determining which team will beat the point spread), and instead substitutes a more easily calculated heuristic attribute (which team will win).

As for Levitt's study, however, Simmons and Nelson's initial research investigated only the percentage of bettors backing favourites, but did not consider actual money wagered. In a follow-up study¹⁰⁵, the authors attempted to address this by tracking the wagers bettors submitted to a website created for the purposes of the experiment. The 178 participants were asked to assign one of five possible wager amounts (\$0.50, \$1.00, \$1.50, \$2.00, or \$2.50) to each of the 226 Sunday games during the 2007 NFL season. For every game, the size of the point spread presented to the participants had been artificially tilted in favour of the underdog, meaning favourites lost (124) more games than they won (98). Participants were also assigned to different experimental conditions. Some were asked to predict which team would beat the spread. Of those, about half were actually informed that the spread price had been manipulated. Others were instead asked to simply predict the match winner, along with an estimation of the actual point differential of the game. Simmons and Nelson found the same bias towards point spread favourites as in their initial research. For nearly 90% of games, participants uninformed about the price manipulation wagered more than 50% on the favourite. That figure was only marginally smaller (83%) for those who were told that the spread had been increased. Frustratingly, however, the authors haven't reported total volumes of money wagered on favourites versus underdogs. Nevertheless, their findings are indicative of an unwise crowd, even when given added hints about teams that held value.

Perhaps more intriguingly, however, those bettors asked to predict a match winner and point differential elicited superior betting choices. By

converting the average predicted point differential for each game into predictions against the point spread, the authors found that this crowd more wisely (given the manipulated spreads) predicted the underdog 83% of the time. As the authors suggest, paradoxically it seems that, asking people to **estimate** the point differentials directly may cause them to focus on the very dimension (the point differential) that receives insufficient weight when making cognitively simpler **choices**. These findings would appear to support the authors' theory of intuitive confidence. When forced to apply System 2 to a more cognitively taxing task (estimating the actual point differential) bettors are less likely to make lazy and biased judgements based on intuition. Thus, the authors conclude that "*although systematic biases may ruin the crowd's judgments when judgments are elicited in a manner that encourages intuitive responding, those biases may be absent from logically identical methods of eliciting the same information, and the crowd may emerge wiser.*" It would appear, then, that the presence of crowd wisdom is highly context-dependent.

If Simmons and Nelson are right, the implication would be that moneyline and fixed odds markets, where bettors are faced with a simpler task of predicting a winner only, will more readily exhibit crowd wisdom, on the grounds that there is less scope for attribute substitution. Presumably, the same should be true for total goals and total points markets. Yet Pyckio has reported a distinct preference by bettors for 'overs' in comparison to 'unders', with a 70:30 split from a sample of over 100,000 betting picks. Furthermore, the 'unders', in aggregate, experienced significantly smaller losses as a percentage of turnover, a feature that would imply a supply side price manipulation akin to the favourite–longshot bias in fixed odds and moneyline markets, as well as that proposed by Levitt for point spreads. If bettors are making systematically biased judgements in a simple market like this, surely attribute substitution cannot provide an explanation. What, then, might be? One possibility could be straightforward loss aversion.

Consider the nature of an over/under bet. By its very structure, at the start of a match, one side, the 'under', is automatically a winner, and remains so until such time as enough goals (or points) are scored to transform it into a loser. In contrast, the 'over' starts out as a loser and remains so until such time as enough goals (or points) are scored to transform it into a winner. Given everything that prospect theory has taught us about the psychology of

winning and losing, bettors, conceivably, prefer the possibility of seeing a loser transformed into a winner rather than the other way around. If losing hurts about twice as much as winning is enjoyed, the 70:30 bias in favour of ‘losers → winners’ over ‘winners → losers’ would seem to make sense. Daniel Mateos, Pyckio’s co-founder, revealingly explains why he, too, suffers from the ‘overs’ bias.

“I don’t want to be caught, that is, if you go for the ‘overs’, you feel positive feelings with every goal; if you go for the ‘unders’ you feel negative feelings with every goal. It’s stupid, but that is how I feel, and I am supposed to be a more rational bettor than average.”

Daniel’s account brings back memories of the stresses I suffered watching televised games for which I had backed the ‘unders’ option on the total corners market, something I played frequently many years ago.

We might also speculate whether the influence of loss aversion offers an alternative explanation to intuitive confidence for the bias towards favourites in the point spread market. By definition, bets on a spread favourite start games as losers, whilst those on underdogs begin as winners. Of course, the evolution of the result is potentially more complex than for an over/under wager, where changes in the score could repeatedly transform losers to winners and back again several times. Nevertheless, perhaps the key feature is how the bet starts its journey once the game begins. Mirroring Daniel’s thoughts, every goal or point scored by a favourite will induce positive feelings in the bettor who’s backed them on the spread, whilst conversely inducing negative feelings in the bettor who’s backed the underdog. The same, of course, could be said of scores made by the underdog, but given that favourites, in general, score more than underdogs, this would account for why the majority of bettors prefer to chase a handicap than to defend one. One might suppose loss aversion could impact the home-draw-away market in football, too. By definition, ‘homes’ and ‘aways’ start life as losers, compared to draws which begin as winners. Traditionally, reasons for why bettors dislike backing the draw option include the fact that it is seen as boring and random. Perhaps, instead, it might simply be disliked because it is seen as something to defend.

Without data on real money wagers placed at the bookmakers, much of this discussion on crowd wisdom (or its absence) in sports betting markets must surely remain somewhat speculative. Of course, given the commercial

sensitivity of such information, no bookmaker is going to release it, for doing so would potentially reveal clues as to how it sets and manages its lines. If nothing else, however, both the favourite–longshot bias (which I discussed earlier in the book) and the point spread bias reported by Levitt, Simmons and Nelson, suggest how easy it might be for a betting crowd to be intentionally deviated away from wisdom by active market makers looking to exploit bettors’ propensity to exhibit a bias. Give bettors an opportunity to make biased judgements and it appears that, often, they will. By comparison, where a market maker is passive, for example, a betting exchange, the crowd will more readily remain wiser.

Discounting the presence of these systematic biases, whose expression may well be the result of supply side manipulation and not sufficiently exploitable anyway, wise markets, it would appear, are designed to make pretty good predictions about uncertain futures. Indeed, according to Nate Silver, that is exactly what the stock market is: a series of predictions about the future earnings and dividends of a company. Surowiecki tells the story of how, within minutes of the space shuttle *Challenger* disaster in 28 January 1986, the stock market had begun dumping stocks of NASA’s four major shuttle contractors, in particular Morton Thiokol, which built the solid-fuel rocket booster. Six months later, a Presidential Commission on the disaster revealed that the O-ring seals on the boosters had failed due to the cold weather and Thiokol was held liable for the accident. A subsequent study into the swift market reaction concluded that insider information was probably responsible for the swift collapse in Thiokol’s share price. Yet Surowiecki sees it differently. On the contrary, he considers it equally plausible that the reaction was simply a manifestation of the wisdom of a diverse and decentralised collection of independent traders aggregating privately held opinions.

But why stop at betting and investing? Wise prediction markets could also be used to actively form and implement political and economic policy. Why rely on a narrow field of expertise (for example the UK’s monetary policy committee in the setting of interest rates) when arguably a larger crowd of less informed individuals would do a better job? Information or decision markets, speculative markets created for specifically the purpose of making social, political and economic predictions, do just that. One of the most well-known is the Iowa Electronic Markets (IEM) project,

operated by the University of Iowa's Tippie College of Business, an online not-for-profit futures market where contract payoffs are based on real world events such as political election results and economic indicators. The Hollywood Stock Exchange offers a virtual market to trade in the success of films, actors, directors and more, including even Oscar nominations and awards. The Policy Analysis Market (PAM), part of the FutureMAP project, was a proposed futures exchange developed by the United States' Defense Advanced Research Projects Agency to assist with the prediction of political developments in the Middle East. In 2003, the idea was shelved in the face of opposition questioning the morality of allowing trading in events such as coups, assassinations and terrorist attacks, arguing further that money would be better spent on 'real world' intelligence in the hands of experts. Well, more than a decade on, where has this 'real world' intelligence led us?

Despite obvious shortcomings, the wisdom of the crowd is a beautiful idea to explain the mystical accuracy of a collective prediction. It works best in domains of uncertainty lacking obvious predictive patterns that yield to learning and feedback, where individuals are independent, opinions diverse, coordination decentralised, where action can be aggregated, where there is potentially a 'correct' or 'true' answer, even if in a probabilistic sense, and where forecasters can avoid expressing intuitive confidence or experiencing attribute substitution. Social influence or social proof, through herding and groupthink that encourages individuals to participate in information cascades, rejecting private information in favour of what others believe in or are doing, as well as other heuristic biases that lead to non-rational behaviour, have the potential to undermine its applicability. Nevertheless, money markets, including many betting and financial ones, are arguably less affected than other social systems simply because money is the objective currency of self-interest. If the market was collectively unwise, smart people would then consistently be able to exploit it for profitable ends. Why then, as Nate Silver puts it, have a lot of smart people failed miserably when they thought they could beat the market? The corollary of collective wisdom is a market that is said to be 'efficient.' This is a word I've used quite a lot already. It's time to put some flesh on the bones.

The Efficient Market Hypothesis

Market efficiency implies that buyers and sellers (or backers and layers) have all the information they need to agree a price, and as a consequence it is impossible to make predictions that will beat the market because prices exhibit a random walk. The market, essentially, is the ultimate expression of the paradox of skill. Buyers and sellers all have private talents at predicting the future, talents that manifestly have improved over time, but which collectively cancel each other out. When the relative difference in skills between buyers and sellers is zero, all that's left behind is statistical noise, the random dancing of prices or odds. The 'Efficient Market Hypothesis' is the embodiment of Ginsberg's theorem.

The origins of the hypothesis can be traced to the Ph.D. thesis of a French mathematician Louis Bachelier. In his *The Theory of Speculation* (1900), Bachelier was the first person to model a random process. In applying his mathematics to evaluate the price movement of stock options, he concluded that, since trading is a zero-sum game with winners and losers, rational self-interest will inevitably push the price at which it takes place to the 'right' number for both parties. Essentially, this is exactly what Vernon Smith observed with his price-clearing experiment. Furthermore, since prices change because of unexpected news, Bachelier deduced that it must be impossible to predict price changes. News, by definition, must be random and unexpected. If it was systematic and predictable it wouldn't be news.

Bachelier's ideas were largely ignored until after the Second World War, when a British statistician named Maurice Kendall developed the random walk hypothesis in his 1953 paper *The Analytics of Economic Time Series, Part 1: Prices*, in which he suggested that the movement of shares on the stock market was random with prices as likely to go up on a certain day as they were to go down. As for spins of a roulette wheel or rolls of dice, a random process is one where the probability distribution of the next state depends only on the current state and not on the sequence of events that preceded it. Random processes are memoryless. The idea was later popularised by Burton Malkiel in his series of books entitled *A Random Walk Down Wall Street*. Malkiel, an economics professor at Princeton

University, tested how easily people can be fooled by random walks. His students were given a hypothetical stock initially worth \$50. The closing stock price for each day was determined by a coin flip. If the result was heads, the price would close a half point higher, but if the result was tails, it would close a half point lower. Cycles and trends were revealed and Malkiel showed them to his friend who was a chartist, someone who seeks to predict future movements by interpreting past patterns (a form of technical analysis) on the assumption that history tends to repeat itself. The chartist told Malkiel that he needed to buy the stock immediately.

The concept of market efficiency was developed by Eugene Fama, who we met earlier when investigating the evidence for skill amongst mutual fund managers, with the publication of his Ph.D. thesis in 1965¹⁰⁶. Fama set about testing the random walk hypothesis by, first, looking for independence in successive price changes and, secondly, by testing whether the probability distribution to which these price changes belonged was one built on randomness. Studying the returns of mutual funds over the period 1950 to 1960, he found exactly what he was looking for. Yes, over the period the mutual fund investments saw a gross return of 14.1%, but this compares to the market benchmark of 14.7%. Furthermore, the funds that performed well one year were no more likely to beat the competition in the next; no fund performed consistently better than any other. Finally, their price movements conformed to a quasi-normal distribution that implied little else other than chance was at work. In the words of Fama:

“[A] situation where successive price changes are independent is consistent with the existence of an ‘efficient’ market for securities, that is, a market where, given the available information, actual prices at every point in time represent very good estimates of intrinsic values.”

Fama was seeing exactly the same thing as we saw with Betfair’s football odds whose implied probabilities matched almost perfectly the intrinsic probabilities as measured retrospectively from outcomes. The implication of this conclusion is stark: if price movements are independent, technical analysis, a methodology for forecasting the direction of prices through the study of past market data, is as worthless and unscientific as alchemy. He also seemed to recognise the paradox of skill that Stephen Jay Gould subsequently discovered in MLB hitters, remarking that the existence of

many sophisticated analysts actively competing with each other to take advantage of any dependencies in a series of price changes ultimately defeats its own purpose. Absolute skills are nullified, average relative skill is zero and what is left behind is just luck and independent price changes engaged in a random Brownian dance.

The version of market efficiency that Fama was describing is known as the weak form. The basic tenet, as Fama described, is that future prices cannot be predicted by analysing prices from the past. This is as much applicable to stock valuations as it is to betting odds. Excess returns cannot be earned in the long run by using investment/betting strategies based on historical prices or other historical data, without taking on more risk than has been assumed in the benchmark market. The efficient market hypothesis pertains to returns on a risk-adjusted basis. Gamblers can choose to invest in riskier stocks or back longer odds to achieve above average returns of course, but the downside is the possibility of doing much worse instead. Whilst technical analysis techniques will not be able to consistently produce excess returns, it is believed that some forms of fundamental analysis which look at the nuts and bolts behind intrinsic value (for example, assets, earnings and liabilities for an individual stock, interest rates, employment and GDP for foreign exchange markets, team and player form, injuries and weather in sports) may still provide excess returns. Fundamental analysis is essentially a bottom-up process that seeks to identify the 'true' value of an asset or probability of outcome from first principles, known as the firm foundation theory. Value betting is an example of fundamental analysis, although some of the methods used to calculate intrinsic probabilities rely on analysis of past trends of data (for example, the goal supremacy rating system I described in my first book *Fixed Odds Sports Betting: Statistical Forecasting and Risk Management*) which arguably could be considered a form of technical analysis.

The weak form of the efficient market hypothesis recognises that prices do not need to be at or near equilibrium, but only that market participants are not able to systematically profit from market inefficiencies. Efficient does not necessarily mean correct, just not consistently and predictably incorrect. Malkiel defends the hypothesis against the challenge presented by financial market bubbles and crashes on the ground that almost no one sees them coming. On the contrary, such phenomena are arguably little more

than expressions of hindsight bias. Furthermore, as Nate Silver's insight reminds us, "*a market that makes perfect predictions is a logical impossibility*," since there would be no market. Well, perhaps not for perfectly rational *homo economicus*. But the efficient market hypothesis doesn't insist that all agents should be rational. Wisdom of the crowd shows us that a market can be correct even if not one market participant is. All that is required is that individual errors are either randomly distributed such that the average error over the long term is zero or, if distributed systematically, with sufficiently small bias that the costs of playing in the market preclude their exploitation. Furthermore, everyone can still overconfidently and irrationally believe in the possibility of beating the system even if hardly anybody is actually managing to do that. Why else would we still choose to play in casinos? Games like roulette and craps, by definition, are the quintessential models of market efficiency, and yet millions still flock to play them. Of course, as we now know, rational utility to a gambler means much more than simply positive expectation measured by money.

There are two stronger versions of the efficient market hypothesis. The first, semi-strong version, postulates that prices (or odds) adjust so rapidly to new information and in an unbiased manner that trading on fundamental information will not be of any benefit either in producing excess returns. Not only will technical analysis prove to be useless, but so, too, will fundamental analysis. Everything that might relevantly influence the value of a stock or the probability of a result will already be reflected in the number on the board. As Malkiel notes, all of this does not imply that price movements are insensitive to changes in fundamental information. On the contrary, price movement is so sensitive that no one has time to benefit. In an age of nanosecond high-frequency trading, what possible chance does the average investor have? Furthermore, since the arrival of news, as Bachelier observed over a century before, is entirely random and unpredictable, price movements will be too. The strong version goes even further, suggesting that prices instantaneously reflect not just public news but private, insider information as well. Evidently, the illegality of trading on insider information renders the strong version more of a hypothetical extreme than something to be taken literally. If one couldn't take advantage of even insider knowledge, there would be little point in legislating against it. Conceivably, casinos regard the business of card counting in blackjack or

wheel clocking in roulette as forms of inside knowledge, refusing permission to exploit it. Similarly, bookmakers take measures to limit or refuse the custom of ‘sharps’ who arguably they regard as gaining unfair advantage trading on private information. Indeed, early explanations of the favourite–longshot bias in horse racing assumed a supply side price shortening to counter the effects of insider information.

The efficient market hypothesis has fallen out of favour in recent years following the research of Kahneman and Tversky into systematic cognitive biases, in particular overconfidence, loss aversion and herding, which we reviewed earlier in the book. Such biases, as we noted, give rise to irrational and suboptimal investment behaviour like the dumb-money effect and over-trading (particularly in men). In fixed odds markets, for example, we witness the over-betting of low probability outcomes (the favourite–longshot bias) and teams on winning streaks (the hot hand fallacy), whilst for point spreads we prefer backing the favourites. Proponents of the new behavioural finance point to recent economic bubbles and crashes in the stock market as examples of self-sustaining irrational exuberance, driven by investors thinking too much about what other investors are thinking about, and which as John Maynard Keynes says “*stay irrational longer than you can stay solvent.*” Sceptics of the efficient market hypothesis offer three ways of disproving it. The first seeks to find evidence of predictability; the second attempts to demonstrate that some players consistently managed to beat the market; finally there are those who argue that the hypothesis isn’t really a hypothesis at all because it cannot be properly tested. Let’s look at these in turn.

Price/Earnings (or P/E) ratios are commonly put forward as predictors of future stock returns, at least over fairly long time horizons. Specifically, the smaller the P/E ratio the bigger the predicted returns, the assumption being that companies with low P/E ratios are undervalued according to a fundamental analysis of intrinsic company data. Nobel Prize-winning economist Robert Shiller¹⁰⁷ found that as much as 40% of the variance in the 10-year annualised returns from the S&P 500 index during the 20th century could be explained by the average 10-year trailing P/E ratios. The correlation was even stronger for 20-year returns. Stocks with low P/E ratios of about 10 have typically produced an annualised return of about

9%. By contrast, stocks with P/E ratios of 25 have, on average, struggled to show annual growth of just 2%. Malkiel, however, suggests that such inefficiency might largely disappear once we include the influence of interest rates. Interest rates and inflation tend to be negatively correlated with P/E ratios. When inflation and interest rates are low, there is a greater opportunity for higher real earnings growth, increasing the amount people will pay for a company's earnings. The more people are willing to pay, the higher the P/E. Conversely, when inflation and interest rates are high, investors demand more for their money to maintain their purchasing power, and consequently P/E ratios are lower. During the final two decades of the 20th century, for example, US interest rates fell from historic highs of 15% to just 4%. At the same time, average P/E ratios rose from around 7 to peak at over 40 during the Dotcom bubble. During this extended period of economic prosperity (with only a brief hiatus during the 1990-91 recession), investors became increasingly happy to pay more for company earnings (although one might reasonably wonder whether this is merely evidence of herding). Consequently, the strength of the 'P/E effect' is not consistent over time. Furthermore, it is not entirely clear whether excess returns predicted by lower P/E ratios are entirely risk-adjusted. Finally, the undervaluing of a company might sometimes be entirely justified, reflecting not a failure of proper fundamental analysis but a real concern about its viability. If it does subsequently go bust, conceivably its low P/E earnings data disappear from future analysis leading to survivorship bias.

Warren Buffett is frequently paraded as an example of how the efficient market hypothesis must necessarily be flawed. How else could he have amassed an estimated worth of over \$70 billion, whilst his Berkshire Hathaway investment company has outpaced the S&P 500 benchmark index by a full 10% over the past half a century. Furthermore, he's done it through being a disciple of fundamental analysis and the firm foundation theory on which that is based. However, Warren Buffett began his investment career in the 1950s, just as the era of maximum structural advantage for investing in low-P/E markets began. Most of his investment career has been conducted in an environment of higher interest rates. Would he have been as successful had he been investing in the first half of the 20th century? Nassim Taleb (in *Fooled by Randomness*) puts it even more strongly than that. In a large population of random investors, there will inevitably be a

few who can produce an equivalent track record just by luck. With enough monkeys tapping away on keyboards, eventually one will reproduce something that looks exactly like skill.

Patrick Veitch is horse racing's equivalent of Warren Buffett. The *Racing Post* once described him as the bookmakers' 'public enemy number 1', which became the title of his 2009 book telling the story of how he took £10 million off their hands during an 8-year stretch beginning in 1999. Such performance is undeniably impressive, and according to Patrick even accounted for changing the way the UK bookmakers operated in an attempt to stop him winning. I first met him at Cambridge University in 1989 when he was already scalping 4-figure profits on a weekly basis, using runners to place bets and collect winnings in an attempt to hide the identity behind the betting activity. Throughout his betting career, he has relied on the services of as many as 200 agents to place his bets. I have little doubt that Patrick has proved to be a consistent winner. Nonetheless, it's worth taking a closer look at his numbers and methods to see how it could have been achieved.

When he wrote his book, Patrick said that his career profit over turnover was 16.7%. In Asian handicap or point spread markets this would represent a truly phenomenal, nay even unbelievable, performance. In horse racing, however, punters necessarily have to take bigger risks betting longer odds, on account of the greater number of runners compared to other sports betting markets. Consequently, volatility in returns will be much greater. This can mean large positive yields if you win but also large negative yields if you lose. Unfortunately, Patrick's book does not make available his full history of betting. He does, however, describe many high-priced horses he has backed all the way out to 100/1 and admits that most of his profitability has been gained at odds of 5/1 and longer. Based on selected wagers reported in his book, a rough estimate for a typical stake would be in the region of £5,000 to £10,000, although there has clearly been considerable variance with much smaller and much larger wagers than this. From these numbers we can estimate that Patrick placed in the region of 1,000 bets per flat racing season, the brand of racing that he preferred to specialise in. A particularly heavy day's betting saw 14, so this figure does not seem to be unreasonable. It's also in the same ballpark as many of the well known online racing tipsters. If we then assume average odds of 10/1 (again typical of many racing tipsters), it is possible to estimate that his performance

could be expected to happen by chance somewhere in the region of once in a few hundred thousand to a few million. Adjusting for risk (as the t-test does), this is comparable to a yield of about 5% from the same number of even money bets (in Asian handicap football or NBA spreads, for example). According to the 2010 British Gambling Prevalence Survey, 16% of the adult UK population, or around 8 million, bet on horse races. Based on this back-of-the-envelope analysis, a character like Patrick can conceivably be expected to show up simply because of good fortune. Of course, as Taleb regards Buffett, I'm not saying Patrick isn't skilled; rather that such a performance as his is not beyond the realms of lucky possibility.

My personal view of Patrick is that he's more than just lucky. From an early age, he was a prodigious talent, gaining admission to study mathematics at Cambridge University at just 16. He has also demonstrated a quite remarkable discipline and methodology in his wagering that would put most mutual fund managers to shame. Furthermore, the fact that bookmakers have queued up to refuse his custom is testimony to their acknowledgment of his sharpness as a value hunter. If bookmakers regard you as lucky, they'll more than likely just let you keep playing until you start regressing to the mean. That's something that between 1999 and 2006 Patrick didn't do, showing 6 to 7-figure annual profitability year after year. So how did he manage to stay so consistently successful for so long? Arguably, much of his advantage has been achieved through information not available to the general public. Being the owner or co-owner of over 130 winners during this 8-year period, Patrick has been able to exert considerable influence over the trainers who train his horses and the jockeys who ride them. On a number of occasions, he describes the business of talking tactics with the jockey prior to the race to take maximum advantage of the draw or the conditions of the track. Financial investment markets and their regulatory authorities would not permit such insider dealing. In betting, provided a jockey is not being instructed to lose a race, there are no such restrictions; the whole point of horse racing, after all, is for the benefit of gamblers. Instead, however, bookmakers take their own measures to eliminate those whom they perceive as gaining unfair advantage although, given the way some bookmakers ruthlessly target winning, it's hard not to come to the conclusion that they regard any kind of customer profit, including unskilled, as unacceptable. According to Patrick, both Ladbrokes

and William Hill reprogrammed their online trading systems to restrict the stakes for every internet customer betting on horses they believed were associated with him.

Of course, enhancing the probability of success through tactical discussion with jockeys and trainers necessarily means you have something worthwhile to tell them. Patrick regards horse racing as a ‘multi-layered conundrum’, and one that has taken him years and years of painstaking research to solve. For him, there is no short cut to betting success, no magic bullet. The only way to win is through sheer hard work, and few people will have the necessary diligence to apply themselves in that respect. Of course, if the paradox of skill is correct, it wouldn’t be of much benefit if everyone did. Betting, remember, is a relative skills contest. Patrick keeps videos of every race run over the previous 5 years. Indeed, the way he uses those videos, studying how the horses are ridden even down to the slightest detail of watching the precise movements of the jockey’s hands and reins, is strikingly similar to the learning process, via pattern recognition, of a chess grandmaster studying past positions or a professional baseball player watching the throwing motion of the pitcher. Like Buffett, he is steadfast in his conviction that the acquisition of such deep knowledge, and not luck, has provided him with the tools to succeed where almost everyone else fails. And yet Patrick also understands the principles of the efficient market hypothesis.

“I could provide you with methods that currently identify the right horses. Once I’d published such a document, the horses concerned would no longer be underestimated by the market.”

Essentially, Patrick’s existence as a ‘super-smart’ punter trading on years of deep knowledge and insider information is the exception that proves the rule, at least for the less strong versions of the efficient market hypothesis. Once Patrick’s money shows up, the markets start to move, sometimes in a very big way, most memorably for the 2:45 at Nottingham on 16 August 2004, when his horse Exponential was backed from 100/1 all the way down to a starting price of 8/1. It duly won, netting him a cool profit of £235,133.71 (Patrick’s record keeping is as meticulous as everything else he does).

According to Karl Popper’s principle of falsifiability, hypotheses cannot

be proven, only disproved. As Nassim Taleb reminds us, even with hundreds of thousands of white swan sightings and no black one, it is never possible to prove the statement ‘all swans are white.’ One single sighting of a black swan would immediately disprove the statement. As described, sceptics utilise the weapons of predictability and consistency to disprove the efficient market hypothesis. Others, however, attack it on the grounds that it can’t even be disproved, and therefore can’t be a meaningful hypothesis at all. To begin to test it, we first need to know what market efficiency looks like. For economists, this traditionally involves modelling expected returns with some equilibrium asset pricing model, against which abnormal returns can be measured. However, therein lies the problem. Our model is simply making an assumption of what market efficiency looks like. Consequently, when a model yields a return significantly different from the actual return one can never be certain if there exists an imperfection in the model or if the market is inefficient. This is known as the joint hypothesis problem, essentially because we are testing two hypotheses: 1) efficient markets look like this; 2) the market we are testing is efficient. If a market is found to be inefficient, it is impossible to know which hypothesis was proved false. Any anomalous market returns may reflect market inefficiency, a bad asset pricing model or both.

It’s tempting to conclude from this that economists enjoy splitting hairs. However, perhaps the importance of the hypothesis lies not in attempting to prove or disprove it, but rather as an idea (like the wisdom of the crowd) that reflects how and why prices move in real markets, and specifically why pricing errors (like Bayesian priors responding to new information) should be autocorrecting. Things are manifestly much simpler in betting where there are clearly defined probabilities (the odds), regular market closure (the results) and hence the opportunity for retrospective testing of efficiency (the odds versus the results). Consequently, it’s much easier to model explicitly what market efficiency should look like: the relative weighting of money risked on outcomes to realise desired returns is proportional to the probability of those outcomes. The scatter plot correlation of implied Betfair probabilities versus actual outcome frequencies I showed earlier was, in my opinion, as good a test of the efficient market hypothesis as we’re likely to get.

So are markets efficient? Theoretically, the hypothesis has significant

appeal. In practice, as Malkiel rightly points out, markets trading in uncertainty will not conform perfectly to the mathematician's ideal of complete independence of present price movements from those in the past. Furthermore, whilst markets may be continually and dynamically autocorrecting, the process will typically be neither instantaneous nor completely accurate. Sometimes, the market might overshoot whilst at other times it will be slow to respond. I have observed, for example, that odds for English football league matches (during 2010 to 2012) which experienced significant pre-match steaming (shortening) or drifting (lengthening) still fell short of achieving theoretical efficiency (as implied by results). That is to say, steamers and drifters didn't do enough steaming and drifting. The reason: we are not dealing with robots but emotional human beings who spend much of their time worrying more about what other people are thinking than fundamental value assessment. Consequently, the price of a stock or the odds of a team do not just reflect a rational assessment of expected profitability but also irrational market sentiment driven by competing psychologies of fear and greed.

However, just because markets might not behave purely like a random walk, or crowds might not always be wise, does not imply that they are exploitable. Most systematic biases (for example, the favourite–longshot and point spread biases) are typically not large enough, after the costs of playing are considered, to be reliably and consistently profitable. Where they might be (for example, the hot hand fallacy) their exploitation would require immense discipline. Furthermore, even when inefficiencies are found, it is more often than not in hindsight only, with little predictable before the event without the benefit of what amounts to insider information. Opportunities may exist, but as Terry Burnham says, our "*lizard brain is not built to be able to see them.*" Finally, where persistent deviations from classical stochastic behaviour can be found they are liable to self-destruct anyway after a given time. How long would a \$100 note remain unclaimed on the street pavement? As we saw earlier, it doesn't take very long for a crowd to become wise, where the vast majority of individual opinions becomes less accurate than the collective point of view. In a telephone interview with Glenn Croston (*The Real Story of Risk*), Aaron Klein, CEO of the investments advisory service *Riskalyze*¹⁰⁸, perhaps described best the

elementary absurdity of perpetual inefficiency:

“We’d all love to be able to see the future, to have a magic button we could push that would tell us which investments will go up or down. But then everybody would push it and it would not be accurate anymore. The magic button does not exist.”

Essentially, this is just a restatement of what Patrick had said about his forecasting methodology. In the face of this randomness, efficiency and unpredictability, we are left to consider two related questions: why is it so hard, from the point of view of prediction, to learn anything meaningful in such environments; and why do so many continue to believe (almost always wrongly) that their success is not simply a matter of luck? Fundamentally, the answers can be found in what Daniel Kahneman calls the ‘illusion of validity’.

The Illusion of Validity

Whilst serving in the Israeli Army’s officer training programme, Daniel Kahneman was tasked with measuring the leadership qualities of soldiers by means of an obstacle course challenge. Specifically, a group of 8 candidates was required to work together to haul a log over a six foot high wall without either the log or any of the soldiers ever touching it. Scoring the soldiers according to their performances, he developed a coherent story that he felt confident would be able to predict the future: taking over in moments of crisis was a predictor of leadership quality. Feedback sessions of how cadets were subsequently performing in officer-training school, however, revealed that the validity of the original predictions was not much better than blind guesswork. Yet despite this global evidence, he continued to feel and act as if his judgements were valid. Kahneman likened the dissonance this created to that imparted by the Müller-Lyer illusion (which I showed earlier in the book), where we **know** that two lines are the same length and yet continue to **see** them as unequal. He was so struck by the analogy that he labelled what he had found an illusion: the illusion of validity.

The illusion of validity is a clear instance of substitution of a cognitively more complex task with a more easily calculated heuristic (or short cut),

and the representativeness heuristic in particular, where, as Kahneman explains in his seminal (and personally favourite) publication, *On the Psychology of Prediction*¹⁰⁹, people predict or explain outcomes that appear most representative of the available evidence. Consequently, intuitive predictions are insensitive to the reliability of the evidence or to the prior probability of the outcome, known as the base rate. When ‘what you see is all there is’, our pattern-recognition engine (as for pigeons) is primed to jump to invalid conclusions about causal relationships. Just because something is more representative does not make it more likely. For Kahneman, all he saw was one hour of a soldier’s behaviour in an artificial setting. It was far easier for him to associate performance in that short time with future leadership skills than to begin to unravel his ignorance of factors that would ultimately determine a candidate’s performance.

The representativeness heuristic is nicely summarised by the familiar duck-test maxim: if it looks like a duck, swims like a duck, and quacks like a duck then it probably is a duck. The fallacy that arises is in assuming that similarity in one aspect leads to similarity in other aspects. This applies as much to objects as it does to predictions and outcomes. People in my home town of Buxton were guilty of fallacious similarity judgement when choosing to swim in a disused limestone quarry lake known as the ‘Blue Lagoon’. The lake actually gains its inviting turquoise colouration from the leaching of calcite crystals into the water rather than simply from the preferential scattering of blue and green light commonly seen in shallow tropical seas or swimming pools. Such complex information, however, is not available to the average person deciding to take a dip, who instead associates the colour with stored memories of similar instances. In fact, the ‘Blue Lagoon’ has alkalinity levels approaching that of bleach and is so toxic that the Government has refused permission for it to be drained. In an attempt to deter people from swimming in it, the local council decided to dye it black. The representativeness heuristic also encourages people to expect similarity between causes and effects (and in the medical profession, treatments as well). For example, it has commonly been believed that stomach ulcers are caused by stress, presumably because both stress and ulcers elicit similar sensations. In fact, bacteria are the culprits. Similarly, the prevalence of sunburn is frequently misattributed to temperature, where

how warm the day is might dictate how much sun cream is considered necessary. Of course, sunburn is caused by the sun's ultraviolet radiation, and is just as likely to happen on a cloudless day in summer irrespective of whether the air temperature is 20°C or 40°C. Furthermore, people tend to judge the probability of an outcome by finding a comparable known event and assuming that the probabilities will be similar, whilst neglecting base rate information that isn't as easily accessible. Kahneman and Tversky beautifully demonstrated the base rate fallacy by means of their taxicab problem.

Originally published in 1982, Kahneman has summarised his taxicab problem in *Thinking Fast and Slow*. He asks you to consider the following scenario and then give your intuitive answer to the accompanying question.

A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. 85% of the cabs in the city are Green and 15% are Blue. A witness identified the cab as Blue. The court tested the reliability of the witness under the circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colours 80% of the time and failed 20% of the time. What is the probability that the cab involved in the accident was Blue rather than Green knowing that this witness identified it as Blue?

This is a standard problem of Bayesian statistical inference in which Bayes' rule is used to update the probability for a hypothesis as evidence is acquired. We examined the technique early in the context of evaluating the evolving probability that a bettor might be skilled given his sequence of betting. Bayesian inference derives the posterior probability as a consequence of two pieces of information: a prior probability also known as the base rate and specific information derived from observation. In other words, the methodology calculates the conditional probability of A given that B is true.

In the taxicab problem, the base rate is 15% since that is the probability, assuming no other details are known, that a taxicab will be blue. However, the specific information of the case informs us that the witness can identify the cab's colour correctly 80% of the time. Given this information, what was your answer? The most common one that respondents give is 80% and most suggest probabilities over 50%. In fact, following Bayes' rule, the probability that the cab was blue is only 41%. To see why, it helps to consider frequencies, which most people find much easier to think about,

rather than probabilities. Suppose there are 100 taxicabs. 15 are blue, 85 are green. If the witness identifies the right colour 80% of the time, this will imply the following:

- 12 of the blue cabs will correctly be identified as blue (true positives)
- 3 of the blue cabs will be incorrectly identified as green (false negatives)
- 68 of the green cabs will correctly be identified as green (true negatives)
- 17 of the green cabs will be incorrectly identified as blue (false positives)

Consequently, the witness will identify 29 taxicabs as blue. In fact, only 12 of them are, so the probability that a taxicab is blue, given that the witness said it was, is 12 divided by 29, or 41%. The reason the actual answer is so much lower than the intuitive one is because it's easy to forget about all the false positives – the identification of a green cab as blue. Given that there are so many more green cabs than blue, these false positives contribute significantly to the total identified as blue.

The neglect of base rate information and the influence of false positives in statistical inference are worryingly prevalent in the medical profession. A typical Bayesian problem concerns the misdiagnosis of breast cancer and the use of unnecessary intervention. About 1% of women in their 40s develops breast cancer (the base rate). 80% of the time a mammography will correctly identify it. However, 10% of the time it will incorrectly return a positive result even if there is no cancer. If a woman has a positive test result, what is the probability that she has breast cancer? Typically, most people, including even physicians, give answers much higher than the true Bayesian posterior. As for the taxicab problem, the explanation is the same: base rate neglect. Let's turn it into a frequency problem again. Suppose we have 1,000 women; 10 will actually have breast cancer. Consequently:

- 8 of the women with breast cancer will correctly show a positive mammogram result (true positives)
- 2 of the women with breast cancer will incorrectly show a negative

mammogram result (false negatives)

- 891 of the women without breast cancer will correctly show a negative mammogram result (true negatives)
- 99 of the women without breast cancer will incorrectly show a positive mammogram result (false positives)

Thus, the real probability of a woman with a positive mammogram result having cancer is about 7.5% (8 divided by 107), less than 10 times the most typical answer given by respondents. Again, as Nate Silver reminds us, when the underlying incidence of something in a population is low, false positives are liable to dominate the results.

It is clear now why gamblers in markets of psychology (betting, investing and poker) may vastly overestimate their chances of succeeding. They are simply ignoring the base rate that value expectation in zero-sum games is zero less the commission paid to play, when presented with individual winning performances, either their own or someone else's. People are reluctant to infer the specific from the general, particularly when the latter involves unintelligible statistical concepts (performance regresses to the mean) and the former presents a far more interesting narrative (I predicted X and X happened). Bettors and investors who have beaten the market express a misplaced confidence in the probability that they have really done so. As Kahneman explains, however, such confidence arises not through reasoned evaluation that the probability is correct, but through the coherence of a story and the cognitive ease of processing it. For 'winners', the possibility of profitability that exists in games of unknown unknowns is substituted by an inevitability of profitability. What is the probability that I am a winner? This is a difficult question. It's much easier to answer another one that can take its place: have I won? The corollary of this illusion of validity is the self-serving bias that I've previously discussed, where people attribute their successes to internal factors like skill, but attribute their failures to external ones like luck. This is the illusion of skill. In turn, this fosters an illusion of control, the tendency for people to overestimate their ability to cause things to happen. 'I made money from my predictions \Rightarrow I am skilled at prediction \Rightarrow I caused my wealth to increase' is a far more coherent, appealing and psychologically rewarding narrative than 'I was

merely chucking darts.' Finally, survivorship bias, where only the 'winners' are left on display, will enhance the sense of overconfidence.

The representativeness heuristic also influences our perceptions of randomness. Fundamentally, our 'interpreter' is evolutionarily designed to seek patterns, sometimes even when none actually exists. Things that appear to lack any logical sequence are regarded as more representative of randomness, and therefore more likely to occur, than more orderly sequences. For example, a sequence of 5 consecutive blacks on a roulette wheel would be considered less probable than a sequence of 2 reds, 2 blacks and a red, whereas in fact both sequences have exactly the same chance of occurrence. The coherent narrative, however, is that the former sequence is more interesting as a pattern, and in the absence of any further understanding about probability theory, 'more interesting' substitutes for 'less likely'. 'More interesting' is also more memorable, making such patterns much easier to recall from our associative memory. Remember, the tendency for the human mind towards such selective reporting is described by the availability heuristic. For example, the story of 26 consecutive blacks at the Monte Carlo casino is far more memorable than the countless millions of other sequences where that didn't happen.

A closely related tendency is for people to underestimate the influence of randomness in accounting for sequences. The Monte Carlo example is a case in point. It's easy to assume that such a rare event must have some other causal explanation, for example an unbalanced wheel or some other underhand influence. However, if one properly considers the number of times a roulette wheel is spun anywhere in the world, frankly we should expect far longer sequences to have occurred in the past 100 years since that happened. Indeed a quick internet search uncovers an example in an American casino in 1943 of 32 consecutive reds. A simple back of the envelope calculation guessing the number of worldwide casino roulette wheels (10,000) spinning once per minute for 10 hours per day and 300 days per year would suggest that a sequence of 33 has probably been witnessed somewhere in the world in the past 10 years. Part of the problem in underestimating the influence of randomness is that we do not properly take into account the number of opportunities for something to occur, as this example illustrates. Consequently, we are often surprised when random chance produces coincidences. Consider another example. How many

people do you think we would need in a group before there was an even chance (50%) of two of them sharing the same birthday? With 365 days in a year, most people will opt for a fairly large number. In fact, the actual figure is just 23. It's not intuitively obvious that with such a small group there are still 253 possible pair combinations. By the time you get to just 105 people, there is more chance of you winning the UK national lottery than there being no shared birthdays. Such underestimation of randomness probably accounts for Top-Tipster.com refusing to accept that his best tipsters were probably just guessing like all the rest. If you have enough combinations to choose from, almost anything is possible just by chance.

We are also particularly insensitive to sample size, or what Kahneman and Tversky called a 'belief in the law of small numbers'^{[110](#)}. Judgements made from small samples are inappropriately perceived to be representative of the wider population. Tversky and Kahneman^{[111](#)} presented subjects with the following scenario, accompanied by a question:

A certain town is served by two hospitals. In the larger hospital about 45 babies are born each day, and in the smaller hospital about 15 babies are born each day. As you know, about 50% of all babies are boys. However, the exact percentage varies from day to day. Sometimes it may be higher than 50% sometimes lower. For a period of 1 year, each hospital recorded the days on which more than 60% of the babies born were boys. Which hospital do you think recorded more such days?

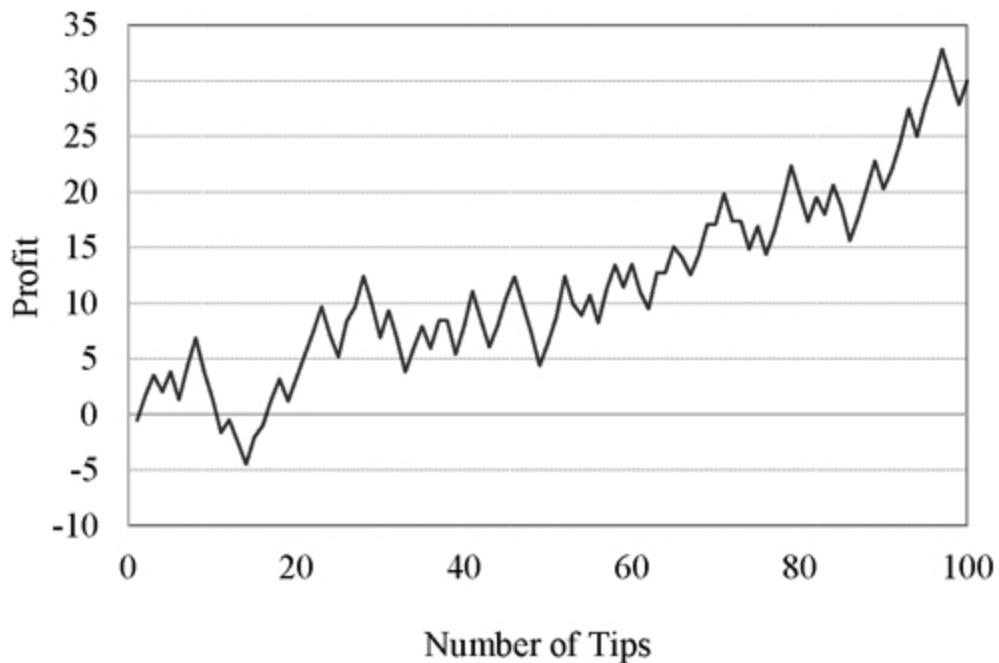
- *The larger hospital*
- *The smaller hospital*
- *About the same (within 5% of each other)*

According to binomial theory, the number of days where boys born outnumber girls by at least 6 to 4 will be nearly three times greater in the smaller hospital compared to the larger one, simply on account of the larger volatility in birth ratios. A larger sample is less likely to stray very far from 50%. Yet only 22% of respondents gave the correct answer. Evidently, thinking about the implications of sample size is not particularly intuitive.

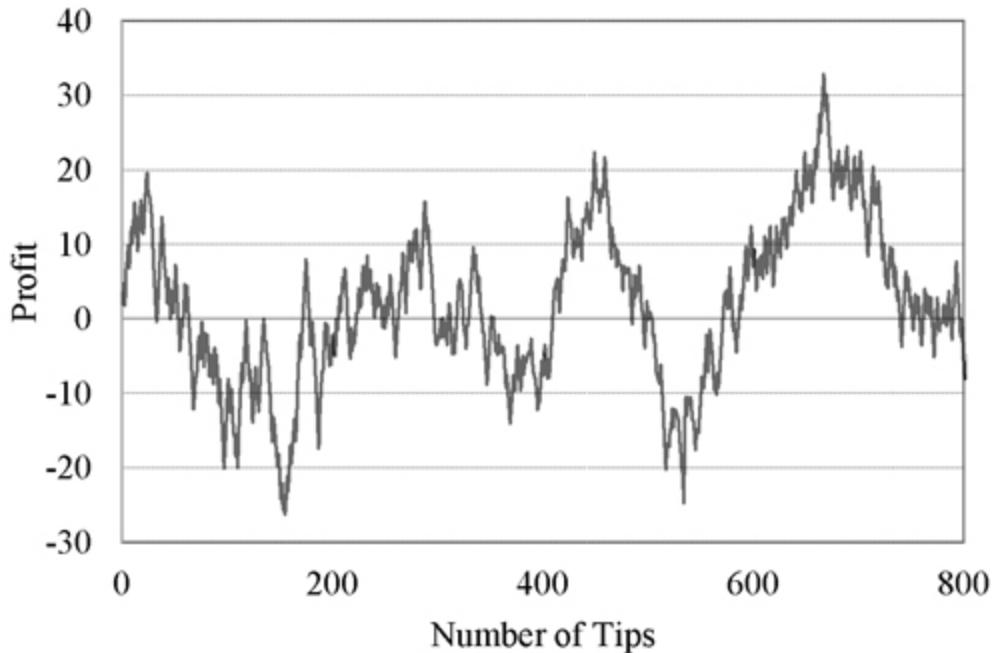
Two phenomena I've previously described that arise out of a belief in the law of small numbers are the gambler's fallacy and the reverse gambler's fallacy, sometimes described as the hot hand fallacy. According to the gambler's fallacy, sequences of one outcome, for example, roulette black, must eventually be followed by another outcome, in this case red, in order

for the full sequence to be considered representative of a random process, which most roulette players understand the game to be. This fallacy may also be described by a belief in the law of averages, a misinterpretation of the law of large numbers, in which it is assumed that unnatural short term balancing or things evening out should occur. Of course, things don't have to even out; they just tend to do so over the long run. In contrast, the hot hand fallacy is the erroneous belief that a person who has experienced success with a random event in the short term has a greater chance of further success in additional attempts. Evidently, people expressing the hot hand fallacy reject randomness due to a belief that a streak is no longer representative of a random sample. We are back again to the illusions of validity and skill.

Misplaced belief in the law of small numbers can be particularly damaging to gamblers who misinterpret profitability from small samples as representative of a departure from randomness, and in particular as evidence of forecasting skill over the longer term. Evidently, Burton Malkiel's chartist friend suffered an illusion of validity when urging him to buy his imaginary stock on the basis of a randomly generated time series that bore all the hallmarks of a causally significant pattern. Sports bettors, in particular, are prone to apophenia – the perception of patterns in random data. Consider the next two charts. The first represents a time series of 100 tips submitted for verification by a sports betting advisory service. Impressive, isn't it?



Now look at the next one. This time series is from the same tipster. The first represents just a small sample of his tips, the second the whole record. It shows nothing but a random dance in the bankroll, sometimes rising, sometimes falling, even over extended periods, which on their own might convince people that something predictable exists. As with almost all tipsters, nothing consistently predictable exists at all. Regression to the mean is absolute; (relative) skill is zero.

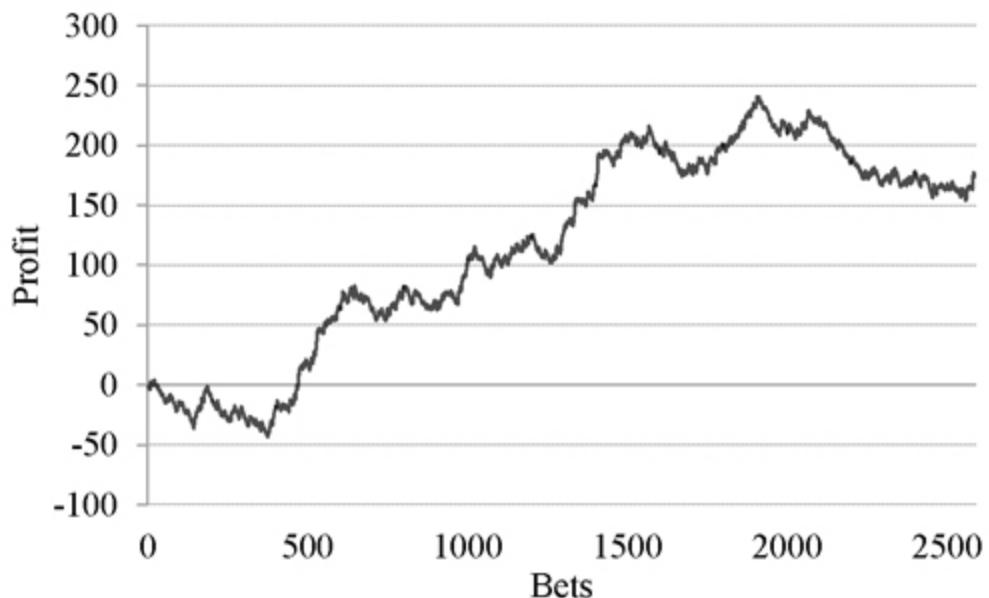


Regression to the mean, like the statistics of small samples and the law of large numbers, is not a concept that can be intuitively understood by a brain that prefers causal interpretations for patterns it sees. Daniel Kahneman, as usual, puts it succinctly:

“When our attention is called to an event, associative memory will look for its cause. Causal explanations will be evoked when regression to the mean is detected, but they will be wrong because the truth is that regression to the mean has an explanation but does not have a cause.”

Incorrect interpretations of regression to the mean are part of a more general confusion of correlation with causation. Even where the distinction is understood, laziness on the part of the statistician looking to find explanatory relationships in data to support hypotheses – for example, I am skilled at betting – can lead to invalid conclusions about what is really happening. Patterns in betting or investment returns can look very meaningful without actually being so. Correlation without causation is essentially worthless. Consider the following time series of profitability from betting on English league football during the 2012/13 season.

Evolution of profits during the 2012/13 English football league season



Most punters would probably agree that this betting time series looks pretty healthy. It has a yield of 6.7% and a p-value of 0.02 (significant at the 95% level) implying such a result could be expected by chance about 1 in every 50 seasons. Unfortunately, I have to tell you that such a series was generated by betting on every away win in the Premiership, Football League and Football Conference to best market prices. The proverbial monkey got rich simply by betting everything blindly. Of course, unless something systematic had gone terribly wrong in the way bookmakers had been pricing these away wins during that season, the strong correlation between accumulated profit and the number of wagers ($R^2 = 0.8$) that might imply skill, is completely spurious with no underlying causation at all. Betting on every away win in English league football does not ‘cause’ a bettor to become rich. If it did, it would represent the mother of all market inefficiencies and ultimately there would be no bookmakers left to facilitate it. Of course, given what we’ve learnt about markets in this chapter, such an inefficiency would quickly self-destruct, assuming it ever existed in the first place. On the contrary, this time series just represents a lucky season which hadn’t fully regressed to the mean by the end of it. Consider, however, that this sample was made of 2,588 wagers. In the context of sports betting, that’s a considerable sample size. Indeed, a typical football bettor would

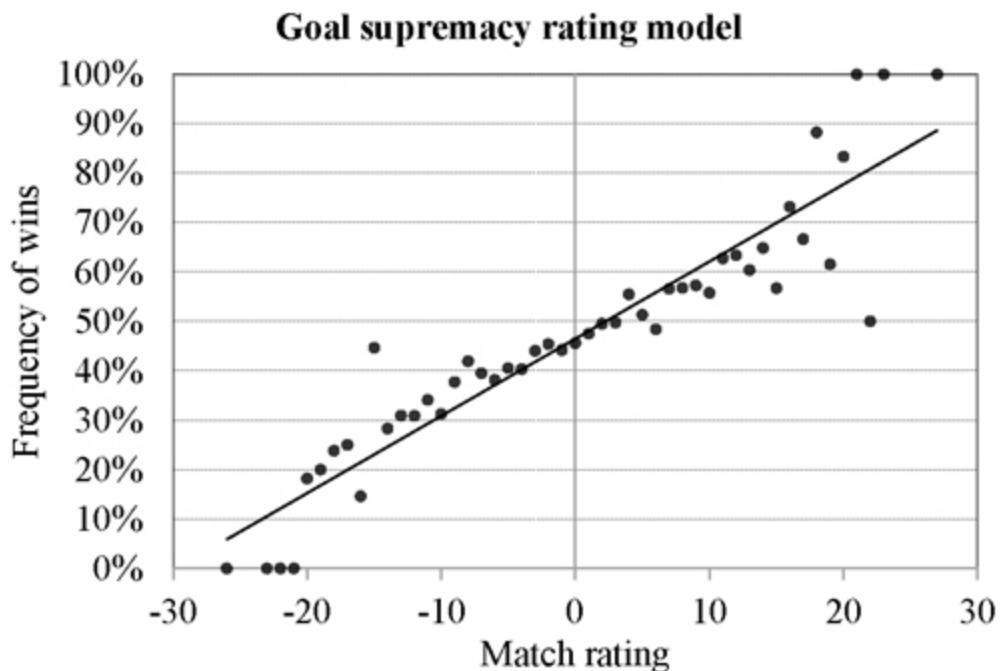
take about 3 to 5 years to turn over such a number. If such samples don't fully regress to the mean, arguably there is no sample size too large that will help encourage a fallacious belief in the law of small numbers.

I have little doubt that, if I looked long and hard enough, I could find statistically significant, yet meaningless, correlations in data sets 10 times this size.¹¹² We rely on 'big data' to provide us with validity at our peril. Google did precisely that when they published a paper in the world's most respected scientific journal, *Nature*, on the detection of influenza epidemics using search engine query data¹¹³. Without access to any data for medical check-ups, Google predicted the spread of influenza across the US because of a strong correlation between flu-related search queries and physician visits in which a patient presents with influenza-like symptoms. It was also much faster at achieving this than the Center for Disease Control and Prevention, which relied on reports from medical practices. Today, it even has its own global 'Google Flu Trends' website¹¹⁴. The success of 'big data' like this encouraged some of its cheerleaders to reject the need for statistical theory altogether. When one has so much data to analyse, why bother worrying about what causes what, because the numbers speak for themselves. Google Flu Trends was theory-free and much cheaper as a result.

In the winter of 2012/13, however, the prediction accuracy dropped, with Google overestimating the spread of flu-like illnesses by almost a factor of two. The problem, of course, was that because a correlation had been found between what people searched for online and whether they had flu symptoms, the assumption was that getting flu caused people to use the internet to search for information about it. In fact, Google hadn't bothered to develop a hypothesis to account for why flu-related search terms might be correlated with the spread of the disease in the first place. Google didn't know why the correlation existed and what caused what; it had just found a statistical pattern in the data. Tim Harford¹¹⁵, the *Undercover Economist* and columnist for the *Financial Times*, explains that "*if you have no idea what is behind a correlation, you have no idea what might cause that correlation to break down.*" One explanation that might account for an overestimation of outbreaks would be pre-emptive searching about avoidance strategies and treatments by the 'worried well' in response to

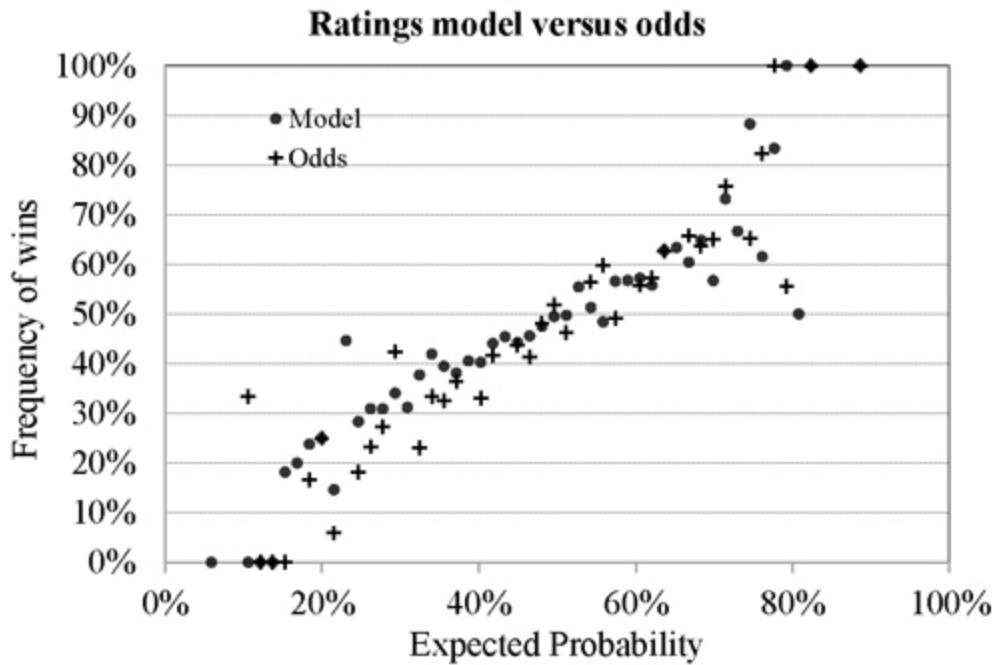
excessive news coverage of influenza. Another, as Tim postulates, could be Google changing the goalposts when it began automatically offering diagnoses in response to search queries. To reiterate, correlation without causation is meaningless, or as David Spiegelhalter, Winton Professor of the Public Understanding of Risk at Cambridge University, says: “*complete bollocks; absolute nonsense.*”

Let’s get back to football. Most serious football bettors, of course, don’t do silly things like bet blindly on all outcomes; they develop forecasting models that attempt to predict the outcome of football matches. Here’s a match rating system I developed and published in my first book *Fixed Odds Sports Betting* in 2002. It is based on a goal supremacy model which measures the relative strengths of two sides based on their goal difference for the previous 6 games for English league football. The distribution of win frequency versus match rating is shown in the chart below, based on data for the 1993/94 to 2000/01 seasons for four professional English divisions. The match rating is simply calculated by subtracting the 6-game goal difference for the away side from that of the home side. Positive/negative ratings imply a larger/smaller than average expectation for a home win.

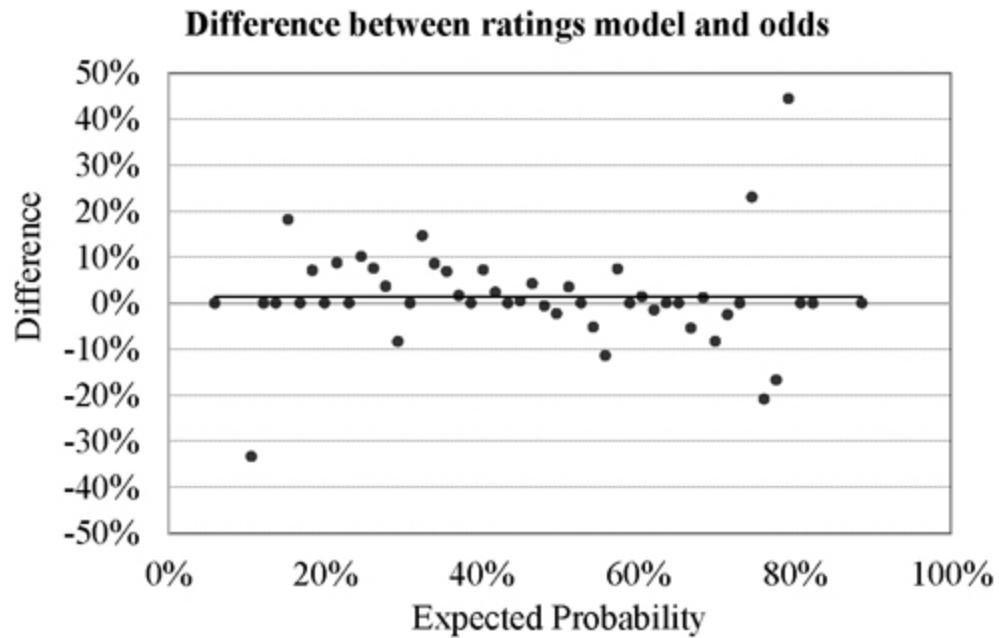


Despite some variance for extreme positive and negative ratings (due to the far fewer number of matches with those values), the relationship between recent goal supremacy and home wins seems pretty solid. Indeed the model, as depicted by the trend line, accounts for as much as 86% of the variability of outcomes. The regression equation for this model allows us to compute fair betting odds, which can then be compared to published ones to identify value opportunities. Applying the model to the following 2001/02 season, a profit over turnover of 2.1% was achieved (to level stakes) from 526 wagers (compared to a loss of 3.7% blindly betting on all home wins).

Manifestly, this model does a great job of identifying teams that are more likely to win. Unfortunately, however, this does not guarantee profitability. Naturally, many others who are betting on football are developing similar forecast models. The nut that needs to be cracked here is not whether this model works, but whether it works better than those used by the bookmakers, or rather those used by other bettors which shape the betting odds the bookmakers will publish. Remember, the betting odds (at least with bookmakers who don't interfere with their markets) to a significant degree represent the collective wisdom of everyone's opinion expressed with money. If we are to find value in the published odds, our model must produce more accurate forecasts than most of the others. The bigger the bookmaker's take, the greater our model superiority will need to be to overcome this disadvantage. From what we now know about wisdom of crowds, this will be no easy task. So how did my goal supremacy model perform relative to the market? Take a look for yourself. The next chart compares the distribution of win percentage versus expected probability as implied firstly by the model (the circles) and secondly by betting odds for the 2001/02 season (the crosses). Evidently, there's not a whole lot of difference. My model might be good at forecasting wins, but it's doing pretty much the same thing as the forecasts implied by the betting odds. It might be accurate, but it's not more accurate than other forecasting methods used by other bettors.



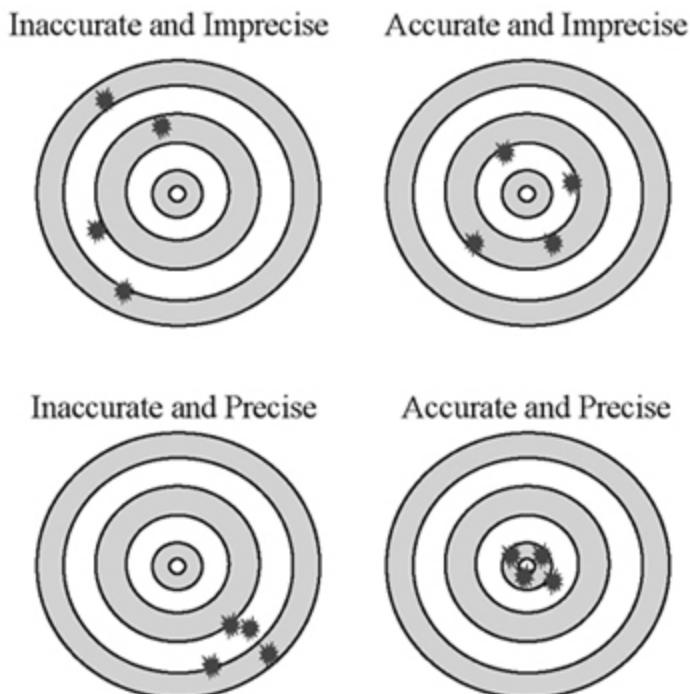
The final chart shows the difference between the two. If you look carefully, you'll see the trend line; $R^2 = 0$.



It's now clear that the home win profitability generated in 2001/02 by my recent goals supremacy rating was not built on any valid correlation. My model predictions did not 'cause' those profits because it was not more

accurate than other models doing the same thing. Lest you need reminding again, betting is not an absolute skills contest but a relative one. My model was fairly accurate in an absolute sense, but relatively speaking it appeared to be worthless. The profit it generated in 2001/02 was simply achieved with a bit of luck. If relative skill or predictive accuracy is zero, there will be no persistence in outcomes and regression to the mean will be inevitable. By way of confirmation, yields for the 2002/03, 2003/04 and 2004/05 seasons that followed were -9.3%, -2.5% and -12.6% respectively, compared to -8.6%, -5.9% and -9.0% from blind betting on all home wins. I had simply suffered an illusion of validity.

A recent goals supremacy football rating system is hardly the most novel of forecasting ideas. Nevertheless, those engaged in developing more 'intelligent' models should heed the dangers of confusing correlation with causality. Another way to think about the difference is by means of the distinction between precision and accuracy. Precision means that measurements are close to each other, that is to say, they are reliable (and reproducible). Precision or reliability, however, does not guarantee accuracy. Accuracy is a measure of how close you are to the 'true' value. A useful analogy is with throwing darts at a dartboard, as the next diagram illustrates.



Precision is associated with random error, accuracy with systematic ones (otherwise known as bias). My dart throwing will not be very precise but, if I throw enough of them, averaged together they will be centred on the bullseye, and consequently quite accurate because of minimal systematic error. We might also describe such a performance as one of low skill and high luck, in much the same way that markets arrive at efficient solutions through collective wisdom and the ‘assistance’ of Adam Smith’s invisible hand. A measurement is considered valid if it is both accurate and precise; a valid measurement implies some causal explanation for the outcome being measured. In the winter of 2012/13 Google’s Flu Trend data, which thus far had been pretty precise, missed the target. Its accuracy had been compromised, probably on account of the fact that it hadn’t concerned itself with causality. Validity is a measure of whether what we think is the cause is actually the true cause, and whether our measurement repeatedly points to that conclusion. Validity therefore implies both persistence and predictability. Sadly, bettors and investors often misinterpret precision, accuracy and validity when studying their outcomes, confusing correlation and causation in the process. The standard fallacy is in believing that excess returns above the market expectancy were ‘caused’ by a player’s predictive skill. Everything I have talked about thus far has hopefully demonstrated that for the vast majority bettors and investors causation is just an illusion, formed by our lazy ‘belief engine’ that prefers to see patterns rather than undertake complex statistical thinking. Betting and investing is largely taking place in a zero-validity environment.

The Environment of Skill

What causes the illusion of validity in gambling markets? Evaluating the quality of a football team or the balance sheet of a company involves a great deal of work and commitment that surely responds to a learning process and the acquisition of skill. After all, the more chess or tennis I play, the better I’ll get at chess and tennis. Surely, the more research I do into sports or financial markets, the better at forecasting I will become. The problem here is that we’re not answering the right question; we’ve substituted it with an easier one. Skill in stock picking or sports betting is

not simply a matter of picking winners; rather, it's about whether we're better at picking them than everyone else. For fear of labouring the point, speculative gambling, where opinions are traded (betting, investing and poker), is a zero-sum, relative-skills competition. Many players exhibit considerable skill in forecasting the future; the trouble is that you're not necessarily rewarded for doing so. When buyers and sellers, backers and layers are competing with each other, the skill that really matters is evaluating whether the available information about the market is already incorporated into the price, and being able to do that consistently. If markets are mostly efficient most of the time, the prospects for outperforming the wisdom of the crowd appear to be severely restricted. Not only do most players lack such a skill, but as Daniel Kahneman says "*they appear to be ignorant of their ignorance.*" System 1, the automatic, fast and intuitive brain has substituted the much easier question without System 2 even realising, creating a coherent narrative about prediction skill – "I made my profits happen" – that underpins a subjective and self-serving confidence. To understand the difference between merely evaluating the future and outperforming the market, we need to remind ourselves of the mechanisms behind the acquisition of skill, and specifically what lies behind intuition.

2014 PDC World Darts Champion Michael van Gerwen was asked by *Telegraph* columnist Jim White how he throws his darts^{[116](#)}. "*I don't know... it's just natural.*" I imagine that many professional sportsmen and women, as well as chess players, musicians, artists, air force pilots, surgeons and others engaged in activities that evidently require skill, would say a similar thing. After many years of repetitive learning doing the same thing over and over again, the correct decision or behaviour just feels intuitive. When psychologist Gary Klein^{[117](#)} probed a fire lieutenant about how he came to make a life-saving decision, evacuating his team from a building that's ablaze, he put it down to extrasensory perception, a feeling that something wasn't right. Moments after they left the building, it collapsed. From such descriptions, Klein developed a theory of decision making called the recognition-primed decision model. The model can be applied to all those activities listed above. Fundamentally it is a model of intuitive decision making built on pattern recognition. Earlier in the book, I referred to two types of learning process: a faster emotional one with neural circuits

quickly built by the dopamine-serotonin anticipation-reward feedback mechanism; and a slower one that lies behind the acquisition of expertise via pattern recognition. Intuition, then, is not extrasensory or magical, but merely the recall of stored memories in the form of patterns.

When embarking on the learning of a new skill, there is a lot to learn. Instead of memorising every unit of information, it is grouped into larger chunks to aid storage and recall. ‘Chunk’ is really just another word for pattern, and ‘chunking’ is the name of the theory that describes the learning process. For example, people usually store and recall their phone number in two blocks of numbers rather than each digit individually. Chunks are typically meaningful to the person engaged in the learning process. Chunking lies behind the ability to recall large sequences of numbers, for example in competitions to recite the number pi to hundreds, thousands or even tens of thousands of decimal places. Chunking is the mechanism behind learning to read, grouping letters into syllables and words, words into clauses and, in some rare examples of speed reading, whole pages of text. Similarly, chess grandmasters treat game positions as chunks or patterns, having acquired a mental database of thousands of different configurations. Ask a grandmaster to memorise and recall a particular game configuration and they’ll manage it in a few seconds. However, set the pieces in random positions and she’ll fail as miserably as any other rank amateur, because there is no pattern to recall from the memory database.

Such failure of pattern recognition in sport was wonderfully demonstrated when Jennie Finch, a gold medallist with the US softball team at the 2004 Olympics, struck out future Hall of Fame slugger Albert Pujols, two-time All Star Brian Giles, and future Hall of Fame catcher Mike Piazza at the 2004 Pepsi All-Star Softball Game. Their failure had nothing to do with reaction times. Indeed, it couldn’t. Pitching at about two-thirds of the speed of a MLB pitcher but from about two-thirds the distance, the ball was arriving at the bat in roughly the same time. Standard MLB pitches take about 400 milliseconds to arrive at the strike zone. That means it’s already halfway there by the time a batter has even initiated a muscular action to respond to it. Fastest reaction times are generally in the region of 200 milliseconds, the time it takes for nerve impulses to jump numerous synapses from retina to brain and then from brain to muscle. Fundamentally, humans don’t have a visual system fast enough to track the

pitch. What the top MLB hitters do have, however, is high visual acuity, and a huge database of patterns based on past pitches. When they hit a ball, they are not reacting to its flight but rather matching the positional characteristics of the pitcher with one of the memorised patterns from previous pitches. Essentially, expert hitters are focusing on visual cues before the ball has even been pitched. As David Epstein says in *The Sports Gene*, “*the only way to hit a ball travelling at high speed is to be able to see into the future.*” It’s now clear why Pujols, Giles and Piazza failed so miserably. They had nothing in their database to call upon. Finch’s action, of course, is underarm, and thus presents a whole new set of visual data that professional MLB hitters will never have acquired. “*When a baseball player faces a softball pitcher, he is stripped of his crystal ball.*”

What sets experts apart from the rest of us then is an accumulated database of patterns gained over long periods of deliberate practice. When engaged in a task that involves decision making, often we know where to look for the best solution, it’s just that most of us lack the cognitive database needed to extract the information from it. As individuals practice a skill, the mental processes involved move from the slow, conscious and rational processing System 2 areas of the brain to the intuitive, subconscious and automatic System 1 areas. Think of when you first learnt to drive a car. Everything you did – acceleration, breaking, gear shifting, signalling, checking mirrors – all had to be thought about in advance. The more you did it, however, the more intuitive the processes became, until eventually you found yourself doing all those things without even being aware you were doing them. The accumulated time of driving had built up a large database of memories, associating particular patterns to particular outcomes, which assist with future decision making. Following Epstein the ‘software’ matters much more than the ‘hardware’. Without the database of stored patterns, everyone will be a chess grandmaster facing a random chess board.

This model of the acquisition of expertise or skill is what is called Naturalistic Decision Making. It is significantly at odds with the work of Kahneman and Tversky on psychological heuristics and biases. The former focuses on the marvels of intuition, the latter more on its flaws. The likes of Gary Klein are more preoccupied with how experts become experts; disciples of prospect theory, by contrast, shift attention to where expertise is

illusory. This is understandable given the types of ‘experts’ each has had in mind. Klein has spent much of his time studying fire fighters, nurses and air force pilots engaged in behaviour with obvious feedback mechanisms for learning. Kahneman, on the other hand, has concerned himself more with stock pickers, clinicians and political pundits engaged in forecasting under uncertainty. For him, decision making under uncertainty may not even yield itself to expertise. The debate gave rise to a long and fruitful collaboration between Klein and Kahneman that culminated in a joint publication¹¹⁸ to settle an outstanding question: when can you trust an experienced professional who claims to have intuition?

Klein and Kahneman settled on the following answer. The conditions of acquiring skill are two-fold. 1) An environment should be sufficiently regular and stable to be predictable, where cause and effect are linearly related – the condition of validity. 2) There should exist opportunities to learn these regularities through prolonged practice and feedback. For people learning to drive, play the piano or become a grandmaster in chess, the rules of the game remain the same (stable) and outcomes are clearly linked to the things that those people do (linear). Consequently repeated practice, supported by immediate and unambiguous feedback, cultivates the acquisition of skill. Arguably, betting and investment markets lack both these conditions, although the ‘deep study’ methods applied by Warren Buffett and Patrick Veitch might suggest not entirely. They are complex and mostly random because the news that drives the movement of prices arrives to the market randomly. We’ve already seen that both stock pickers and sports bettors fail basic tests of consistency and validity. This implies that the environments within which they operate are neither regular nor predictable, but instead dominated by luck. The corollary is that there is limited scope for feedback. Since feedback is the oil that drives the machinery of deliberate practice, accumulating playing experience in markets will not deliver expertise. Imagine trying to practice a game of roulette, using previous wheel spin outcomes as feedback. It’s as absurd a proposition as trying to lose such a game of luck. Similarly, in betting and investing the price handicapping process that is implicit in the balancing of supply and demand, backing and laying, necessarily means that establishing causal relationships between decisions and outcomes is more akin to

guesswork. This is not to say I cannot become an expert in sports or business. For example, I know nothing about the sport of handball but if minded to study it, gaining familiarity with teams, players, their histories and televised performances, I imagine after a time I would be able to intuitively know as much about the sport as I do about football. In a competitive market environment, however, this learning process does not offer me feedback in determining whether I am better than the market at forecasting an outcome. My absolute skills may be significant but relatively speaking if they are no better than the rest of the market, and I have no way of knowing whether they are, I might as well be guessing. Success in markets is not measured by whether you agree with them, but by whether you can beat them, and beat them consistently and predictably. In a betting market, there are no good or bad teams, just correctly or incorrectly priced ones.

The problem here is that collectively wise markets, and success in them, are inherently probabilistic in nature, not deterministic like many of the activities which markets speculate about. What feedback is available from backing a team or player and winning my bet when I don't even know what the 'true' probability was that they should win, and what events in the game ultimately caused its 'probability wave function' to collapse in my favour? As Michael Mauboussin says:

"[W]hen your undertaking involves a dose of luck, the link between cause and effect is broken. In the short term, even when you do everything right, the outcome of your effort can be bad. Moreover, you can succeed even when you do everything wrong."

Of course, it's worse than that. In the short term, we don't even know whether our effort is right or wrong. In such environments, there is little to be gained by focusing on outcomes. Rather, and as I'll examine more thoroughly in the final chapter, it pays instead to focus more on the process. Furthermore, even if we did know, the strong version of the efficient market hypothesis would argue that such fundamental information would be worthless anyway, because the market would already be reflecting it.

At least in betting, however, the probabilities are explicitly stated in the price. Furthermore, given the availability for market closure there is at least some opportunity for feedback, via watching the matches we have bet on. Of course, given the many path histories such games can follow, trying to

identify the reasons why things happened the way they did, and whether they correlate consistently with the things you predicted would happen and for the same reasons, will be no simple task. In poker, too, there may well be sufficient statistical regularities that can be exploited by accumulated practice and the development of skill. Not only do players gain regular and repeated closure with the revealing of competing hands (offering feedback for improving playing strategy) but the size of the market one plays in is typically much smaller than for betting. The smaller a market is, the greater the scope for inefficiency to be found within it, since smaller markets will tend to be less wise. Hence the better the chances for establishing linearity and validity with regards intuition.

Given the vast size of financial markets and their lack of closure, investing would appear to represent the hardest game of all. Perhaps, then, it is surprising it still attracts the superior cultural status it does when compared to other forms of speculative gambling, since, of all the forms, it appears the closest thing to a game of craps. I suppose that is due, in part, to the purpose for which it was originally designed as a moderator of risk, a role that continues to have positive social and economic benefits despite the recent turmoil in the global economy. Perhaps also because of the illusion (through the engine of capitalism) that it's easier to win, if winning simply means making a profit (although arguably, if trading costs are properly accounted for, even that is debatable). However, financial markets are now so large, so sophisticated, so wise, and consequently most of the time so efficient that it's the hardest gambling domain, outside games of pure chance, to exert any meaningful expertise. This is not because investors lack skill. On the contrary, more of them now know more about businesses and reading their data than ever before. It's because of the paradox of skill in which the variation in skills between players has narrowed to the point where relative differences are negligible and drowned in a sea of randomness. The more investors have become skilled in evaluating the business prospects of companies, the less opportunity they have for exploiting the inefficiencies that markets have to offer, in much the same way that the 0.400 MLB hitters have disappeared. To a lesser extent, the same is probably true in sports betting. With increasing popularity, more and more liquidity is available, trading on more and increasingly sophisticated opinions. Liquidity increases efficiency, and efficiency

decreases the validity of forecasting strategies (the links between cause and effect) and the probability that anything other than luck will account for above average, risk-adjusted returns.

Klein and Kahneman also agreed on one other important principle: the confidence that people have in their intuitions is not a reliable guide to their validity. As Kahneman says, “*do not trust anyone – including yourself – to tell you how much you should trust their judgment.*” Kahneman traces people’s confidence in their own judgements to the ease with which they form a coherent narrative, that is to say, how an explanation for an outcome seems to make sense. Coherence, however, does not guarantee that the explanation for the outcome – its causality – is true. This is most relevant for decision making under uncertainty, and specifically forecasts about the future. Nate Silver has suggested that the confidence people express in their predictions may be inversely correlated with their validity. Nowhere is expression of confidence as groundless as in the speculative gambling of betting and investing.

Yet our mind has evolved to see only that which is there to see, too easily ignoring other things which it does not know. Confidence in intuition is the default position. As discussed previously, bold forecasts (overconfidence) help to counterbalance timid choices (loss aversion). For our ancestors, being predisposed to find patterns rather than discounting them was certainly essential for survival. As Nigel Turner¹¹⁹ of the Centre for Addiction and Mental Health in Ontario says:

“[I]f a person was walking in the jungle and saw a pattern of light and dark stripes in the shadows, it would be prudent to assume that the pattern was a tiger and act accordingly. The consequences of incorrectly assuming that the pattern is not a tiger far outweigh those of incorrectly assuming that it is. But when applied to random events, this survival ‘skill’ leads to errors.”

The ease with which causal narratives are formed from meaningless patterns in random data is matched in intensity only by the difficulty of changing those beliefs. Try to do so at your peril. “*What do you think I am; a nobody?*” People are so bad with probability and randomness that luck is rarely even considered as an explanation for winning by those heavily invested in trying to beat the system. Presentation of evidence, of monkeys throwing darts, is met with self-serving denials to guard against the

cognitive dissonance or anxiety which contradictory explanations for a person's gambling success produces. Denial usually feels more rewarding than change. Ginsberg may argue that we can't win, break-even or even get out of the game, yet the rewards for playing, as I've argued before, are evidently more than purely financial. The ego would much prefer to remain in control through believing that it can shape its own future, to win with wits, than to be forced to delegate to chance. Control requires determinism, with things happening for a reason. Chance and randomness have no place in such a world. Not only are most of us fooled by randomness, we probably also deny that we are, and if forced to accept it deny that it matters anyway. In a sense, it doesn't matter. For much the same reason why many express Divine faith, others choose to gamble to feel in control. If randomness means that this is illusory, so what? A feeling is often just as rewarding as the reality.

⁹⁷ Ericsson, K. A., Krampe, R. & Tesch-Romer, C., 1993. The Role of Deliberate Practice in the Acquisition of Expert Performance. *Psychological Review*, **100**(3), pp.363-406.

⁹⁸ <http://www.pinnaclesports.com/en/betting-articles/betting-strategy/how-good-is-a-betting-tipster>

⁹⁹ We can check to see how Betadvisor's tipster ranking pages looked in the past via the Wayback machine (<http://web.archive.org>).

¹⁰⁰ Rohleder, M., Scholz, H. & Wilkens, M. Survivorship Bias and Mutual Fund Performance: Relevance, Significance, and Methodical Differences. *Review of Finance* (2011), **15**, pp.441-474.

¹⁰¹ LaPlante, D. A., Kleschinsky, J., LaBrie, R. A., Nelson, S. E., & Shaffer, H. J., 2009. Sitting at the Virtual Poker Table: February 2005 through February 2007. *Computers in Human Behavior*, **25**(3), pp.711-717.

¹⁰² Smith, V. L., 1962. An Experimental Study of Competitive Market Behaviour. *The Journal of Political Economy*, **70**(2), pp.111-137.

¹⁰³ Levitt, S.D., 2004. Why are Gambling Markets Organised So Differently from Financial Markets? *The Economic Journal*, **114**(495), pp.223-246. Levitt is also the co-author, with Stephen Dubner, of the critically acclaimed book *Freakonomics*.

¹⁰⁴ Simmons, J.P. & Nelson, L.D., 2006. Intuitive Confidence: Choosing Between Intuitive and Nonintuitive Alternatives. *Journal of Experimental Psychology: General*, **135**(3), pp.409-428.

¹⁰⁵ Simmons, J.P., Nelson, L.D., Galak, J. & Frederick, S., 2011. Intuitive Biases in Choice versus Estimation: Implications for the Wisdom of Crowds. *Journal of Consumer Research*, **38**(1), pp.1-15.

¹⁰⁶ Fama, E. F., 1965. The Behavior of Stock-Market Prices. *The Journal of Business*, **38**(1), pp.34-

105.

[107](#) Shiller, R., 2000. *Irrational Exuberance*. Princeton: Princeton University Press.

[108](#) <https://www.riskalyze.com>

[109](#) Kahneman, D. & Tversky, A., 1973. On the Psychology of Prediction. *Psychological Review*, **80(4)**, pp.237-251.

[110](#) Tversky, A. & Kahneman, D., 1971. Belief in the law of small numbers. *Psychological Bulletin* **76 (2)**: pp.105-110.

[111](#) Tversky, A. & Kahneman, D., 1974. Judgment under uncertainty: Heuristics and biases. *Science*, **185(4157)**, pp.1124-1131.

[112](#) It is entirely possible that the statistically significant ‘hot hand’ bias I found evidence for in my data set of 36,126 football league matches (discussed earlier in the book) represents such an example.

[113](#) Ginsberg, J., Mohebbi1, M. H., Patel, R. S., Brammer, L., Smolinski, M. S. & Brilliant, L., 2009/ Detecting influenza epidemics using search engine query data. *Nature*, **457**, pp.1012-1014.

[114](#) <https://www.google.org/flutrends/>

[115](#) <http://www.ft.com/cms/s/2/21a6e7d8-b479-11e3-a09a-00144feabdc0.xhtml>

[116](#) <http://www.telegraph.co.uk/sport/othersports/darts/10514734/How-to-throw-darts-like-a-pro.xhtml>.

[117](#) Klein, G., 1998. *Sources of Power: How People Make Decisions*. Cambridge MA: The MIT Press.

[118](#) Kahneman, D. & Klein, G., 2009. Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, **64(6)**, pp.515-526.

[119](#) Turner, N. & Powel, J. *Probability, Random Events, and the Mathematics of Gambling*. <https://www.problemgambling.ca/EN/ResourcesForProfessionals/Pages/ProbabilityRandomEventsandtheMathematicsofGambling.aspx>

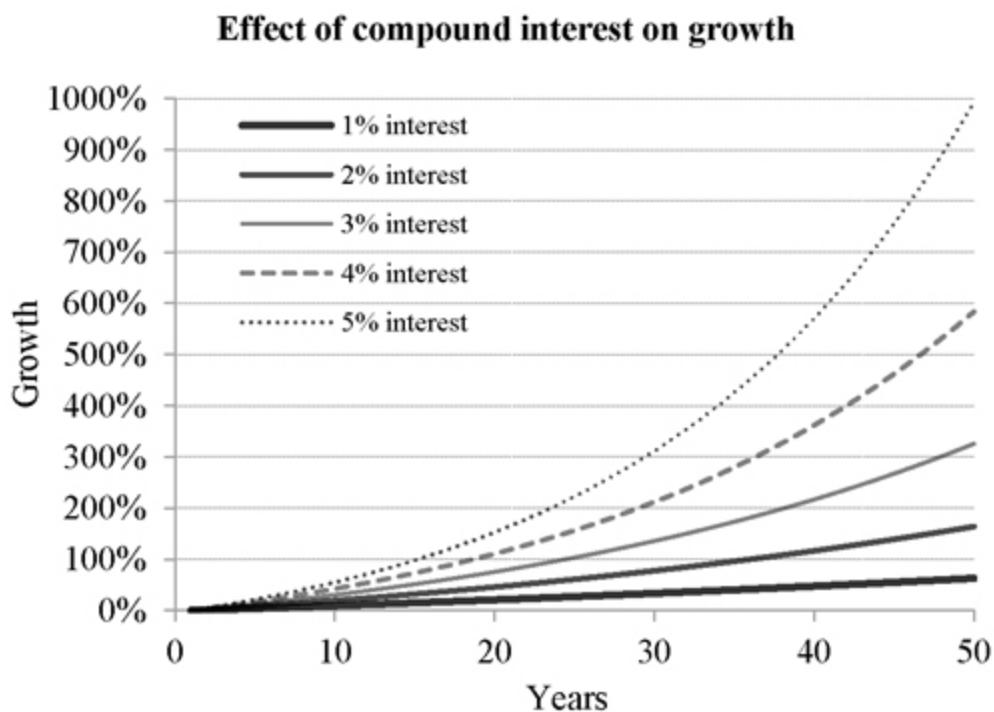
WINNER TAKES ALL

Despite the wisdom of crowds, random walk theory and the efficient market hypothesis, we still have to admit, as Burton Malkiel does, that markets probably don't conform perfectly to these ideal models. After all, if they did, prices wouldn't move or at least would have to move instantaneously. The fact that prices do move, and take time to do so, implies there is no such thing as perfect market equilibrium. Rather it is dynamic, always looking to eliminate inefficiencies but never perfectly doing so as a constant stream of news is received and interpreted by the players. Whether it is possible to consistently predict those inefficiencies sufficiently to overcome the cost of playing in the market is a moot point. Proponents of behavioural finance and bettor/investor irrationality would argue in favour of its possibility. Efficient market advocates, however, would remind them that irrational players don't necessarily mean irrational markets. Nevertheless, given the existence of 'super-smarts' like Warren Buffett, Patrick Veitch and others, including George Soros (the man who broke the Bank of England), Zeljko Ranogajec (blackjack and racing expert and the world's biggest gambler for 2010), Haralabos Voulgaris (NBA king and sportsbook owner) and Matthew Benham (owner of Smartodds, one of the world's biggest betting syndicates), survivorship bias aside, we should at least consider the possibility that some professionals are capable of making a living through beating the market. Despite the complexity and non-linearity of markets, it may be that some individuals are genuinely talented or hard working enough to uncover sufficient regularity to overcome the costs of playing and be consistently successful. If that is accepted, what are the consequences for the rest of us? Sadly, it makes our prospects even worse than if it was all just a game of chance. In zero-sum games like betting, investing and poker, someone has to pay for the 'winners'. The market facilitators won't do it; inevitably, that means the job falls to the rest of us, the 'losers'. Furthermore because small differences in relative skill, over time, are compounded, the resulting distribution of profitability can be

distinctly non-linear. That is to say, the winner takes all.

Compounded Advantage

Of compound interest, Albert Einstein is believed to have remarked: “*it is the most powerful force in the universe.*” The chart below demonstrates why he was probably right. Compound interest is essentially interest earned on interest. Consequently, the greater the interest rate, the faster the growth rate of your wealth. A small difference in interest rates, when compounded over a long period of time, results in a dramatic difference in final wealth.

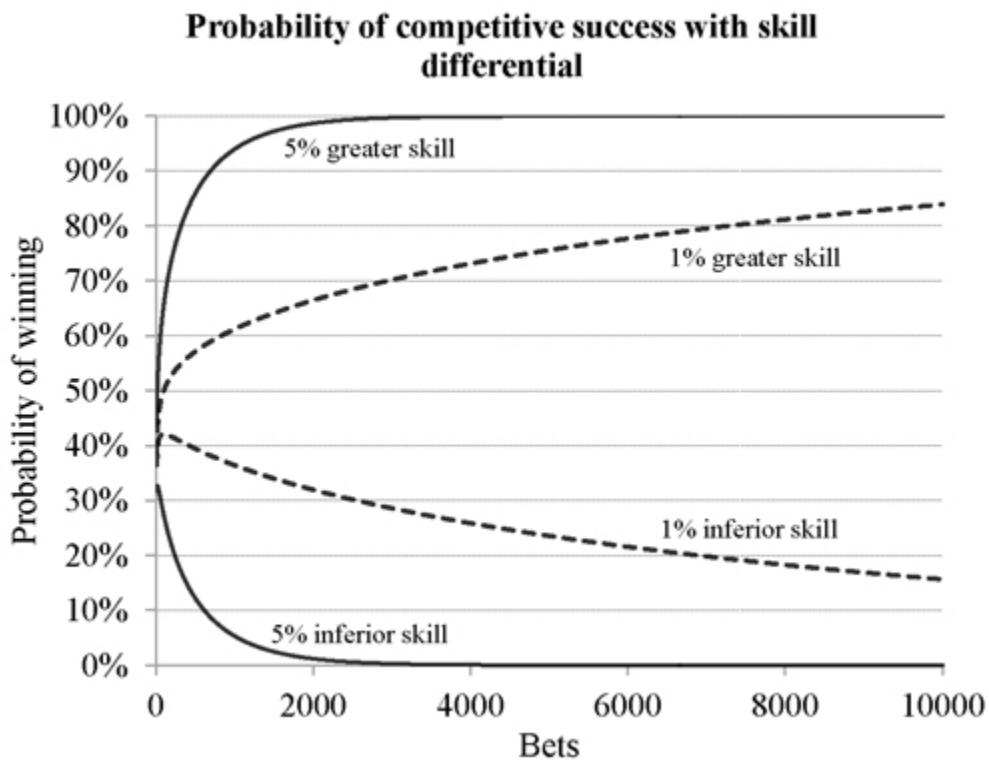


A similar compounding effect exists in iterative relative skills competitions. By that I mean a competition within which there are many mini contests and where one player has a slight advantage over the other. I’ve previously referred to such competitions in 2-player sports. Tennis is the obvious example, played over perhaps 150 to 250 individual points. Accounting for serve, if one player is slightly better than the other on a per-point basis, over the long haul that will translate into a much higher probability of overall victory.

Possibly the most famous exposition of compounded advantage is Darwin's theory of natural selection. For him, survival could be considered as a numbers game, with the odds stacked against most creatures competing for finite resources in a zero-sum competition, where rewards are based on relative rather than absolute performance. Genetic traits are tested for suitability towards their environment, with beneficial adaptations gradually selected – the winners – and maladaptive traits gradually discarded – the losers. Darwin implicitly recognised that small advantages ultimately translate into large differences in success: "*The slightest advantage in one being... over those with which it comes into competition... will turn the balance.*" The statistician Ronald Fisher, who subsequently synthesised many of Darwin's original ideas on evolution, showed that a random genetic mutation giving an advantage of just 1% in fitness to the organism would spread via inheritance through the entire population within 100 generations.

Betting, investing and poker can be seen as taking place in a similar competitive environment. Relative to the market, these gambling games are zero-sum with no wealth being created, but players competing for differential advantage through monetary transfers that reflect the rewards and punishments of success and failure respectively. For poker it is intuitively obvious that players compete against each other; in betting and investing less so, because typically the other players are hidden from consciousness. Squares tend to think of betting as a competition against the bookmaker, when really it is a competition of opinions amongst players. As I've already explored at length, the odds largely represent the public face of all those private opinions^{[120](#)}. When you back a price you are effectively competing against everyone else who laid it, either explicitly if betting via an exchange or implicitly (through backing the other outcome) if betting with a bookmaker. Investing, too, is simply a competition between buyers and sellers. The buyers (and backers) think the price is too cheap, the sellers (and layers) believe it's too expensive. The players compete to see who is right. Of course, being subject to so much luck, a single contest doesn't tell us very much, but over many hundreds or thousands of contests, as for a tennis match or a poker game, if there is any differential in forecasting skill between players, this can be expected to show its influence. To see how,

consider the following chart, which shows the evolving probability of success for two competing players engaged in successive even money wagers, where one of them is slightly more skilled at prediction than the other.



The dotted lines show two players separated by a 1% differential in skill. For even money propositions this is equivalent to 50.5% and 49.5% success rates respectively for each wager. Despite the tiny relative skill differential, the players' respective chances of overall success (as measured by greater profitability) are seen to diverge, with the probability of overall success for the weaker player steadily decreasing. The solid lines show the evolution of the probability of success for a 5% relative skill differential, equivalent to 52.5% and 47.5% forecasting success on a wager by wager basis. This time, the divergence is far more rapid. Indeed, after just 1,000 iterations, the weaker player will have just a 6% probability of overall victory, and by about 11,000 iterations he's got more chance of winning the national lottery. When the difference is as much as 10%, representative of the sort of strike rates the best sports handicappers appear to be capable of, the weaker player might as well just pack up and go home. Even after just 500 bets, he has

little more than 1% probability of out-scoring his opponent, and past just 1,000 bets the game is as good as over.

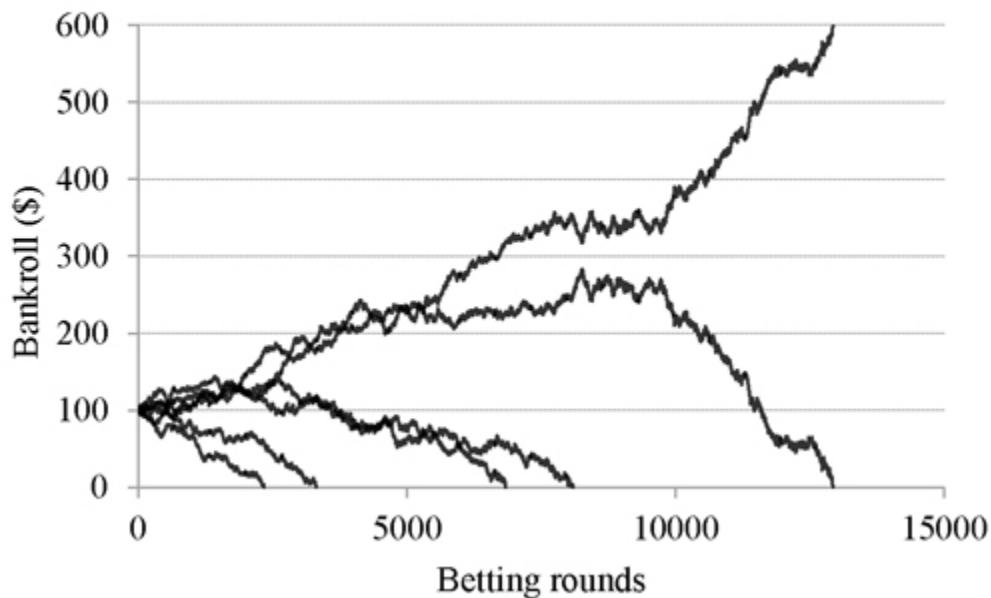
Readers may have noticed that this chart is reminiscent of an earlier one I showed depicting hypothesised probabilities of being in profit when betting against a bookmaker's margin. This is not a coincidence. Essentially, they are showing the same thing. The margin can effectively be regarded as the relative skill differential the bookmaker has over its customers; the bigger the margin, the greater the relative skill differential. Of course, no skill is involved here; rather it arises simply as a consequence of the shortening of prices relative to fair value to ensure that the bookmaker earns his commission for facilitating the action between players. The same is true in investment markets, where the relative advantage the market maker has will take the form of the buy-sell spread and any other transaction costs that are incurred by traders and investors. Similarly, in poker, the rake taken by the poker room can be regarded as a type of super-skilled player. Evidently, then, the presence of sharps in a market magnifies the problem for squares. Not only do they have to overcome the disadvantage imposed on them for playing but also the challenge of beating competitors who may be better than they are.

What happens when there are many players competing in a market? Sadly, my modelling skills are not up to such a complex task, but I've made an attempt at examining the question by considering the interaction of six of them, sufficient, in my opinion, to provide a meaningful representation. As in the example above, in each round one player competes against another chosen at random. Each round consists of a fair even money wager (there is no bookmaker) with stakes of \$1. The winner pockets the \$1, received generously from the loser. All six players begin the competition with \$100. Between each player is a relative skill differential of 4%. Thus, if two players selected to compete are closest to each other in terms of skill, the weaker one has an implied forecast probability of 48% whilst the stronger one has one of 52%. Consequently, the relative skill differential between the weakest (player 1) and strongest (player 6) is 20%, with implied forecast probabilities of 40% and 60% respectively. Wagers are settled by means of a random number generator with the relevant skill differentials applied for each contest. After each round of betting, the total money amongst the six players remains constant at \$600; this is a zero-sum game. There is just one

rule to the game: once any player loses his bankroll, he is eliminated from the contest. The chart below illustrates how the game played out.

As it turned out, the players dropped out of the contest in ascending order of relative skill level, although for a considerable period of time, players 3 and 4, and players 5 and 6 respectively, were fairly evenly matched. Of course, this was just how this particular contest played out. In other path histories, the finishing order could well be different, although the most probable order of defeat was the one witnessed here.

Modelling the effects of compounded advantage in a zero-sum iterative even-money 6-player contest



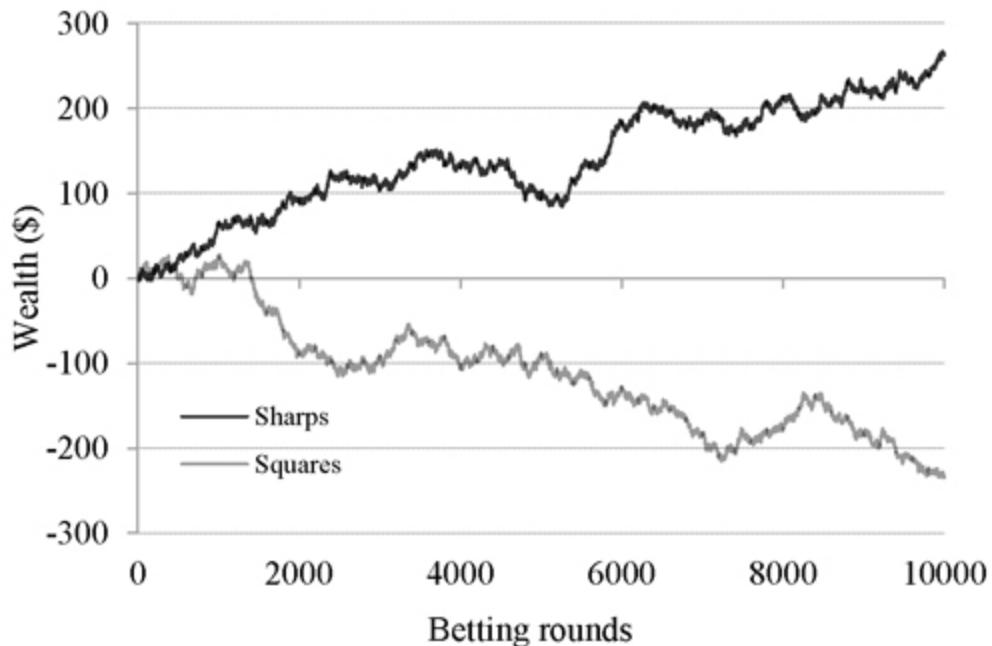
There is an obvious conclusion here: players with the greatest relative forecasting skill will feed off those weaker than they are until they are the last ones standing. It is a case of ‘winner takes all’. In his book *The Signal and the Noise* Nate Silver has done something similar modelling the behaviour of poker players. As here, the elimination of the weakest squares from a table has a cascading effect on other players. The previously next-to-weakest player is now the biggest square, and will start to lose money until he too goes bust. In my model, you can see how players 3 and 4 initially performed quite well in the early rounds of betting whilst weaker players were still playing in the game. Once these were eliminated, however, the inevitable took its course. Ultimately it’s simply a matter of time until the

strongest player has beaten all the rest.

A more complex model based on up to 3,000 players engaged in a coin-tossing game was developed by Michael Miller¹²¹ at the Santa Fe Institute to specifically address the question of the evolution of wealth distribution in strategy-based betting. Two versions of the model were tested. In the first instance, a biased coin with unchanging bias was used to mimic the influence of skill. Secondly, the bias was allowed to vary in response to the betting patterns of players, probably a better representation of reality given that true outcome probabilities in betting are never known exactly. In both cases, play evolved towards one stable equilibrium wealth distribution, where one player with the best strategy (closest to the true coin bias) was left with all the wealth or multiple players if they shared the same strategy. Of course, in real markets, there will be a constant stream of new players taking the place of those who have either gone bust or more usually lost sufficiently to dissuade them from continuing further. Nevertheless, the same ‘survival of the fittest’ environment will persist, with sharps surviving longer and squares disappearing earlier.

Roughly what proportion of sharps to squares could co-exist in a zero-sum prediction market? A little hypothetical thought experimenting might help us here. Since losers are paying for winners, one possible scenario might involve an exact 50:50 ratio with the gains made by sharps balanced exactly by the losses incurred by squares. The chart below illustrates the evolution of wealth for a typical higher-skilled and lower-skilled player in such an environment for a fair even money game (wagers \$1) where the relative skill differential between the two types of player is 5% (implying 52.5% versus 47.5% forecast success respectively). Where two sharp players or two square players meet, the wager becomes a game of chance. Where a sharp meets a square, the former has a 5% advantage. Gain expectancy for the sharp is +2.5% whilst loss expectancy for the square is -2.5%.

Wealth evolution in a betting market with 50:50 squares and sharps

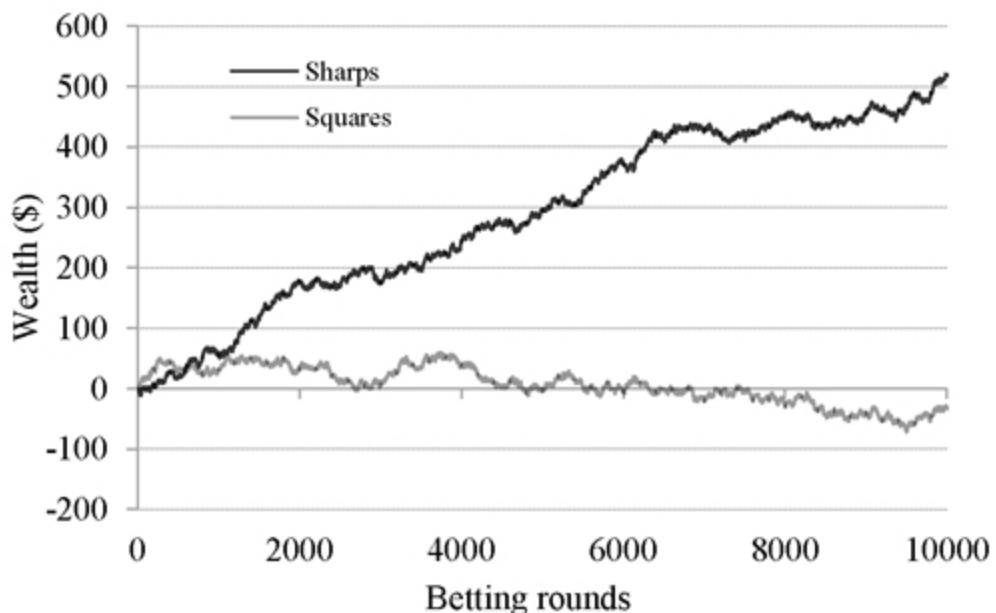


Of course, such a scenario seems wholly unrealistic given the low validity environment of markets and the compounding of any relative advantage, ensuring its non-linear correlation with profitability. Presumably, having such a large proportion of sharps in a market would be completely unsustainable in a real market setting anyway, unless squares were infinitely wealthy and charitable, or there was a limitless supply of new ones waiting to replace those already gone bust. On that last point, some might rightly observe that the large numbers of people who choose to bet with negative expectancy imposed by the bookmaker and the longevity of the industry imply that such a proposition might not be quite as absurd as it first seems. With each passing year a new set of people becomes eligible to gamble, and provided the numbers are sufficient to replace those who quit, the pool of squares will be eternally recycled. However, supporting the lifestyles of a few bookmakers is a far cry to funding the wealth acquisition of a huge number of sharps. Undoubtedly, as sharps become wealthier, they will start to increase stakes. Unless the squares do likewise, something would have to give. In such an inefficient and irrational market the squares would quickly disappear, leaving the sharps to fight it out amongst each other. Small differences in skill levels between the remaining sharps would give rise to

big differences in profitability, with the more skilful ones feeding off those less so, just as the dynamics of ‘winner takes all’ imply. Alternatively, if they were equally skilled at prediction (an unlikely proposition) all that would be left would be a random market, with outcomes simply settled through luck, or no market at all since, perceiving no more advantage, they would all stop playing. Really, this is just a restatement of the paradox of skill. As absolute levels rise (in this case because the squares have been eliminated), the variance across players tends to diminish, pushing markets towards increased efficiency and a paucity of valid and consistent positive expectation.

An alternative and more plausible scenario that fits better with observed outcomes for players would suggest a much bigger ratio of squares to sharps. Not only does such a balance allow the squares to last longer in the game, since they spend more time ‘tossing coins’ with each other, it also helps sharps (if there are any) to avoid too much contact amongst themselves. Such a scenario is illustrated again below. This time the ratio of squares to sharps is 100:1. The profit expectancy for the former is now almost 5% whilst for the latter it is only marginally less than 0%.

Wealth evolution in a betting market with 100:1 squares and sharps



Perhaps more importantly, and as the ‘winner takes all’ thesis has attempted to demonstrate, even slight disadvantages in relative skill will be costly, meaning a hierarchy of performance is the most likely outcome. Such is the influence of compounded advantage that, even for a fair game without playing costs and where the variance in skill across players is small, the most stable equilibrium involves a large number of small losers supporting a much smaller proportion of big winners. Even if we supposed that relative skills were linearly distributed amongst players, their playing outcomes would not be. As philosopher George Bernard Shaw once remarked: “*in gambling the many must lose in order that the few may win.*”

The Pareto Principle

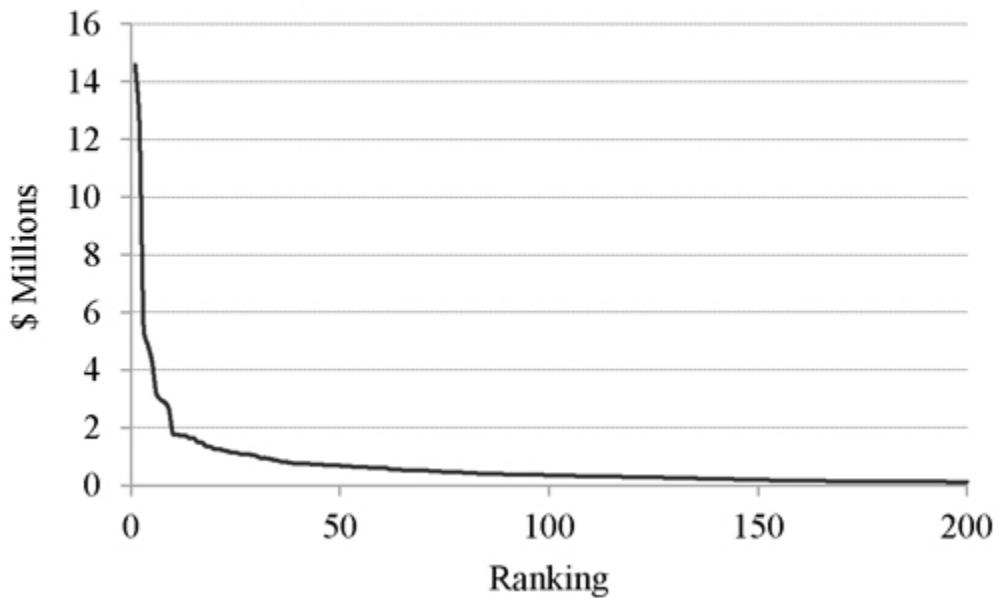
Situations where small differences in starting conditions lead to big differences in outcomes are frequently described by power laws. A power law is a functional relationship between two quantities, where one quantity varies as a power of another. Written words, cities and earthquakes are three examples of things that follow power laws. Linguist George Zipf, after whom Zipf’s law is named, demonstrated that in written texts the frequency of any word is inversely proportional to its rank in the frequency table. Typically, the most common word is ‘the’, followed by ‘of’ and then ‘and’ appearing just half and a third as often. I’ve checked this for my own text to date and it is pretty close. Zipf’s law also mysteriously holds true for the populations of cities. The Gutenberg-Richter law is a similar power law correlation between the number of earthquakes and their magnitude. Generally speaking, there are about 10 times as many earthquakes for a particular magnitude as there are for the next higher one. Earthquake magnitude is itself a logarithmic scale. With each increasing magnitude of earthquake, the earth displaces about 10 times more. Hence the logarithm of the size of the earth’s displacement is inversely proportional to the logarithm of the frequency of such an earthquake.

One well known power law relationship is called the Pareto Principle, sometimes known as the 80/20 rule, named after Italian economist Vilfredo Pareto who showed that, at the close of the 19th century approximately 80% of the land in Italy was owned by 20% of the population. Mathematically,

the 80/20 rule is roughly followed by a power law distribution called, unsurprisingly, the Pareto distribution. Many natural phenomena have been shown empirically to exhibit such a distribution; for example Facebook posts (80% of your likes come from 20% of your followers), business sales (80% come from 20% of clients) and complaints (80% are made by 20% of customers). Furthermore, this pattern is recursive; within the top 20% of a system that exhibits a Pareto distribution, the top 20% of that slice will also account for disproportionately more of whatever is being measured, and so on and so on. There is nothing special about the ratio 80/20, although it does appear to be a particular common pairing. Other Pareto-type distributions will show different ratios. For example, the top third ranked countries in the 2015 Eurovision Song Contest collected 80% of the votes, whilst 80% of the global wealth (as measured by national 2014 GDP figures from the International Monetary Fund), was monopolised by just 10% of the world's nations. Wealth, in particular, appears very prone to such power-law distribution. Such disparity is sometimes also referred to as the Matthew effect, after a verse in the Gospel according to Matthew: *"For unto every one that hath shall be given, and he shall have abundance: but from him that hath not shall be taken even that which he hath."*

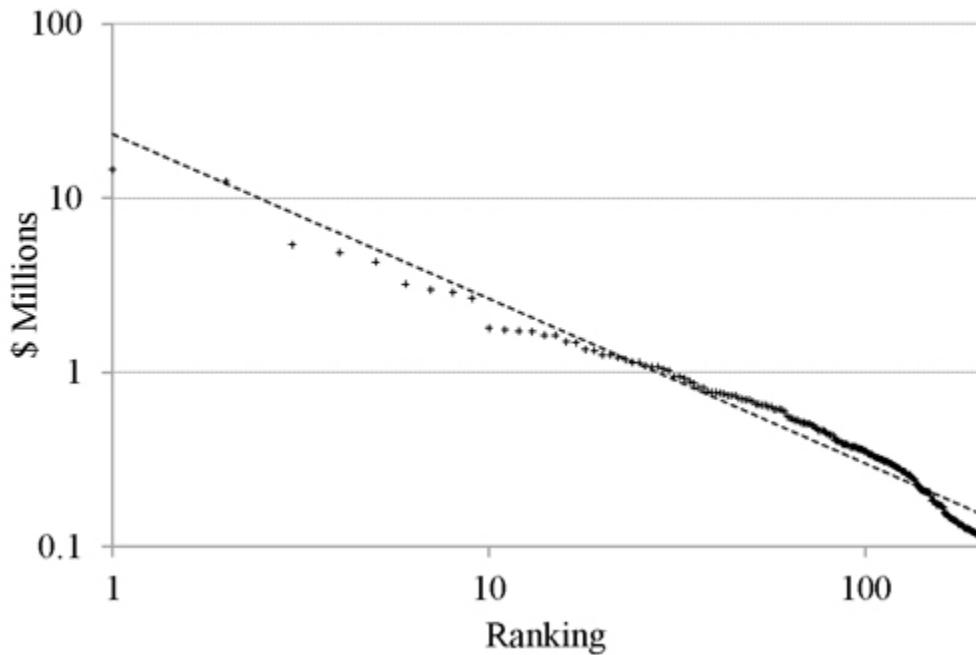
Predictably, Pareto distributions of wealth can be found throughout sports. Consider the distribution of prize money for the world's top 200 tennis players (singles and doubles) on the men's ATP tour from 2013 which is shown below. Just a handful of the world's best take the lion's share of the tournament prize money. In fact, Nadal (ranked 1st at the end of 2013) and Djokovic (ranked 2nd) took nearly a fifth of the whole lot in 2013, whilst the top 10% took half of the total prize pot.

Prize money distribution for the top 200 tennis players on the 2013 ATP tour



Such a plot is characteristic of a power-law distribution, with a long tail to the right containing the majority observations with relatively low scores and a tall peak to the left containing a few observations with much bigger scores. In such a distribution, the overwhelming majority is below the average. The next chart plots exactly the same data but with logarithmic scales. This time we see a fairly strong linear relationship. A straight line on a log-log plot like this is strong evidence for a power law relationship.

As above but scales logarithmic

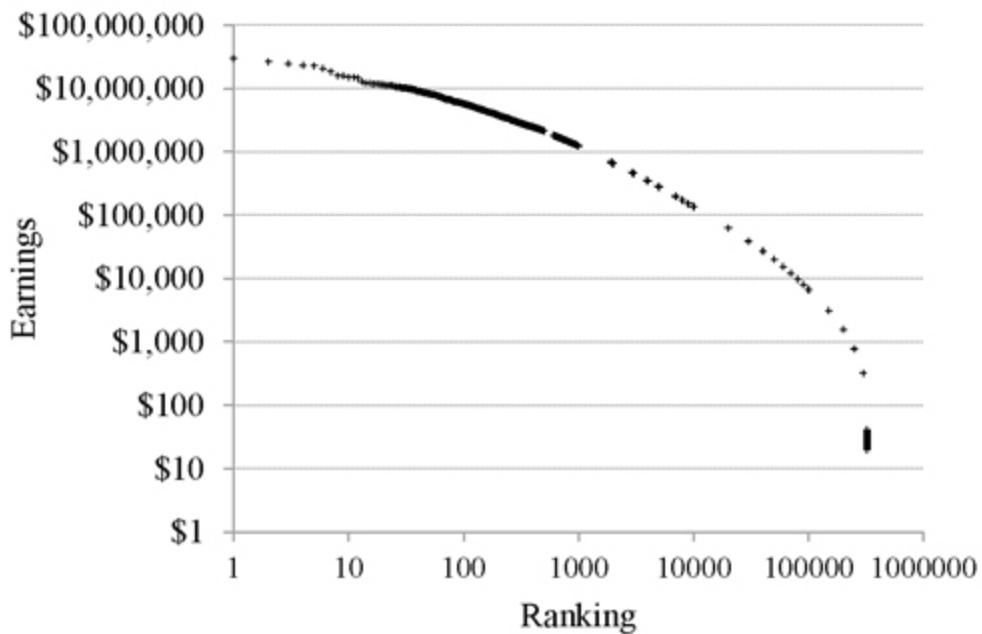


Clearly professional tennis is a ‘winner takes all’ competition. The differences in skill between these top players are probably very marginal but sufficient to translate into huge disparities in prize money. In his book *The Winner-Take-All Society: Why the Few at the Top Get So Much More Than the Rest of Us*, Robert Frank identifies a number of explanations for how such Pareto-type distributions of wealth can arise. A key requirement is that relative, not absolute, performance will differentiate the winner in a competition. Another is the potential for compounded advantage, which we might otherwise describe as leverage. Better tennis players can leverage their slight premium in talent against an opponent over the duration of a game to dramatically increase the chances of victory. Furthermore, because of the knockout nature of professional tennis, such leverage will translate into a dramatic increase in games played, success rate and hence prize money. Analysing 14 seasons of professional tennis match data (2000 to 2013) from the top level ATP and WTA circuits reveals how strong this leverage can be for the best players compared to the rest. Just 10% of players exhibited a positive net win record (more wins than losses). Of those, barely a fifth (so just 2% of the total population) showed net win figures of over 100 games, the highest being Roger Federer at 683 followed

by Rafael Nadal on 507. The rest had net win figures of between plus and minus 100 games with 70% between plus and minus 10. The top players got to play and win so many games simply on account of their leverage. The fact that so few players actually managed to win more games than they lost clearly demonstrates the strength of the ‘winner takes all’ phenomenon in tennis. In this case, as for other Pareto-type distributions, the overwhelming majority had net scores below the average, which of course in this example is 0.

Poker, like tennis, is a ‘winner takes all’ market. The game may be subject to a much larger amount of luck but is played over a potentially much larger number of rounds; any tiny advantage in skill will increasingly assert its authority over weaker players. Similarly, as in tennis, over long periods of competition better players can leverage their earnings potential simply by virtue of playing more and staking more. The chart below shows a log-log plot of all-time poker tournament earnings versus money list ranking based on data made available via the Global Poker Index. As in professional tennis, almost all the prize money goes to just a small proportion of players, with the vast majority below the average. Indeed, whilst the biggest winner, David Negreanu from Canada, had collected prize money of over \$30,000,000 as of 29 June 2015, the median player of the 321,708 in the database had won just \$2,674. Presumably, once entrance fees are taken into account, the implication must be that the majority of players in the list will show negative net returns from their poker tournament play.

**All time poker money list earnings by rank
(as of 29 June 2015)**



Betting does not exhibit the same kind of knockout competition that tennis does although, viewed over the longer term, elements of it are similar. In particular, we could view the cessation of wagering by square bettors on account of their disenchantment of losing as a kind of knockout, leaving sharper bettors to continue to play. Additionally, winning players can be expected to leverage their stakes as their accumulated profit increases, giving rise to a wealth concentration that reflects a Pareto-type distribution. Do we see any evidence of the 80/20 rule in sports betting data? Examining the bwin data set I discussed earlier in the book, it would certainly appear to be prevalent. 20% of customers placed 89% of the stakes. Indeed, the highest staking 1% was responsible for over 40% of all the money wagered, a group that the original research team identified as a cohort of heavily involved bettors. Interestingly, total money bet by a customer did not correlate strongly with the total number of bets they placed, with only 12% of the variance in the latter accounting for the variance in the former. Nevertheless, the same Pareto rule was observed for the number of bets too, with the most active 20% of customers accounting for 79% of all bets. 20% of winning customers collected 83% of net profits. Indeed 9% of all the money paid out by bwin to sports bettors during that 8-

month period in 2005 went to just 6 customers (out of a total of 5,444 winners) with profits of over €10,000. Similarly, 20% of losing customers were responsible for 80% of the net losses, with just 20 of them (out of a total of 35,045 including those who broke even) handing over 5% of all money received by bwin. Without access to the betting prices for the wagers the biggest winners made, it is impossible to determine whether compounded advantage will have contributed to their excessive profitability. In all likelihood, however, and given the similar Pareto-type distribution in losses, large differences in player wealth and risk preference (staking), as well as good and bad luck, were probably the major contributing factors towards these ‘winner takes all’ and ‘loser has to fall’ phenomena, rather than leveraged skill. That’s not to argue that marginal skill differential won’t give rise to a ‘winner takes all’ market, as my earlier examination of compounded advantage demonstrated. Rather, and as Robert Frank recognises in his book, exactly the same phenomena can arise out of marginal differences in luck as well.

Similar power-law distribution of profits and losses can be seen in bwin’s poker and casino data sets. 91% of net profits were paid to the top 20% of the 383 poker winners, with 40% of it going to just 5 players. Meanwhile, for bwin’s 3,062 poker losers, the biggest 20% handed over 83% of the total. For bwin’s casino games, 85% of all net profits were paid to the top 20% of the 458 winners, with a full third going to just 5 of them. Similarly, the biggest 20% of the casino’s 3,764 losers contributed 86% of net losses, whilst a quarter of the total was paid by the top 1%. Casino games, of course, are purely matters of chance. Given that such disparity in the magnitude of casino profits and losses occurs, we are surely forced to conclude that much of the variation seen in the sports and poker distributions also arises as a consequence of luck and staking (the latter a consequence of differences in player wealth). And yet a closer examination of how the highest-staking gamblers performed in each of the three environments yields some interesting observations. The table below shows the variation of the percentage of winners (as defined by whether they made a profit) and aggregated yields with descending stake size percentile¹²². For example, the highest-staking 10% of sports bettors (the 0 to 10 percentile) lost 6% on total aggregated turnover with 22% of them ‘winners’,

compared to losses of 21% and 8% ‘winners’ for the smallest-staking 10% (the 90 to 100 percentile).

Variation of winner percentage and aggregated yield with stake size percentile

Percentile	Percentage of winners			Aggregated yield		
	Sports	Poker	Casino	Sports	Poker	Casino
0 to 10	22%	20%	14%	-6%	-4%	-2%
10 to 20	17%	10%	14%	-10%	-11%	-3%
20 to 30	16%	10%	16%	-11%	-11%	-3%
30 to 40	14%	7%	13%	-14%	-18%	-4%
40 to 50	13%	10%	10%	-14%	-14%	-4%
50 to 60	12%	11%	9%	-15%	-13%	-4%
60 to 70	12%	11%	7%	-17%	-14%	-5%
70 to 80	12%	11%	8%	-17%	-12%	-5%
80 to 90	9%	10%	9%	-21%	-11%	-5%
90 to 100	8%	12%	8%	-21%	-13%	-7%

For casino play there is fairly minimal difference in success rate between the larger and smaller stakers. What variation does exist probably arises on account of the type of games bigger and smaller stakers are playing respectively. Typically, riskier games (that show a bigger variance in returns) attract larger house edges. Riskier games would also presumably attract smaller stakes on average. Hence, positive correlation between stake size and both success rate and yield could be expected. That is to say, people stake bigger on games with smaller house edges, so more of them will make a profit and lose less as a percentage of turnover. For sports, the variation of winning percentage across average stake size is much more marked, with bigger stakers winning considerably more often than smaller ones, and with smaller percentage losses. This would appear to offer some credibility to the notion that more successful players will bet with larger stakes, and by doing so leverage the amount of money they might win.

Unfortunately, I don't think that causal explanation is correct, or at least it does not account for much of the variation. Larger stakes are rationally bet on propositions with shorter odds, and because of the favourite–longshot bias which is strong at bwin there is a relatively greater chance of success. I think most of what we are seeing simply provides further evidence for bettors to be willingly exploited by this pricing inefficiency.

The figures for poker are possibly the most intriguing. I've already considered this gambling market as the one that possibly offers the best chance of consistent success arising through superior playing skills, given its much smaller size compared to betting and investment markets. You may also remember that players who had played more in this sample had a considerably greater chance of success. Variation in winning percentage and aggregated returns across stakes is fairly limited with the exception of the top 10% of stakers. Double the proportion of these 'high rollers' are winners compared to the rest, and aggregated losses for the group as a whole are only a third in percentage terms. Again, however, I have to be the harbinger of pessimism. High-stakes games attract proportionally lower rakes. The highest-staking 10% of poker players in this sample were wagering an average of €200 per session, compared to less than €50 for the next 10%. Additionally, there was absolutely no correlation between the length of play and the stake size. And lastly, if we did concede that high-staking poker players were winning more often, who are they beating? Presumably, the losers too must be staking big to play in these games, and their data must contribute to the aggregate, unless it just happens to be the case that this sample of players contained a disproportionate number of high-rolling winners who were playing against losers who had registered with bwin before February 2005. If this is so, that is simply a matter of chance.

It is finally worth remarking that the actual correlation between a player's average stake size and his or her yield is very small, with $r = 0.037$ for poker and an almost meaningless 0.006 for sports (remember 0 implies zero correlation). Of course, by far the biggest influences on the amount a gambler will choose to stake will be their attitude to risk and the amount of disposable income they have to gamble with. Given the multitude of different risk preferences and the numerous other ways that people can accumulate disproportionate wealth in a 'winner takes all' society, the

sharpness of gambling play is not likely to show a big influence in this sort of data. Whilst power law distributions account for all sorts of variations in the way people gamble on sports, poker and casino games, I see little evidence within the bwin data set that would lead me to change my earlier conclusion, that almost all of what takes place at online betting platforms and poker rooms is simply a matter of chance. Slight differences in luck may very well account for some of the very biggest wins and losses (and evidently they must do in the casinos), but distinguishing these from slight differences in skill is impossible. To reiterate, that's not to say skill won't lead to leveraged success and profitability; the theory behind compounded advantage is quite sound and the performance of characters like Patrick Veitch would support it. It's just I can't convince myself that, at least for the limited data sets I've studied in this book, I'm seeing it. What tiny proportion of genuine high-staking sharps may exist, it seems, is lost in a sea of squares. Betting and poker might not be entirely zero validity environments but I believe they're pretty close.

The Pareto Learning Curve

In economics, the law of diminishing (marginal) returns states that, as one input of a production process is incrementally increased, all other inputs remaining the same, there will be a decrease in the incremental output. It's a clever way of saying you have to do increasingly more for proportionally less extra reward. The law applies nicely to what Nate Silver calls the 'prediction learning curve', which he describes by the Pareto Principle. The Pareto Principle, remember, states that 80% of the outcomes comes from 20% of the causes. We can visualise the Pareto learning curve by means of the following schematic.

The Pareto skill learning curve



That's great; surely this means I only have to make a little bit of effort to get a lot of reward. In an absolute sense that's true. A little bit of effort will mean you're making similar decisions to experts most of the time. Of course, in prediction markets the problem is that you don't get paid for doing the same as everyone else. In competitive environments it's not a question of how good your absolute skills are, but how they measure up against the opposition. Nate Silver (in *The Signal and the Noise*) hits the nail on the head, using his favourite game as an example.

“In poker, you can make 95% of your decisions correctly and still lose your shirt at a table full of players who are making the right move 99% of the time.”

If most players are already at 95%, you need to go further to find an advantage, but when absolute skills or prediction accuracy are already so high, you can see from the Pareto learning curve that incremental improvements will take proportionally much greater amounts of effort to achieve. Betting and investing present potentially even more challenging propositions than poker, since to achieve positive expectation you effectively have to outperform the collective wisdom of the market rather

than just second guess a few hands of cards. The wiser and more efficient those markets are, the higher your skill level and prediction accuracy will have to be. That in turn implies power-law increases in effort to realise this goal. Achieving consistent excessive returns in betting and investing is undeniably a full time job (despite what some people will tell you) and when even professionals prove to be rather bad at it you can see the problem you're up against. Patrick Veitch gave up his maths degree in order to devote the sufficient time and effort necessary to learn all there was to know about flat season horse racing, and for that he secured an advantage of just 17% profit over turnover in a high-variance environment. It's perhaps not surprising, then, given the shape of the learning curve, that so few players in prediction markets manage to outperform expectation. Small differences in skill might lead to big differences in outcome but they also require big increases in effort to engineer. Perhaps investment strategist Donald Luskin understood this when he said: "*I'd always rather be lucky than smart.*"

The Evolution of Winners

The time series below shows the evolution of profits for a sports betting advisory service – PH Sports Betting – that I verified from 2004 to 2012. During that period it advised 1,461 tips, mostly from football, tennis, darts and snooker, and when it closed had a rather impressive yield of 9.75% from median betting odds of 3.00. Those figures, however, don't tell the full story, as the time series makes clear. The second half of the record shows no profit-taking at all. Why did this happen?



The simplest explanation is PH Sports Betting had been lucky, very lucky; indeed the sort of performance one might expect to see by chance just once in several thousand occasions, but which eventually and inevitably started regressing to the mean. Conversely, the second half represented a period of unlucky performance for a tipster that was sufficiently and consistently skilled to show a profitable expectation over the long term. Neither of these descriptions, however, appears to sufficiently capture the dynamism of markets. An equally plausible explanation might involve the disappearance, over time, of skill. That might happen, as I've described earlier, in environments where larger and larger numbers of sophisticated operators begin to cancel each other out – the paradox of skill. The period during which PH Sports Betting operated certainly witnessed an explosion of interest in sports betting. With markets increasing their efficiency, initially skilled performers could find themselves bewildered in the face of disappearing value expectation. Equally, however, it might happen because of an ‘evolution’ of winners. As in other business environments, winners come and go not because they are changing the things they do but because the market itself, like the natural world, changes through an evolutionary process.

In *The 80/20 Principle and 92 Other Powerful Laws of Nature*, Richard

Koch introduces his concept of the ‘business gene’. Building on Richard Dawkins’ theory of memes¹²³, Koch explains the DNA of business as ideas containing ‘useful economic information’; for example the design behind the internal combustion engine or the Microsoft Windows operating system. Business genes represent the knowledge about how to increase wealth. In the context of a competitive market a successful business gene is one that manages to exploit inefficiency where it can be found, for consistent profitable expectation. Koch explains that, just as animals and plants are the ‘vehicles’ for biological genes, carriers through which the DNA code is able to replicate, business genes are carried by the physical apparatus of economic activity: firms, people, assets, products and services. And as in the natural world, those vehicles best adapted to carry economic value, to exploit economic niches or inefficiencies, will be the ones that are most successful and will flourish. Bill Gates, through his Windows software, has successfully exploited a market niche for decades. Of course, such a process is dynamic because the market is dynamic, always changing and, what Michael Shermer, in *Mind of the Market*, calls a complex adaptive system, where individual parts – the players – interact and adapt their behaviour to changing conditions. Inefficiencies, once exploited, tend to disappear, with new ones emerging for which other vehicles may be better adapted to exploiting and propagating those useful (and profitable) ideas. Microsoft Windows has been a success, but it won’t last forever. Its internet browser didn’t. Once the dominant player with nearly 90% usage, today it doesn’t even make double figures. First Firefox and then Chrome became the market leader. Both will also ultimately be replaced by others doing things better than they can, better adapted to propagate those successful business ideas. The market, like the environment, may always be striving towards efficiency but the constant interaction of its players means perfect equilibrium is rarely if ever achieved, and if it is it doesn’t last.

The success of PH Sports Betting did not last forever. Initial exploitation of market inefficiency, if we concede that it was real (and not simply a manifestation of luck), experienced a gradual and perhaps inevitable process of self-destruction, as the private forecasting methods which its owner used were given public expression within the market, leaving them open to exploitation by others. PH Sports Betting may not have changed

what it did, rather it was simply overtaken by other ‘vehicles’ better adapted to exploit what market inefficiency existed. Most businesses end up failing. As in the natural world, failure is the normal condition. So, too, most people who claim to forecast the future will end up as losers, even those who may, for a time, have genuinely been doing something more than merely replicating chance. The market and its players operate within a process of evolution; the only way to remain a winner is to adapt to the ever-changing environment of the market. That’s something that Haralabos Voulgaris, the world’s top NBA gambler, recognised. Having made a fortune since the late 1990s, in 2004 he started to lose. He recruited the services of a maths prodigy to help him design a new forecasting ‘machine’ they called Ewing, and from 2008 returned to winning ways. Nevertheless, Voulgaris has admitted he faces the same issue that all sharps face: the sustainability of his edge, no matter how sophisticated the model that produces it. In 2010/11, Ewing clocked an ROI of more than 6%. By 2011/12, it had fallen to around 5%. As he told ESPN¹²⁴, *“every time you make a bet, you’re educating the people taking the bets. They’re learning the right way to make a line. They figure **** out based on what you’ve already figured out.”*

In markets, as in life, sometimes you have to evolve even if only to stand still. For a time, winners can grow and prosper through the power law processes I have examined earlier, but those same processes act to replace old winners with new ones. The most beautiful aspect of this evolutionary process is that it happens as if by magic. Complex adaptive systems may appear to be designed from the top down, but in fact they are constructed, like the wisdom of the crowd, from the bottom-up through what Shermer calls functional adaptations – what works survives and what doesn’t disappears – with the arithmetical total of success and failure summing to zero. Perhaps, then, at the close of this chapter, we should consider a rephrasing of its title. Winners evidently do take all, but not indefinitely. In markets where new information about price is endlessly and randomly being assimilated by the players, the predictability of inefficiency and the possibility of consistent success, if present at all, are merely transitory. The market, like life, represents a continuous zero-sum dance between winners and losers, where winners do not remain fixed and where the search for

‘true’ value, as game theorist Oskar Morgenstern said, is like a search for will-o’-the-wisp. Most of what happens in a market is random but even the bits that aren’t are subject to forces that make consistently winning in them one of the most difficult things to do.

[120](#) The evidence from American point spread markets discussed in the last chapter would suggest that this might not always provide a completely accurate representation. Recall, when offered an opportunity to act unwisely, squares may very well take it.

[121](#) Miller, M., 2005. *Strategy-based wealth distributions*. Santa Fe Institute. <http://www.santafe.edu/education/reu/2005/files/michaelmiller.pdf>

[122](#) For poker, stakes are per session of play.

[123](#) Richard Dawkins developed the concept of a meme as a cultural parallel to a biological gene and its evolutionary transmission via inheritance, random mutation and natural selection. A meme represents an idea, behaviour, or style that spreads from person to person within a culture; for example via writing, speech or ritual.

[124](#) http://espn.go.com/blog/playbook/dollars/post/_id/2935/meet-the-worlds-top-nba-gambler

A MARKET FOR LEMONS

In 1970, George Arthur Akerlof, an American economist and University Professor at the McCourt School of Public Policy at Georgetown University, wrote a paper¹²⁵, intriguingly titled *The Market for Lemons: Quality Uncertainty and the Market Mechanism*. In 2001, it won him a Nobel Prize. The hypothesis behind it was simple. If a seller knows more about a product than a buyer, such information asymmetry will decrease market efficiency and the quality of goods (or services), creating a market dominated by dishonesty and greed and infiltrated with gullible buyers and crooked sellers, what we might call suckers and sharks.

A ‘lemon’ is British and more commonly American slang for a faulty or defective item. Its use has become popularised in the car sales industry to describe a new (and sometimes also used) car that is found to be defective after it has been bought. Akerlof’s paper used the market for used cars as an example of the problem of quality uncertainty, hence its title, and provided an explanation for the familiar phenomenon that used cars barely a few months old sell for well below their new-sale price. His model of information asymmetry was simple yet intuitive. Assume that some used cars are defective – the lemons – whilst others are of high quality – the cherries. If buyers could tell which cars were lemons and which ones were cherries there would be two separate markets, one for lemons and one for cherries. Typically, however, relevant information about the mechanical operation of the car is hidden and not easily accessible, so the buyer does not know beforehand whether it is a cherry or a lemon. Of course, the sellers know, hence the information asymmetry. The buyer therefore perceives some probability that the car he buys will be a lemon. Hence he is willing to pay less for it than if he was certain it was a cherry. Sellers of cherries, however, will not be willing to sell at discounted prices demanded by the buyers looking to insure themselves against the risk of acquiring a lemon. They opt not to enter the market. The withdrawal of cherries reduces the average quality of cars on the market, causing buyers to further revise

upwards their expectation that the car they are buying is a lemon, and reduce the price they are willing to pay, and so on and so on. Ultimately, in a lemons market there may be no such thing as a fair price. The consequence of information asymmetry is the driving out of the good by the bad, leaving a market filled with devious sellers – the sharks – looking to profit unfairly from desperate buyers – the suckers – hoping to score a bargain.

Competitive, zero-sum, winner-takes-all markets are havens for lemons, particularly when the incentives are of a financial nature. As champion backgammon player, George Sulimirski, told Joseph Mazur for his book *What's Luck Got to Do with It? The History, Mathematics, and Psychology of the Gambler's Illusion*, “*where there's money there's cheating.*” To my mind, gambling can be a positive and deeply rewarding experience, encouraging the player to explore risk and uncertainty and to help them come to terms with a world that so often seems to lack purpose and causality. Indeed, one of the major themes of this book, which I will develop in greater detail in the last chapter, is that gambling can help us make better judgements under uncertainty by educating us about the process of decision making. Sadly but perhaps understandably, many critics of gambling do not see such positives, instead preferring to highlight the deviant nature of its participants. In ‘To Gamble or not to Gamble: is there a Question?’ I noted that religious opponents in particular regard the business of seeking to gain property at the expense of another through skill and knowledge as an intention to cheat or defraud. It would appear that they perceived the same information asymmetry which Akerlof uncovered, albeit from a less tolerant and more extreme moral standpoint. Evidently there is a fine balance between play that makes use of mental, observational or technical skills to build an advantage and information asymmetry that arises out of secret knowledge. Different gambling arenas treat such advantages in different ways. In betting, for example, the value expectation derived from knowing secrets which are not yet priced into the odds is not considered a form of criminal deception, although that doesn't stop bookmakers from refusing the custom of those possessing them. In the financial world, by contrast, such secret knowledge, often arising out of conflicts of interest (for example, a company employee trading on private information about his company, which is not publically accessible) is typically regarded as insider

information, the use of which is strictly prohibited. Of course, sports professionals betting on their own events would represent similar conflicts of interest, and is similarly criminalised. The casino industry would prefer it that any advantage play be considered cheating, although this view is not reflected by contemporary legislation. Nevertheless, the casinos are free to refuse custom from those they perceive as possessing advantage skills.

More generally, the business of exploitation which arises out of information asymmetry in a market is often seen as a reason for doubting and criticising free trade and capitalism. Such criticism, in my opinion, is unfounded. Capitalism, an economic system built on production and trade of privately owned property, is the inevitable end result of a system in which diverse and independent ideas, privately owned and motivated by self-interest, are shared and traded in a manner that produces unintended but socially desirable ends. This is the invisible hand of Adam Smith. When it works properly, it is the ultimate expression of the wisdom of a crowd. When it fails, it does so not because the system is flawed but because its players are so. Later in the chapter, I'll explore why cooperation, trust and transparency lie at the core of a properly functioning market. Whilst not strictly material that, superficially at least, appears to have much to do with gambling, I nevertheless feel that this book would not be complete without at least devoting some attention to the matter of deception and cheating which often provides gambling (in all forms) with its rather sinful reputation. There are, of course, many examples I could choose from, but I have focused on just a handful, and in particular two stories of lemons trading that I have had a particular personal interest in. They deal with the field of advice, which in finance and betting is now as big as the primary markets themselves. Stories, after all, can do a much better job at conveying concepts than statistics are able to.

Financial Lemons

In the autumn of 2010 a member of my family, whom henceforth I will refer to by the pseudonym Jack, purchased some land from the proceeds of his late mother's estate in the English county of Lincolnshire, for a fee in the region of £25,000, through an investment broker called Goldleaf

Associates. James Matthew Browne, the sales consultant for Goldleaf Associates, had told Jack that a buyer had already been found for his investment and that he stood to make a cool £90,000 profit, receiving payment by 15 December once contract negotiations had been completed. Being a kindly fellow, James suggested to Jack that he would be able to avoid payment of tax on the capital gain if he did the following: firstly, reinvest £50,000 of the profit in a further land purchase; secondly, filter the remaining £40,000 through a trust. James suggested to Jack that the trustees could include him and three other individuals, each of whom would be able to contribute their annual capital gains tax allowance. One of those individuals was me. I put it to Jack that James had not the slightest interest in helping him mitigate his tax; rather, he was simply encouraging him to hand over more money for an investment with a promise of incredible returns that was at best improbable and at worst fraudulent. Not only was he being asked to put up the £50,000 before the first profit had been realised, he was also being encouraged to exploit other people's tax arrangements without any intention of asking for prior permission or presumably allowing them to share in any of the gains. I implored him not to make the further investment of £50,000.

Jack is not a financially astute person. Indeed, whilst it has never been officially diagnosed, he probably suffers from Asperger's syndrome, an autism spectrum disorder characterised by significant difficulties in social interaction and non-verbal communication, making him susceptible to placing trust in entirely the wrong people whilst distrusting others who have only his interests at heart. At his request, I spoke to James by telephone on 17 November. He explained that to date no trust had been set up, but he was of the view that everything that was being proposed was perfectly legitimate from a legal standpoint. Should I have any further concerns I should talk to his legal department which would be able to clarify matters. I did not share James' enthusiasm and told him what I thought of him and his scheme in the most flowery language possible. Needless to say, that was the last time that James spoke with me. He was not in the habit of being disrespected like that; poor James.

The following day Jack's solicitor, another of the suggested trustees for this spoof arrangement, spoke with James. Apparently, there had been no such encouragement on the part of Goldleaf to utilise a trust to avoid capital

gains tax; the idea had come from Jack. That's pretty impressive for a man who forgets to buy home insurance. I spoke again with Jack; no, the idea had definitely been Goldleaf's. James, of course, couldn't have cared less whether Jack had ever bothered to go ahead with a tax fraud of his own. There would never have been any trace back to his company in the event that such a thing unravelled with the Inland Revenue, since there was no paperwork to prove such a conversation had ever taken place. Later that afternoon, I attempted to telephone Goldleaf again; the line was dead, allegedly, Jack had been told, because of rewiring in their Mayfair (London) offices. Those offices, of course, were virtual, through which business post could be channelled and telephone calls forwarded. Similarly unprofessional, their business domain goldleaf-associates.com had only been registered on 24 August, less than three months before and privacy-protected to hide the registrant's details. None of this was evidence of a company with a long track record. James Browne presented the image of a respectable investment professional when in fact he had probably just been operating out of his bedsit. He was running what is commonly known in the industry as a boiler room scam^{[126](#)}.

Despite my repeated protestations, Jack went ahead with the payment to Goldleaf of the second investment tranche of £50,000. Of course, 15 December came and went with no sale of his original land investment. The reason, apparently, was because Goldleaf had insisted Jack make an additional payment to them to cover outstanding legal and accounting expenses that were accompanying the successful distribution of his assets. On 15 February 2011, Jack's solicitor telephoned James (his mobile, since the Mayfair office phone remained 'unwired') and asked whether he could categorically offer an assurance that Jack's investments would be realised and transferred the next day if such fees were paid. James was unable to answer and said he would have to refer to his superiors. He evidently didn't bother (of course, there were no superiors) since no return call was made. Later that day, I spoke with Jack. Since December he had evidently been busy handing more money to Goldleaf. Total investment value now stood at around £140,000, and they would be calling him the next day with some possible news about the sale of his plots and the realisation of his profits. A week later, with no word from Goldleaf, I attempted to access their website.

It had disappeared. Of course, there was a completely innocent explanation, Jack said, because Goldleaf was no longer dealing in land banking but had switched their attention to carbon credits trading. I urged Jack to contact the police but he declined, since he felt that he was obviously a much better judge of character than I was.

Goldleaf Associates never sold Jack's investments. The land he bought was effectively worthless. Unregulated land banking by unregulated firms is a common investment scam. No amount of explaining this to Jack over the period he was dealing with Goldleaf made the slightest bit of difference. This was probably a consequence of the Asperger's syndrome. He was sooner prepared to trust a total stranger with 6-figure sums of money invested in unregulated products with a well-publicised association to fraud than to listen to his own family and solicitor. Indeed, he was initially not even prepared to entertain the views of the police when they contacted him in October later that year to explain that James Browne and his accomplice Philip Victor Everhard had been charged with fraudulent trading contrary to section 993(1) of the 2006 Companies Act. Evidently, another of Goldleaf's victims had called time on James and Philip's grubby little activity.

We might wonder what blinded Jack to the obvious reality that he had been targeted as a sucker from the start. A psychologist would doubtless present a multitude of possible explanations: a Freudian desire to lose as a form of guilt-loaded punishment; an expression of childhood rejection (with such reckless investment activity acting as a displacement vehicle for a proxy form of love), a cognitively irrational mindset unable to judge the probabilities of risk, a desire to validate self-worth, or merely a delusion psychosis. I am not a psychologist and my attempts at encouraging Jack to open his mind have been met with complete resistance. His failure to recognise what everyone else around him could see was taking place, however, might have a very simple rationalisation – greed. It is inherently connected with the motivations of James and Philip for perpetrating the fraud in the first place. Goldleaf was selling Jack a lemon. They knew it was worthless; land investments of this nature always are since they invariably fail to acquire the local planning permission that would be necessary to realise any future profit potential. Indeed, in this case it turned out to be sited on such a useless plot with such a significant gradient that it would have been impossible to support any residential construction that

might conceivably have increased its value. Motivated by greed, such concerns were of course trivial. Jack, meanwhile, promised a risk-free trebling of his investment, was unable to learn the age-old lesson of gambling: there's no such thing as a free lunch. Conceivably, he was as much driven by avarice as those willing to fleece him, although I fear the explanation will sadly remain hidden.

Greed, one of Christianity's seven deadly sins, must be the basis for any transaction taking place in a market for lemons, where unscrupulous sellers are deceiving desperate buyers. Indeed, we might very well implicate a few more of them. A slothful pursuit of money for nothing lay at the heart of Jack's complete abjuration of the central principle of investing: the greater the potential reward, the higher the risks are in achieving it. For Jack, a promised trebling of his assets appeared so easy that he had once suggested I was foolish not to follow his lead. Pride: often considered to be the most serious of the sins, and the source of all others, lay at the root of Jack's denial, his refusal to admit that he had been suckered and the maintenance of the belief that he was still right and better than everyone else who was telling him a different story. Anger: as Neil Isaacs (author of *You Bet Your Life*) might argue, perhaps Jack's reckless gambling with his inheritance even represented a kind of 'ritual enactment of aggressiveness', an irrational desire to beat the system and 'win with wits'. Isaacs, furthermore, might regard the loss chasing which followed each deferment of promised returns as an example of a wrathful process of 'exacting vengeance' on the causes of those losses, and a craving to get even. Whether Jack understood those causes to be the individuals behind Goldleaf or his interfering friends and family advising him to stop gambling with them is anyone's guess. Daniel Kahneman, by contrast, would simply regard it as loss aversion. William Douglas MacKenzie and Robert Henry Charles, our theological critics of gambling, would doubtless be crowing: "we told you so."

In 2014, James and Philip finally pleaded guilty to committing fraud and were sentenced to 24 months imprisonment. Philip's was suspended for 18 months with the proviso that he should use the time to pay back his victims. Jack finally had some of his original monies returned to him in spring 2015. James, on the other hand, had taken it upon himself to abscond prior to his hearing. For his secondary stupidity, he was found in contempt of court and had 9 months added to his sentence, the whole of which he had to serve. In

November 2010, I had warned James that such fraudulent practices, as he was obviously indulging in, have a habit of ‘regressing to the mean’. Well, in this case they did ‘regress’. Your past, as they say, often catches up with you. Of course at the time, James didn’t take kindly to the language I used to describe him. That’s understandable; such self denial in the face of personal criticism is of evolutionary benefit, a conditioned response to raise the mood-calming influence of serotonin when faced with unpleasant or stressful situations we want to send to the subconscious. Furthermore, kidding ourselves about the truth helps to kid others (it certainly fooled Jack, although clearly he didn’t take much fooling), since it does a better job of hiding the behavioural cues of lying and deceit. James’ self denial, however, ultimately made him a victim of his own lemon trading. Through information asymmetry and motivated by a fantasy of expectation, Jack hugely underestimated the probability that he was being defrauded. But James too, equally motivated by greed, evidently underestimated the probability of getting caught. In all likelihood, what started largely as an attempt to make some money in an unregulated market quickly spiralled out of control when the full possibility of what he and Philip could achieve was presented to them, in the shape of a vulnerable individual flush with an inheritance and lacking the wherewithal to manage it sensibly.

Conceivably, we might even formulate a hypothesis here: the strength of information asymmetry in a market for lemons is proportional to the potency of excessive greed for both seller and buyer. Without unrestrained greed in either party it is unlikely that a trade in a lemons market will take place. For sellers, that’s pretty obvious why: they’re not looking to rip people off. For buyers too, however, without the desperation to secure what appears to be such a good deal they are less likely to be fooled by the sales pitch. It’s worth reminding ourselves of the tried and tested maxim: if something seems too good to be true, it probably is. Such a hypothesis seems reminiscent of the Seer-Sucker theory¹²⁷ by J. Scott Armstrong, a forecasting expert at the University of Pennsylvania. According to Armstrong, people are willing to pay heavily for expert advice in a variety of fields, including economics, politics, stock picking and betting. The available evidence, however, implies that this money is almost always poorly spent. Few people, however, pay attention to this evidence,

presumably because patterns look nicer than randomness and denial is emotionally easier than change. Armstrong articulates: “*no matter how much evidence exists that seers do not exist, suckers will pay for the existence of seers.*” Sharks understand this theory well, which is why a market for lemons persists for them to trade in. As Steve Forbes, the publisher of *Forbes* magazine, quotes his grandfather: “*it’s far more profitable to sell advice than to take it.*”

Land banking, is of course, not the only unregulated investment strategy that attracts fraudulent practice of this kind. Carbon credits, which allegedly Goldleaf had branched into, agribusiness (including forestry and sustainable crops like jatropha), renewable energy and rare earth metals are all investment vehicles which are used by lemon sellers to ply their trade. All of them sell a message of a sustainability or technological development which encourages investors to believe they are helping to protect the planet or shape the future whilst at the same time lining their pockets. Sadly, Goldleaf Associates was not the only investment brokerage (if you can call it that) to prey on Jack. Many others have done so and probably most of them are connected, possibly even managed by friends and associates of James and Philip. Jack, in all likelihood, had found himself on a suckers list¹²⁸, probably since he started investing his inheritance. In almost all cases, as with Goldleaf, the websites which presented professional and ethical images of fast-paced industries have long since disappeared, as presumably have the monies that their owners have stolen.

One of them, describing itself as a private client brokerage delivering an innovative outlook on alternative investment markets, even had the same Mayfair holding address as Goldleaf. What makes the story all the more depressing is that, through this brokerage firm, Jack was introduced to an independent financial advisor who helped him to secure a more ‘flexible pension arrangement’. Of course, this rather begs the question: why would a reputable financial adviser fail to carry out due diligence on such a company that was referring him customers? After all, it only took a matter of minutes of googling to establish that it was almost certainly a scam. Apparently, however, he said that was something he hadn’t done. More importantly, where is his independence if he has a referral arrangement with investment brokers in the first place? His use of website testimonials which

have remained unchanged for 4 years tells me everything I need to know about his professionalism. At best they represent the trick of survivorship bias (do you ever see a bad testimonial?), at worst, a deliberate intention to mislead. By their very nature they are unverifiable and hence, in my opinion, of limited worth.

Jack spent the next year with his financial advisor and the advisor's solicitor seeking to bring a civil case against James and Philip to recover his stolen funds, advised by them that the police would either be unable or unmotivated to do so. They were wrong and failed spectacularly and predictably in their own stupid quest but were nonetheless happy to relieve Jack of further 4-figure sums of money for work performed along the way. Should we expect anything better from such regulated individuals? One would like to think so, but perhaps I'm just being naïve. Eventually, they gave up pestering him for more, presumably having come to the conclusion that Jack didn't have anything more to pay them with. In that respect, they can't say they weren't warned. Of course, for my efforts in forecasting the inevitable, our friendly financial advisor thought I was the rudest man he'd ever spoken to. The last person who said that to me went to prison.

Another unregulated broker who has happily been swallowing Jack's money has recently issued him with an interim return, not the full sum his investment has been 'promised' to realise, merely a sweetener to shut him up for a while. Of course, the money won't be coming from his profits; there won't be any. It probably isn't even coming from his original investment; fraudulent brokers like these prefer to spend it on having a good time. Rather, it probably comes from new investors sold the same lies as Jack. This is the business of the Ponzi scheme, named after Charles Ponzi, an Italian businessman and early 20th century con artist in North America, who paid early investors using the investments of later ones, rather than from the profit earned by the investment. With potentially a large number of investors being played, the Ponzi is a particularly insidious form of lemons market. Such schemes, however, are not just utilised by small-time unregulated practitioners. Perhaps the biggest financial fraud in US history was the consequence of a Ponzi, perpetrated by one of the most respected investment businesses on Wall Street: Bernard L. Madoff Investment Securities LLC.

Founded in 1960 by Bernie Madoff, an American stockbroker,

investment advisor and once non-executive chairman of the NASDAQ stock market, it originally acted as a market provider, facilitating direct over-the-counter security trades for institutional investors wanting to bypass the traditional exchange mechanism, before later offering its less publicised investment and wealth management business which ultimately became the focus of an FBI fraud investigation. From 1989 (and possibly much earlier), Madoff used that part of his company to mastermind an elaborate Ponzi scheme that defrauded some 13,500 investors out of an estimated \$65 billion and an estimated \$18 billion in actual cash losses, by offering consistently low-risk, high-return investments that seemed too good to be true. Indeed, such was the level of his consistency that some believed it was legally and mathematically impossible to achieve the gains he claimed to deliver. Many Wall Street firms opted not to invest with him. Yet the complexity and lack of transparency surrounding the firm's statements precluded people from making a thorough investigation. Finally, in 2008, ostensibly secure returns became massive losses for Madoff's unsuspecting clients as he became unable to meet the many billions of US\$ in redemptions. Madoff later admitted that the essence of his scheme was simply to deposit client money into a bank account rather than invest it, and use that same money to pay clients when they requested it, either their own or that which belonged to other clients. He was arrested in December 2008, convicted in June 2009 on 11 counts of fraud, money laundering and theft, and is currently serving a 150-year prison sentence for his troubles.

If a fund's performance is shown to be mathematically impossible, what drives investors to throw money at it? Presumably it is the same set of emotions that encouraged Jack to believe that a short term risk-free trebling of capital was feasible and perhaps something that he was even entitled to. These emotions might include the excessive coveting of profit, the pleasure of reward anticipation, and finally the denial (to ease the dissonance) when it appears that something might be wrong. Such a conclusion could be seen as unfairly blaming the victims rather than the perpetrators of fraud. However, as Geneen Roth, one of the victims of Madoff's fraud, portrays herself in her written account of the episode, her culpability lies not in her failure to see through the man but rather her own lack of consciousness about money^{[129](#)}. In concluding that 'enough' isn't a quantity but rather a

relationship to what you already have, her insight is highly reminiscent of prospect theory's relativity of gains and losses. Furthermore, lest we forget, our biochemical (dopamine) rewards system is designed to be transitory. If we demand too much from it, disappointment can be the expected outcome. If money becomes an end in itself, no amount will ever be enough.

Even when a market is legitimate this does not necessary imply that it isn't a lemons market. This is particularly so where professionals collude to sustain information symmetry at the expense of their customers. For example, in 2014 the UK energy regulator Ofgem launched an investigation into anti-competitive practices within the UK energy sector. Whilst stopping short of accusing the big six energy suppliers of collusion over price setting, Ofgem saw possible tacit co-ordination on the size and timing of price rises.

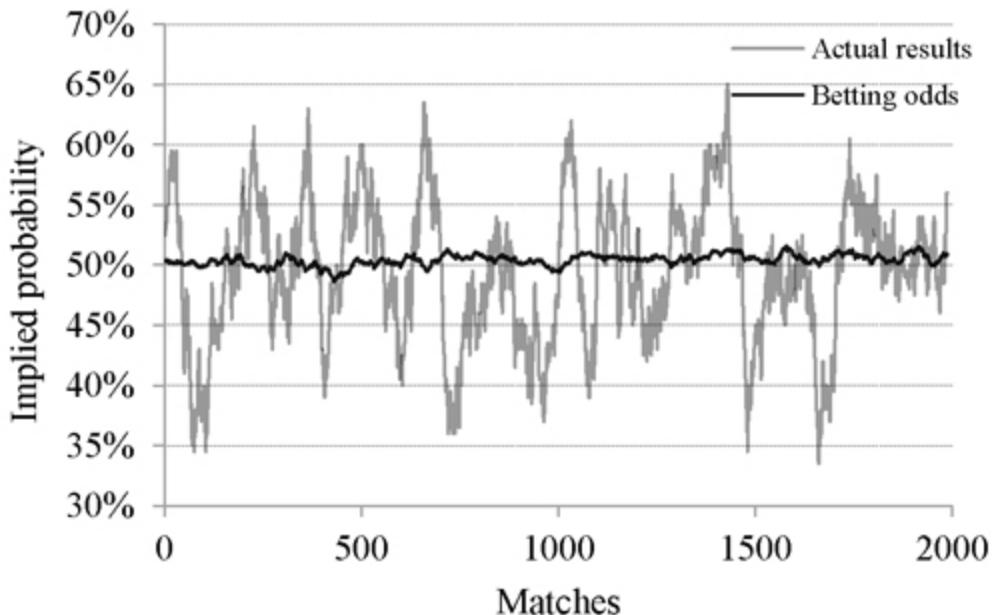
Possibly the biggest example of market collusion by professionals that is still unfolding today concerns something called LIBOR. It has the capacity to make Bernie Madoff look like a petty shoplifter in comparison. LIBOR, or the London Interbank Offered Rate to give it its full title, is an average interest rate calculated through submissions of interest rates by major banks in London that is used for hundreds of trillions of dollars worth of financial products from student loans to credit cards to mortgages to pension funds. Of course, therein lies the problem: LIBOR isn't set by market supply and demand but instead depends on the banks involved accurately reporting the interest rates they'd have to pay to borrow from each other. Hence the LIBOR rate essentially depends on honesty. The scandal arose when it was discovered that banks were colluding to falsely inflate or deflate their rates to profit from trades. Indeed some even described it as a cartel. For example, a bank might purchase an asset that increases in value as interest rates fall and then subsequently report the LIBOR rate at a lower value in an attempt to make a profit from the sale of the asset, effectively trading on inside information. Alternatively, short selling could be a profitable exercise if rates were fixed higher. A move of just 0.01% in the reported interest rate could net a profit of a million dollars or more. Banks also fixed LIBOR at lower values to give an impression of better creditworthiness. In the UK, for example, senior management within Barclays Bank was found to have encouraged staff to lie about LIBOR during the 2008 financial crisis, although there is evidence of rate fixing going back many years earlier. For its misdemeanours the bank has already been fined in the region of \$450

million and may face more. Today LIBOR is based on actual transactions for which fully transparent records are kept. Evidently, it was decided banks could no longer be trusted to set the rates amongst themselves. Collective wisdom, it would seem, had been compromised by a lack of independence. It was not the system that was broken, but its players.

Lemons in Betting

Lemon selling is alive and well in the sports betting advisory industry, too. To my mind, a distinction can be made between two types of information asymmetry: deliberate and blind. I'll look at these in turn. Deliberate or intentional information asymmetry concerns traditional cases of misleading, misselling, deception and fraud, where the seller either offers advice he knows to be worthless or reports a false performance history when his forecasting methods prove to be so. They promote themselves with descriptions like 'professional', 'risk-free', 'sure', 'insider', 'fixed' and 'guaranteed', perhaps endorsed by a few suitably picked testimonials. Many of these bogus tipping services focus on the Asian handicap market where betting propositions are close to even money, and have win percentages anywhere from 70% to 90% (in one case I even saw a record with a 98% success rate). Plainly these win percentages are utterly absurd and statistically impossible given the efficiency of the markets within which they operate; surely, only the most uneducated gambling novices will be fooled. Intrinsically there is just too much randomness in football, as the following chart, comparing a 50-match running average of win probability implied by Asian handicap prices on the one hand, and actual results on the other, demonstrates. Betting outcomes contain lots of noise (luck); short term deviations from expectancy quickly regress to the mean. We know that chaos theory and even quantum mechanics fundamentally make the outcome of a game uncertain. Too much can go wrong for handicappers to do much better than 60% on a long term basis. The same is true for US point spreads.

50-match running average of implied win probabilities for English league matches (2014/15)



Yet given the large sums of money tips are individually sold for, it only takes a few suckers to generate a very decent income. As long as the buyer maintains a net profit after purchasing costs, he will continue to come back for more. Buyers looking for a short cut to financial success will not concern themselves with uncertainty, probability and luck. When a tip wins the causal connection between it and success is made with predictable cognitive ease. Sometimes the seller may attempt to increase return frequency with the guarantee of another free tip if the previous one loses. So the cycle will continue until performance regresses to the mean. When it does, reported betting histories will simply be faked; they are, unsurprisingly, never independently verified. Apparently that's unnecessary since, it is argued, their customers provide the only verification they need. In truth, the tips publicised probably have little relation to the actual tips purchased by buyers, who probably all receive completely different ones to spread the risk and increase the chances of gaining some repeat purchases. Should bad publicity affect business the seller can simply close his website and open a new one anonymously (or perhaps many simultaneously) doing exactly the same thing. As a rule of thumb, any Asian handicap advisory service reporting win rates of more than 65% over long periods should

simply be assumed to be fraudulent, and those between 55% and 65% viewed with suspicion.

Evidently, there are several things a prospective buyer of online betting tips can do to remove the information asymmetry and avoid becoming a victim of fraud. The first, of course, is to never trust a performance that is not independently verified. Secondly, develop a sense of when a performance record looks too good to be true. Gaining an appreciation for the difference between luck and skill and an understanding of market efficiency will be key tasks in that respect. Thirdly, always check the domain registration of an online seller, and ask questions if there are inconsistencies like tips that predate the age of the website or whether the seller has protected the privacy of his identity. Without full transparency the potential for a lemons market will always exist. Fourthly, look out for gimmicks like testimonials which, as I've already argued are largely meaningless, and claims that bets are based on insider information or on matches that have been fixed. Whilst match fixing does take place, thankfully it will represent just a tiny proportion of all events in the world of sports. It's far more probable that the seller is lying than the result being sold has genuinely been manipulated. Even on the rare occasions where the latter is true, success rate is still generally no better than about 75%. Fifthly, use the Wayback Machine at archive.org to see how a seller's website, and particularly his results, looked in the past. Any inconsistencies should be an immediate cause for concern. For example, the football betting advisory service *He Shoots He Scores*, which had been selling since 2010, removed its summer league picks sometime between October 2013 and May 2014. The explanation, of course, is fairly obvious: they hadn't performed very well, and the winter league results looked much better without them. Naturally, the service owner had a different take on it. He used to sell oranges and apples, now he's only selling oranges. A customer who wants to buy oranges doesn't care about apples. Yes, that's true, new ones won't. Old ones, however, who paid for the apples as well, probably do care. When they see a tipster removing picks from a record which lost them money, they may simply come to the conclusion that they weren't actually being sold oranges or apples, but lemons instead.

Another very useful method of testing the trustworthiness of a history of results is to compare the documented betting prices with their closing

market equivalents. As I explained earlier, closing prices should theoretically represent the wisest estimates of ‘true’ outcome probabilities, on account of them representing the largest number of opinions. Consequently, we should expect that sellers reporting significant and consistent profitability should significantly and consistently be beating them. Consider the US sports advisory service Netsportspicks.com. It claims to have made a profit of over \$300,000 since 2005 from \$100 level staking. Probing the owner further I established that this represented close to a 10% profit over turnover from point spreads (which are typically close to even money). With such a consistent performance going back a decade we would, for example, expect an advised price of 1.95 to close at around 1.75, or see an equivalent movement in the handicap. I had the pleasure of verifying the service for about a month at the start of 2015 so decided to see if this seller was appropriately marked by Pinnacle Sports, its primary bookmaker of choice. Even if he personally wasn’t betting the tips, buyers flocking to such a stellar advisory service would cause a market reaction themselves. The results speak for themselves. Of the 102 picks I checked, the average advised betting price was 1.900. This compares to an average closing price for those picks of 1.887, a statistically insignificant drop of less than 1% and nothing like the 10% (plus Pinnacle’s margin) required for a 10% profit expectation over the long term. The price movements were, to all intents and purposes, normally (randomly) distributed. 48 of them shortened (in other words close to half), the rest lengthened or stayed the same. Whilst this represents a small sample one cannot avoid asking the obvious question: if a tipster appears to be so good for so long, why don’t we see a trail of his activity in the market? A look at [Netsportspicks](http://Netsportspicks.com)’ website traffic (via Alexa.com) possibly offers a clue. “*We don’t have enough data to rank this website.*” That’s a kind way of saying it doesn’t get any meaningful traffic at all. Readers can draw their own conclusions from these obvious inconsistencies. I’ve drawn mine: it’s yellow and bitter. Arguably, a bettor capable of generating a near 10% positive expectation in highly liquid American point spread markets from placing several thousand wagers annually at a bookmaker that accepts winners should be able to make a 7 or even 8-figure profit over a 10-year period. He wouldn’t need to waste his time selling his own advice, and in the process increase the risk of giving away his advantage. Evidently, Haralabos Voulgaris didn’t bother,

and he made much of his fortune with an inferior expectation to this one. Netsportspicks also has a testimonials page. It's unchanged in 4 years. Considering the huge profits the owner has (allegedly) been making his clients during that time, isn't it a little strange that no one else has bothered to thank him? For the record, the verification period witnessed a loss of 10.6% on turnover from 192 picks.

Blind information asymmetry concerns those cases where the seller of betting advice isn't aware himself that what he's selling is luck masquerading as skill. Technically, this does not represent a lemons market because the seller is as uninformed as the buyer. From the buyer's perspective, however, in practice it amounts to the same thing. From what we've learnt in earlier chapters about the ease with which illusions of causality, validity and skill can develop in a market which is mostly random, it's unsurprising that it happens so much. Consequently, it is far more prevalent than the deliberate information asymmetry previously described. Its practitioners live in a world of survivorship bias and denial, where performance has not yet fully regressed to the mean, and where suggestions that everything they've done has happened because of good fortune are stubbornly refuted. In that respect they are little different to the mutual fund managers whose performance Daniel Kahneman quietly deconstructed, and yet who continued to believe they were 'somebody'. For them there is usually no intention to mislead, deceive or defraud, merely an unwarranted self-serving overconfidence in their own abilities to predict the future. As the maxim of Albert Venn Dicey (the British jurist and constitutional theorist) suggests: "*A man's interest gives a bias to his judgment far oftener than it corrupts his heart.*" Sadly for a buyer, learning that the seller was not corrupt but merely blind to his own biased judgement will be of little comfort when losses start to accrue.

Denial, however, can take an innocent seller of luck a long way. For the onlooker it becomes increasingly hard to judge whether he really is just blind or intentionally trying to deceive. A common practice in the sports betting advice industry matches closely the tagline of Tom Cruise's film the *Edge of Tomorrow*: "Live, Die, Repeat." The description is pretty self evident. The tipster will first play the character of Nostradamus, the 16th century seer, today famous for his prophecies. Following regression to the mean his role changes to that of Harry Houdini, the legendary illusionist

capable of amazing disappearing acts. Finally, he re-emerges as a phoenix from his own ashes to start again, usually with no admission of previous failures and everything that went before lost in the ether. In my capacity as an independent verifier of betting advisory services, I've witnessed several individuals following this path. Their excuses for disappearing and reappearing under new guises are too numerous to list and most of them are probably lies, but the one that is common to almost all is an underlying refusal to accept that they have no influence on the outcome of a bet. In that respect, despite the overriding impression we might have that they are knowingly selling lemons, this one appears genuinely to represent an expression of an honest blind spot bias. After a succession of failed football betting advisory services and a public unveiling to bring an end to his stupidity on a betting forum, the response I received a couple of months later from one notorious repeat offender was truly revealing.

“Last year I found an interesting pattern on football handicaps that from August until February has delivered beyond 70% winning rate from 140 picks. The pattern I talk about is the same pattern I saw when I was in college and I was betting just for fun. Today I regret why I didn’t take it seriously 16 years ago.”

Not only do we see someone in denial about his own abilities, but evidently the hallmarks of a gambling pathology as well.

Another increasingly popular practice in sports tipping involves the use of advisory service networks, collections of tipsters, services or subscription packages under one umbrella. Earlier I discussed the Betadvisor.com network in the context of survivorship bias. Consider now another collection of tipping services owned by Goran Krljanovic, which includes Betting-Advices.net, Pro-tipsters.net, Besttips4ever.com and Profittipsters.com. Each website adopts a similar business model by offering a variety of subscription packages or tipsters to follow. As for Betadvisor, this allows Goran to spread his risks. Naturally not all services will perform well all the time, but having lots of different ones allows him to rotate his advertising strategy as and when one of them is winning. Typically, he shows up on the Bettingadvice.com and VerifiedTipsters.com forums with news of the latest excellent short term run of form, telling prospective customers how his “experts are on fire” and recommending them to “join and be a winner.” It’s not until you begin to look at the results

that you see a very different picture.

To his credit, Goran has at least attempted to have the work of his tipsters independently verified. His four websites have sold many thousands of tips; analysed altogether they have almost certainly lost money. It's hard to put an actual figure on it because Goran (like Betadvisor) drops poorly performing tipsters, losing their histories from his own websites to enhance survivorship bias, meaning one cannot be completely sure that all past data he has been responsible for will be included for analysis. Furthermore, he has verified different tipsters and packages with different monitoring services at different times, including Bettingadvice.com¹³⁰, VerifiedTipsters.com¹³¹ and Mybigpartner.com¹³² making it quite a task putting it altogether. A broad estimate based on what's more easily accessible and not duplicated across monitors would be around -1 to -2% from close to 7,000 tips (as of the end of June 2015). Perhaps, like TopTipster.com I discussed earlier and presumably other tipster networks like Betadvisor and Tipstersplace.com¹³³, Goran appears unable (or unwilling) to distinguish between luck and skill, arguing that it's completely inappropriate to aggregate performance together like this. It goes like this. Yes, there are losers but we are only interested in the winners, and they are winning because they are winners and that's all there is to it. Causality is obvious and my statistical analysis demonstrating absence of consistency and validity is irrelevant. When presented with such evidence it's far easier just to deny it. After all, there's money to be made selling dart-throwing monkeys, because, as Armstrong's Seer-Sucker theory showed, so many buyers deny the evidence too. Is this blindness or deception? In this case I think probably the former, but you're welcome to make up your own mind.

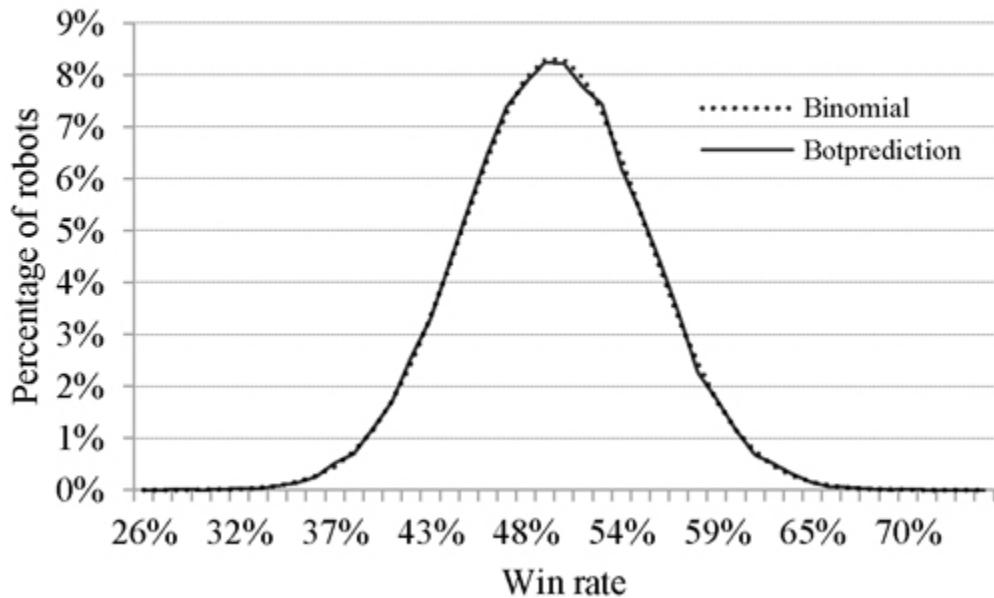
Another case for which it is arguably easier to determine an answer to that question is Botprediction.com, offering "*Successful Soccer Predictions Every Day.*" To be specific, it sells customers tips from football over/under 2.5 goals betting markets with odds between 1.7 and 2.1, offering a choice of 40,000 prediction robots to follow. Its overview provides a useful summary of what it does.

"Our 40,000 soccer prediction robots can be compared to a group of gamblers numbered at 40,000. Statistically, there are always big losers and big winners in that group. Our prediction software enables you to follow those few winners with win rates above 70%! Access our software

and see all the future sports predictions! Follow the predictions, place your bets, enjoy the game and earn your profits! The analyzing software is entirely based on mathematical algorithms and theory of probability - something many people do not take seriously enough, but which, in the long run, always prevails. With our simple and intuitive search platform you can easily monitor all our robots: their stability, trends and most importantly, the predictions of future soccer matches! Join us today and take advantage of our prediction robots and daily winning predictions!"

That, as they say, is the sales pitch. Here is the truth. The mathematical algorithm is just a simple random number generator replicating a coin toss. As such, its robots reproduce the binomial distribution. Having identified a match to bet on (it generally picks about 20 per week) each robot will then randomly assign an over or under prediction. To prove this I compared the performances of its robots to those which would be expected by tossing a coin. In March 2015, I accessed the win rate data for all 40,000 robots, a fairly time consuming task given that the owner wouldn't provide the single database via spreadsheet. We can but wonder why. Each robot history consisted of the same 91 matches played from 29 January to 6 March 2015. The average betting price for these 91 games was 1.88. The average win rate for the 40,000 robots was 49.98%. From that we can estimate that, on average, a robot had a loss expectation of 6%, pretty much in line with the typical bookmaker's profit margin for over/under betting. Obviously some robots did better than average and some did worse. Below is the distribution of the performance of the 40,000 robots, alongside the distribution that would be predicted from the binomial distribution for coin tossing.

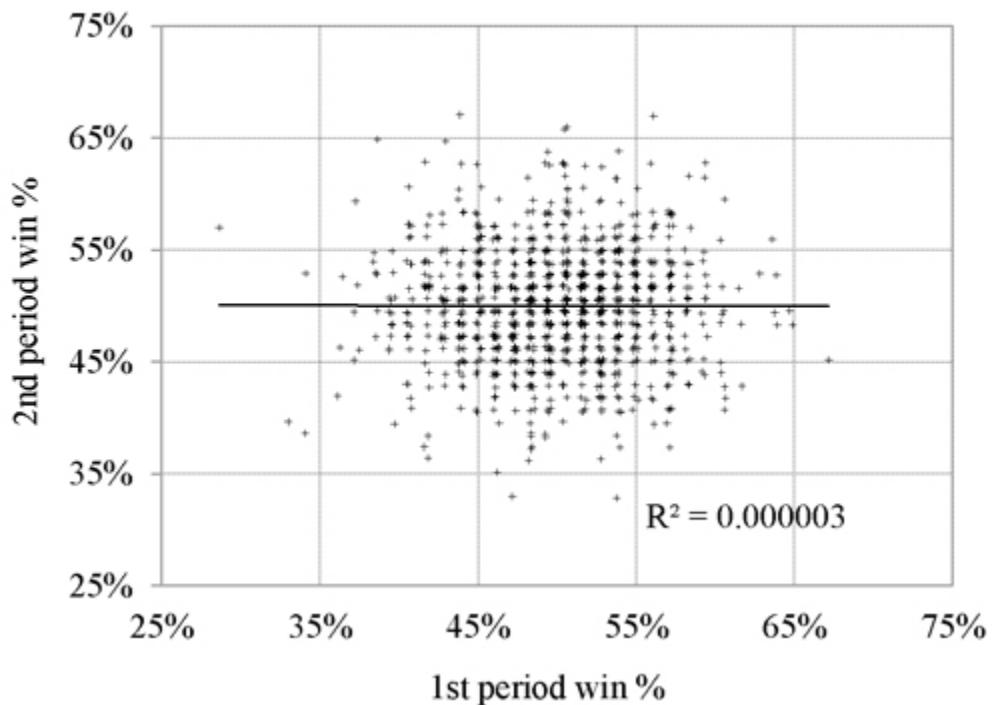
Botprediction's distribution of robot forecasting success versus binomial coin tossing



It's another case of spot the difference. Presumably constraints to server space have precluded having more robots. Theoretically, with 2 to the power 91 of them (about two and half octillion), it could have had one with a win rate of 100% (and one with 0% for that matter), but I suppose then the nature of what was actually going on here would be clearer. Essentially what Botprediction has created is a huge coin tossing experiment that verifies wonderfully the binomial theorem and nothing else. I'm struggling to understand what Botprediction thinks customers will be buying here. It claims that its top robots generate stable and long run win rates that are more than enough to make fair profits. But these profits have happened through nothing more than chance. Beyond the 91 picks, the robots with 70%+ win rates are no more likely to perform better than the robots with less than 30% win rates since, being a Markov process, the next pick in any of the bots has no memory of the last pick. As such, the future performance of any robot will be completely independent of its past performance implying zero consistency, zero predictability and inevitably zero profitability in the long run. To prove this one could repeat this analysis and compare performances between first and second periods. I was not minded to collect all the data again, so the chart below shows the correlation plot for just 1,000 of the robots which is more than sufficient for this exercise.

The second period of analysis is for 91 matches played between 22 May and 1 July 2015. Best and worst performers in the first period showed almost perfect regression to the mean during the second. The correlation coefficient, r , was effectively zero.

Correlation between 1st and 2nd period win %



Is Botprediction intentionally misleading customers, dazzling them with a mathematical sales pitch about algorithms and probability, fully aware that it is really selling chance? Or can we concede that its owner, like the suckers he's selling to, might be similarly blinded by all of this nonsense, seeing causal patterns where none exists? Does he actually believe that the robots with higher win rates are better than those with lower win rates and can be used to predict the future to consistently make a profit? In my opinion, anyone clever enough to replicate a binomial process will surely recognise its worthlessness as a forecaster of betting outcomes. To encourage people to believe otherwise, as the website's summary unmistakably does, is nothing short of deceptive. I asked him to comment about what he thinks he is doing. I never received a response. Silence, as always in cases of honesty and transparency like this, is deafening.

There are numerous other stories of lemon selling in my last book *How to Find a Black Cat in a Coal Cellar: the Truth about Sports Tipsters*. They include tipsters who, like *He Shoots He Scores*, have used ‘creative’ results selection to inflate publicised performance, others who have lived, died and repeated, and finally a few who shamelessly cheated by removing losing results altogether. I don’t intend to repeat them here. Instead, the story I’ve saved until last for this section concerns an independent monitoring service I’ve already referred to: VerifiedTipsters.com. The reason is because I think it represents the most sophisticated example of lemon selling I’ve come across in sports tipping, not so much because of its methodology – that is fairly simple – but in the way that it appears to have been carefully concealed, in contrast to most fraudulent sellers of betting advice where their stupidity makes detection easy. Indeed, it was not for over three years after I first became aware of VerifiedTipsters that I even suspected something might be going on. In the end, however, it was its own statistical data which gave the game away. From that a more detailed investigation uncovered so much circumstantial evidence that in my opinion, as well as the opinions of many others, it is inconceivable that it has not been engaged in what amounts to a scam operation for many years. To summarise, I believe that it has control over many of the tipping services it verifies, manipulating their performances for financial gains in a way that would be classically described by the theory of lemon trading. In other words, it is not independent.

A tips verification service, like my own Sports-Tipsters.co.uk, exists to allow a betting tipster to demonstrate that the performance histories he claims to be responsible for are a true reflection of the advice he gives. Essentially, this involves the seller of betting advice submitting his picks to the monitor alongside his customers. The idea that customers themselves provide the only verification needed is of course absurd; customers cannot be regarded as genuinely independent since they have a material interest in the success of the tipster. Furthermore, tipsters cannot be relied upon to report fairly on their views. Testimonials, as I’ve previously argued, have limited value. The significance of third party independence in verifying the validity, transparency and honesty of a betting record is self evident. Without independence, verification is completely pointless. Of course, it’s not essential for a seller of betting advice to demonstrate these qualities. He

can assume that his buyers will invest in him on trust alone, but as the Russian proverb counsels, there's nothing quite like having that trust verified by other means. Indeed, the maxim 'trust, but verify' seems to make such obvious common sense that Ronald Reagan adopted it as one of his signature phrases during the resolution of the Cold War in the latter part of the 1980s. 'Other means', of course, implies methods that have nothing to do with the party asking to be trusted; otherwise we end up with a kind of trust tautology. You can't completely trust a second party with a vested interest in the first. In seeking to verify the riskiness of mortgage backed securities, Lehman Brothers, for example, turned to the ratings agencies. Of course, by benefiting financially from issuing positive ratings, those agencies were not properly independent and their advice evidently not wholly trustworthy. The lemons Lehman Brothers bought as a consequence of this information asymmetry ultimately led to its collapse in 2008 at the height of the global financial crisis.

As the first to begin verifying the performances of sports betting advisory services in 2001, my intention from the outset was to provide such third party independence for those motivated by the principles of trust, transparency and honesty. Since then, many other monitoring services have followed my lead. Sadly, in my opinion, not all of them have stepped up to the mark. Mybigpartner, for example, appears to earn a share of the subscription fees for some of its verified sellers, and to my mind this compromises its independence. Blogabet.com permits verified tipsters to reset their records an unlimited number of times if they are free. Surely, however, this is completely contrary to the principles of transparent verification, since conceivably it's just another version of 'Live, Die, Repeat'. The monitoring service Verifybet.com (now closed) was shown to have strategic partnerships with some of its monitored services. One of those was, for a time, part owned by a serial scammer. The other owner has since been managing a betting investment fund (Savingonsports.com). In January 2016, customers began reporting the freezing of their assets, some of which were 5 figures. It had all the hallmarks of a failed Ponzi. Others like TipsterConnection.com and Verifiedbets.com leave such obvious clues to their fraudulent activity, including monitored services with 'impossible' profitability and shared common Google Analytics account IDs and/or domain IP addresses, that it's a wonder how anyone falls for it. Evidently

they do since TipsterConnection has been operating since 2009. Whilst we should never underestimate the power of denial, it would appear that this also applies to stupidity as well. I'm unclear as to who is the biggest offender in that respect: the monitor or the suckers falling for these scams. The blindness induced by the seer-sucker effect means the necessary due diligence prospective buyers of 'expertise' should always carry out is often overlooked.

Provebet.com, similarly, is simply just an advertising vehicle for its own advisory services which all share the same IP address. How ironic given that it urges purchasers of betting advice to check that services have been independently proven. The handicapping review service Cappertek.com formerly shared its IP address with the 'independent' monitoring service WagerPolice.com and several of its verified tipsters. There are numerous accounts to be found on the internet referring to Cappertek's dishonest activity including record manipulation and extortion for more favourable reviews.

Ironically, I was first introduced to the work of VerifiedTipsters in 2011 when reporting on their forum the activities of a fraudulent seller of betting advice, Goaloverunder.com¹³⁴. Indeed, VerifiedTipsters devote a specific subsection of their forum to allowing users to report cases of suspected scamming and fraud. The report I filed about Goaloverunder dealt with the manipulation of results compared to those VerifiedTipsters had published, and subsequently the creation of a fake monitoring service, Checkinsports.com, as an attempt to hide further duplicity. Like TipsterConnection, Goaloverunder gave the game away by using the same Google Analytics account ID for its 'independent' monitor. In retrospect, allowing its users to report on cases of fraudulent tipping offers the perfect cover for a monitor doing the same thing. Giving the impression that you care about honesty and transparency must surely mean that you follow those principles yourself, right?

It was not until December 2013 that my suspicions began to be aroused when I came across a forum post¹³⁵ where the management of VerifiedTipsters had proudly announced that it had reached the landmark figure of 60,000 verified tips since 2007. I decided to ask for the aggregated yield for those tips. The figure was 5.1%. This is considerably larger than

the equivalent figures for data sets of tips I reviewed earlier. Pyckio's average was around -2% (on account of using just one bookmaker), whilst my own and that of Oddsportal were around 1% (since the use of all bookmakers was permitted). As argued, those figures were predictably an expression of the efficient market hypothesis. Across a large number of bettors, the vast majority simply tossing coins randomly, value expectation should be somewhere close to zero where bettors have access to best market prices. Naturally, there are three possible explanations for this significant discrepancy: 1) the services VerifiedTipsters attracts for verification are disproportionately better than elsewhere; 2) this sample has been exceptionally lucky; 3) the figures are manipulated. The first option seems improbable, for we would need a good explanation as to why VerifiedTipsters alone has managed to achieve what no other network has (without cheating), given that it passively accepts tipsters for verification (like Sports-Tipsters did) rather than actively recruiting them. Regarding the second option, given that my own verified records contained a period between 2005 and 2008 of 70,000 picks with an aggregated yield of 3%, I was not entirely prepared to discount it. Others at the time, however, were less tolerant. BettingXpert.com, one of its monitored services, had already departed arguing that some of its tipsters, belonging to the same umbrella groups, were reporting performances that appeared to be beyond the realms of possibility.

The issue was revisited in August 2014 when VerifiedTipsters reached the landmark of 80,000 verified tips¹³⁶. The aggregated yield was 4.8%. This time I decided to probe a little further by accessing VerifiedTipsters' own records it makes available through their main website. I counted 78,617 picks (the difference, presumably, can be explained by bets that had been void and not included in the summary figures) from 295 service records which showed a yield of 6.8%, 2% higher than that reported by the management. (It subsequently acknowledged that it might have made a mistake reporting the lower figure.) When standardising for the variable stake sizes that different tipsters use, that figure actually rises to 7.1%. For me it was no longer conceivable that this represented something honest. It was time to do a little detective work.

VerifiedTipsters claims to have grown out of a Ph.D. thesis completed by

the then website manager, a Mr. Greg Wilson, on internet sports gambling under the supervision of Professor Neil Keegan, back in 2007¹³⁷. One of the issues raised in Greg's thesis was the increasing number of fraudulent services that were selling picks over the internet. I asked him if I could access a copy of his thesis, given that it was work close to my own interests, and perhaps talk with him and Professor Keegan about their research. For my efforts, I was banned from the forum. As confirmed by other leading academics I approached in the field of gambling and betting economics, there is almost certainly no Professor Keegan and no Ph.D. thesis. In my opinion it has all been concocted to give credibility to VerifiedTipsters' project. In fact, the reference to Mr. Wilson's Ph.D. was not actually added to the website until the spring of 2013, nearly 6 years after it was first opened. Why the wait? Furthermore, academics embrace opportunities to talk about their work, not censorship.

I continued my investigation into VerifiedTipsters publically via the BettingAdvice.com forum¹³⁸, inviting others to join me in following various lines of thought to test the allegation that was now being made, that in fact it was composed of two parts: its own (in-house) paid services that manipulate records when they have no subscribers and some honest services to dilute the effects of the first group and provide an air of credibility to the project. Being a monitoring service, such a strategy would be almost impossible to prove without admission that it was taking place. When one customer joins any manipulated service its performance will be reported honestly. When they leave, the monitor is free to embellish the results until new ones are attracted. The services I now suspected to be involved – BetAttitude, EliteTipsters, EuroPunters, PanosKnowsBest, ProBetTrader, Tzogosteam and VIPTipsters being the main ones – invariably had high subscription fees and more often than not charged by the tip or small bundles of tips. Consequently, subscribers would be few in number and the length of a typical subscription short, providing plenty of opportunity to enhance the results when no one is following. So long as there are enough services participating in the scam, this will guarantee a very healthy income.

In addition to their high fees, suspicious services had a number of other aspects in common. Firstly, they each had several subscription packages

available to choose from, usually defined by the specialist tipster who was running it. Consequently, I counted at least 58 different records that could have been manipulated; there may have been others. Secondly, the number of monthly picks was almost always small, averaging around 15 per month compared to nearly double that for the rest. Presumably, when tipsters are simply guessing this helps to increase the variance in reported monthly yields and ensures more of them will show profitable returns even for periods during which services have subscribers and reporting of results is honest. Thirdly, every one of these 58 packages was in profit. Aggregated together they accounted for a yield of 14.5% from 34,543 tips, or about 44% of the total number verified at the time I performed the analysis (14 September 2014). By anyone's reckoning such an independently verified return on investment from average betting odds that were close to evens would be truly unprecedented. The remaining 237 services, of which 116 (or 49%) were profitable, exhibited an aggregated yield of 1.5%, much more in line with other data sets I have reviewed in this book. Conceivably, if there are still a few more manipulated records I've not accounted for within this sample, removing those would lower that figure further. Fourthly, first month performance for a large proportion of the tipsters I suspected was profitable. New advisory services generally take a bit of time to acquire some customers. Many bettors like to wait and see how a tipster performs before taking the plunge with money. The first month is the one least likely to have any, and hence is the month most likely to see manipulation. In fact, 53 of the 58 records exhibited profitable first months. *A priori*, if tipsters are guessing, we should expect close to half to be so, as was the case for the other 237 records (with 129 or 54% showing positive first month returns). The probability of this happening by chance is about 1 in 50 billion. Even if we accepted that these 58 tipsters were consistently capable of returning a 15% yield from betting 15 even money picks per month, the odds that so many would be profitable in their first month are still in the region of about 700 to 1. Finally, the start of verification almost always coincided with the initiation of the service. Around half of the tipsters I have verified already had pre-existing records. Indeed, this survivorship bias is, as I've explained previously, precisely the reason why so many then opt to seek verification in the first place. What is more, almost none of these 58 services showed their own results, but instead

directed users straight to VerifiedTipsters.

A valid counterargument to this circumstantial evidence is that I've simply mined the VerifiedTipsters database to find excessively profitable services to create a case. Of course, we could do precisely the same in any large data set of tipsters and find samples with above average performance. Presumably, however, such samples would lack the other similarities and consistencies I've highlighted. Moreover, we are still left having to explain the extraordinary 7% aggregated yield that flies in the face of standard efficient market theory. Nevertheless, if other evidence existed, this would surely support the case being built. Fortunately, there is plenty of other circumstantial evidence that doesn't simply rely on statistics.

All the websites suspected of participating in this alleged scam have similar website designs. Furthermore, most have low traffic (as estimated by Alexa.com) and just two significant inbound links. The first is the BettingAdvice forum page where I was conducting my investigation, the second was VerifiedTipsters. Presumably, services genuinely independent from their monitoring company, as others do, would show more inbound links than this. There were similarities in content too. Like VerifiedTipsters, ProBetTrader.com also indulged in some academic eulogy with all four of its tipsters completing degrees in either Greek or UK universities in economics, finance and accounting. Sadly, British data protection rules have meant I cannot determine whether these degrees were ever completed by the named individuals. Almost certainly, like Greg Wilson's Ph.D., all of it is a massive fabrication. The owner of the website, Athens-born Mr. Gregory Pappas (note the first name) is alleged to have graduated from the University of Piraeus and the London School of Economics, before working as a 'stoke-broker' (his spelling, not mine) and more recently a finance manager. Like Professor Keegan, who became a grandfather in September 2013 (as announced on the VerifiedTipsters forum), Gregory Pappas evidently also liked talking about his family, being married with three children. Are these coincidences? Hmm.

Gregory Pappas, in fact, is the domain registrant for ProBetTrader.com. He's also the registrant for BetAttitude.com, a partner website to ProBetTrader, and is listed as residing in Athens, Greece. Both websites were registered in the summer of 2006, just over a year before VerifiedTipsters.com was anonymously registered. Both services were some

of the earliest to be monitored by VerifiedTipsters and they continue to be so today. Another was called DailyPunter, who at the time used a Blogspot domain. He was later to join Tzogosteam. ProBetTrader has been built with the Website5x.com design software. Nothing interesting in that you might say, until we discover that so were EuroPunters.com, PanosKnowsBest.com and Tzogosteam.com. All were registered anonymously between 2008 and 2009 and all continue to be monitored today. Furthermore, EuroPunters.com and PanosKnowsBest.com use the same company to protect domain privacy: Privacy Protection Service INC, Nobby Beach, Queensland, Australia, exactly the same as for VerifiedTipsters.com. Tzogosteam.com uses a different one, but the registrant is from Athens. ProBetTrader also uses a forum supplied by Simplemachines.org, exactly the same as VerifiedTipsters. Are these coincidences? I'm really struggling to believe so.

Possibly a more conclusive piece of evidence came by way of Netcomber.com, a “*fingerprinting tool... mapping owner networks all over the internet [to] tell you what other sites are being owned by the same person.*” One of VerifiedTipsters’ monitored services goes by the name of Jeff’s Horse Racing Selections¹³⁹. Over the years, horse racing appears to have been a specialist favourite of some of VerifiedTipsters’ participating advisory services, presumably to attract the UK customers. Both EuroPunters and ProBetTrader offered it as a package. EuroPunters dropped their racing service when ‘Brian’, an elderly gentleman from the UK in charge of the service, decided to retire. Given that his yield performance was superior to Patrick Veitch’s, I wanted to get in touch and try to encourage him out of retirement and work for me. Sadly, EuroPunters weren’t able to assist; funny that. Anyway, Jeff’s been going since December 2013, and performing fairly well. He doesn’t use a website, but rather operates simply via email. That’s a pattern that’s become increasingly common at VerifiedTipsters since 2014. Today, he prefers horsenugger@gmail.com, but has previously used horsenugger@mail.com and horsenugger@hotmail.com. One day in January 2015, I happened to be putting VerifiedTipsters.com in to the Netcomber machine and outputted four email addresses sharing the same email account:

- billing@verifiedtipsters.com

- support@verifiedtipsters.com
- service@verifiedtipsters.com
- horsenugger@hotmail.com

Well, well; is that the smoking gun? It's hard to say for sure, since it's unclear just how reliable Netcomber is and exactly what 'shared account' actually means. Nevertheless, at the very least it is further strong circumstantial evidence for a lack of independence. Furthermore, horsenugger@mail.com also shows up on Boxwind.com¹⁴⁰, a search engine that collects email addresses associated with every website. If you choose to register with them, you can find the following addresses associated to VerifiedTipsters.com:

- horsenugger@mail.com
- tradingnotbetting@gmail.com
- support@wontimes.net
- soccersyndicate@live.com
- mathematicianbet@hotmail.com
- infosocceradvice@gmail.com
- support@verifiedtipsters.com

The first six are all addresses connected to services that have been or still are monitored by VerifiedTipsters. None of them has its own website. Search for any of them in Google and the only place you will find their email address is on VerifiedTipsters. In my view, these are evidently all accounts created by the monitor, or individuals associated with it, for the purpose of selling tips, and cannot therefore be independent of it, as is claimed.

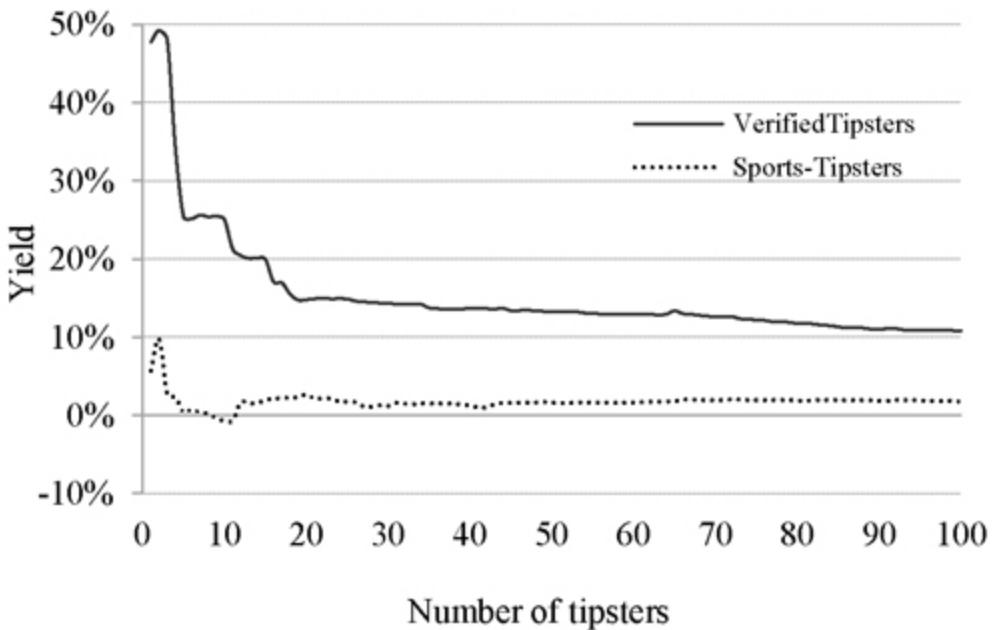
Jeff also says he's open-minded to business proposals. I got in touch on several occasions asking if I could buy his tips and whether he'd like to work for me. Evidently, he was not too open-minded towards my proposal. The only reply I ever received from him was to tell me that he'd found all my mail in his spam folder and that he didn't accept new customers. I wonder why he didn't say that on his VerifiedTipsters profile page. In fact, I attempted to contact a number of VerifiedTipsters' suspect services with

requests to purchase tips. Only one other ever bothered to respond. That was Tzogosteam, which wondered why I would want to buy from a site that I believe to be a scam. Of course, I had never told Tzogosteam that I did. Evidently someone else had told them. How probable would that have been if it had been truly independent from its monitor I was investigating?

Judging by all this circumstantial evidence, it would appear that the independence of VerifiedTipsters has been compromised from the very beginning. Indeed, I think it is not unreasonable to suggest that its very rationale came, not out of a Ph.D. researching online scamming and fraud in the world of sports tipping, but to commit it itself. An internet search yields links to other pages, some dating from several years ago, where rumours of VerifiedTipsters' dishonesty are discussed. Additionally, other parties have informed me that the email they used to register their service for monitoring appeared to have been immediately forwarded to other services for the purposes of spamming. And one individual contacted me bemoaning the fact that he'd lost as much as €10,000 purchasing and betting tips from services monitored by VerifiedTipsters. The story was always the same: their "*effectiveness*" dropped as soon as he'd made a purchase. He also speculated why the odds he backed did not shorten if these services were so good: "*because nobody except me was playing them.*"

Evidently, it took some time for the honest part of VerifiedTipsters' verification service to develop sufficiently to conceal the clues of its manipulation, and presumably before that nobody had been bothered to look. Perhaps it only dawned on the management slowly that having other legitimately verified tipsters amongst their own offered an excellent means of hiding the giveaway statistics. The chart below, for example, shows the cumulative yield for the first 100 tipsters verified by VerifiedTipsters, ordered chronologically by when they first began verification, for tips up to 14 September 2014. By comparison, I have shown the same for my own verification service, Sports-Tipsters.co.uk. It's fairly obvious which one is natural and which surely contains an abundance of manipulated data.

Evolution of cumulative yield for the first 100 verified tipsters



If real, just how much revenue could have been realised by such a scam? A back-of-an-envelope estimate reveals a comfortable living can be achieved. Assuming no more than a couple of customers is subscribing to a service at any one time, and for only 6 out of 12 months, 58 different services doing the same thing typically charging €100 for a basic package of tips could realise a significant 5-figure annual income. This may well represent an underestimate. Is it any wonder that the original suspects who joined in the earliest years are still going strong today?

More than a year since I first asked to see a copy of Greg Wilson's infamous Ph.D., there has still never been an admission of guilt (and no Ph.D.[141](#)). There's never likely to be either, given the difficulty in catching the perpetrators of such charlatanism. Denial is always the favoured response. As stated before, it eases emotional disharmony and it encourages others to believe the denier. As with those financial lemon sellers James Browne and Philip Everhard, such is the power of self denial that those responsible for VerifiedTipsters probably no longer even believe that what they've been doing is wrong. The management once said: "*Did you ever witness any change on a tipster's records even years after the time of your subscription?*" Of course, that's simply stating the obvious. Given the

nature of such a scam, we're not likely to see any changes because manipulation only takes place when no one else is looking. As with Lance Armstrong, seven-time Tour de France winner since stripped of his titles, who famously and consistently argued that he'd never failed a drugs test, absence of proof is not the same thing as proof of absence.

Cheating and Greed

Why is it that some players in information markets like stock trading and sports betting are willing to cheat, gaining an unfair advantage at the expense of their opponents? Writing in *Psychology Today*¹⁴², Robert Sutton (Professor of Management Science at the Stanford Engineering School), speculates that the more pressure people face for performance, the more likely they are to cheat. Steve Levitt, whose research into unwise point spread bettors we reviewed earlier, has shown that, when teachers' pay is linked to the performances of their students on standardised tests, they (the teachers) are prone to cheat. Arguably, competition and as a consequence the propensity to cheat is at its most intense in 'winner takes all' environments. Research published in the *Journal of Economic Psychology*¹⁴³ has highlighted the dangers of cheating in competitive environments with 'winner takes all' incentivisation. Asking 65 participants to complete a set of on-screen mazes, the researchers Christiane Schwieren and Doris Weichselbaumer paid those in one half according to the number they completed successfully whilst rewarding those in the other half only if they completed more than five other members of their sub-group. Whilst subjects in the competitive group performed no better than the other they proved much more willing to cheat, not only by taking advantage of maze solutions the researchers made available to them but also by lying about the number of mazes they had completed. The poorer the performance, the greater was the willingness to cheat. Indeed, only one person who won a tournament did so as a result of cheating. Schwieren and Weichselbaumer concluded that "*poor performers either feel entitled to cheat in a system that does not give them any legitimate opportunities to succeed, or they engage in a 'face-saving' activity to avoid embarrassment for their poor performance.*" Given the Far Eastern origin of many internet betting scams,

where honour and shame form the basis of moral thinking, one might reasonably wonder whether there is some form of face-saving contest operating between competing sellers, where cheating others out of hard earned income is of secondary consideration to the avoidance of the shame of failing to appear successful.

Of course, all this begs the question: what motivates people to perform, succeed and be winners? At the risk of sounding too reductionist, the answer is surely to be found in the theory of evolution. This is, after all, the most important of zero-sum competitions where success in the acquisition of food, safety and sex ensures a greater probability of survival, both of the individual and its selfish gene via reproduction. Human beings, like other organisms, are hardwired to seek strategies that will deliver successful outcomes in that respect. Higher order needs associated with self-esteem and self-fulfilment, which include feelings of success, merely take the place of primary physiological needs once those have been met. As we learnt earlier, our endocrine (hormone) system is the chemical motivator and regulator behind this pursuit of winning and success and in particular dopamine, responsible for reward anticipation and realisation of goal-directed behaviour. The transitory nature of our 'happy chemical' response, however, inevitably means that enough is never enough. In the natural world where organisms compete for limited resources, achieving success may be fleeting and sporadic. When 'happy chemicals' dip, organisms suffer the side effects of renewed stress in the form of cortisol, the 'unhappy chemical'. Evolution has ensured that the craving response is adaptive. Essentially it means that organisms are constantly alert to removing stresses via seeking rewards.

In a more comfortable contemporary world, where it's much easier to have what we want, however, demanding more and more is no longer a healthy pursuit of goals to achieve an evolutionary advantage but an excessive and maladaptive craving that has become an end in itself. Such craving can lead to habituation and the stress of disappointment of not getting enough. Activating our happy chemicals more often increases the amount of cortisol. If the stress of disappointment becomes chronic, cortisol levels will remain elevated. In turn, we are motivated to repeat the behaviour we expect to make us happy, since cortisol increases the sensitivity of the nucleus accumbens to release dopamine and reduces levels

of serotonin. Thus, we enter a vicious cycle of inescapable stress, characterised by the repeated feeling of failure to fully actualise our goal-directed behaviour, and a feeling of being out of control. The more we repeat something we think will make us happy, the less happy it will actually make us. Such repetitive but ultimately unrewarding behaviour is commonly called addiction. The list of candidates is endless: drinking, smoking, drug taking, shopping, travelling, socialising, sex, eating, exercise, thrill-seeking and, of course, gambling. Really, we have just become addicted to dopamine. We might also call this greed, or to be more precise, excessive greed. As George Sulimirski tells Joseph Mazur (*What's Luck Got to Do with It?*), “*there's nothing wrong with greed, unless it's uninhibited; nor with risk, unless it's reckless.*”

Of all the things that one might be addicted to, there's nothing quite like the pursuit of financial success, both in terms of the numbers of people craving it and of those willing to play less than fair to achieve it. Characteristically, there is never an end point to a quest for wealth, because the dopamine release with each new accumulation of money is temporary. Moreover, for those who are successful, winning itself as Ian Robertson, Professor of Psychology at the Institute of Neuroscience, Trinity College, Dublin explains can become physically addictive too¹⁴⁴. Success changes the chemistry of the brain, making you more focused, smarter, more confident and more aggressive, and the more you win, the more you will go on to win. It's as if nature primes winners to keep winning and, for that matter, losers to keep losing. Neurologically, increased aggression and an expression of power reveal themselves through increases in the hormone testosterone, the male steroid sex hormone and one typically invoked in motivational behaviour. Higher testosterone levels are typically associated with risk taking, loss chasing, greed, aggressiveness, selfishness, competitiveness, revengefulness, overconfidence, an enhanced sense of entitlement, and even psychosis. John Coates, a trader-turned-neuroscientist at Cambridge University, has found that a trader's morning testosterone level predicts his day's profitability¹⁴⁵. With testosterone levels typically seven times higher in men than in women, small wonder that most gamblers are men and hormones, not sexism, explain why those pursuing a career in finance are predominantly male.

Elevated levels of testosterone in turn boost dopamine and the quest for rewards. This power-primed ‘approach mode’ towards reward, success and winning is, however, moderated by an ‘avoidance mode’, where mood is low, stress and anxiety are high, and aversion to risk and undesirable outcomes elevated. Such approach-avoidance conflict typically occurs where goal-directed behaviour can have both positive and negative outcomes, as is the case for environments where decision making takes place under uncertainty. High levels of the stress hormone cortisol have been shown to be correlated with avoidance motivation behaviours, and can moderate the action of testosterone. Coates similarly found that a trader’s cortisol levels rise with both the variance of his trading results and the volatility of the market. Hence, whilst testosterone may contribute to economic return, cortisol is increased by risk. With both hormones thus implicated in behavioural feedback effects, changes in their levels may shift risk preferences and affect a trader’s ability to engage in rational choice. Robertson suggests that the wild oscillations of financial markets may partly be the result of traders’ brains lurching between these two modes of approach and avoidance. This is the stuff of irrational market sentiment driven by fear and greed. Presumably, fear and greed lie at the heart of other forms of maladaptive gambling behaviour as well, like loss chasing and delusional perceptions of predictive skill.

Jayne Barnard, Cutler Professor of Law at William and Mary Law School, has proposed that an overabundance of testosterone in men may even explain a propensity to cheat or commit fraud¹⁴⁶. The testosterone-induced ‘winner-effect’ just described, where success and the hormonal response are intimately engaged in a self-perpetuating feedback, often leads to rash, ill-considered, and dangerously risky behaviours in animal populations. Barnard believes that a similar progression may be seen in human competitors, including both athletic environments (for example, Lance Armstrong, the disgraced professional cyclist) and in business environments (for example, Bernie Madoff). A sense of power has also been shown to be linked to an increased propensity to cheat, particularly when people feel they are unobserved¹⁴⁷. As Ian Robertson has supposed, a generation of deregulation of stock markets evidently meant that the entire financial industry became locked into the neurological ‘approach mode’,

where a culture of excessive compensation (through bonuses), the mistreatment of customers, the misselling of products (for example, mortgage payment protection insurance) and the deceitful manipulation of markets (for example, LIBOR) was allowed to flourish. Furthermore, when given power, people set ethical standards much higher for others than they do for themselves, in an unhealthy mix of self-serving attribution bias and self denial. The strategy of VerifiedTipsters, chastising others for scamming whilst ostensibly indulging in the very same, would appear to adequately fit that interpretation. Given the relationship between testosterone and power, this relationship should hardly come as a surprise.

Squeezing the Lemons: Cooperation and Trust

Selfish genes, however, do not necessarily imply selfish individuals. More specifically, possessing a genetic propensity to survive is not necessarily best served by gene-carrying individuals behaving selfishly, as lemons sellers seem predisposed to be. For social organisms like human beings, unquestionably the best route to success in this context is through cooperation, not selfishness. Counterintuitive as this may seem, game theory has demonstrated unequivocally that best outcomes for human beings engaged in competitive decision making under uncertainty are achieved through a reciprocal altruism, trust and fairness, not deception, lying and cheating, whilst neuroscience has revealed the evolutionary mechanisms. In a sense, the fair price in an efficient market represents the pinnacle of compromise and agreement between consensually competing players.

In 1980, Robert Axelrod, Professor of Political Science and Public Policy at the University of Michigan, set out to answer a simple question: when should a person cooperate, and when should a person be selfish, in an ongoing interaction with another person? For several decades, game theory had been modelling the interplay of conflict and cooperation between rational agents when faced with making decisions concerning uncertain outcomes, and where the behaviour you adopt can influence the behaviour of your opponent and *vice versa*. The simplest way to represent this type of situation is to use a game called the Prisoner's Dilemma. In such a game,

there are two players and each has two choices, either to cooperate or defect (cheat). The traditional single-play game reveals why two purely rational individuals might not cooperate, even if it appears that it is in their best interests to do so. The game is constructed as follows. Two people, Smith and Jones, suspected of committing a crime together have been arrested and imprisoned, with no means of communicating with each other. Each cares far more about their personal freedom than the welfare of their accomplice. The prosecutors do not have enough evidence to convict the pair on the principal charge, only a lesser charge, so decide to entice each co-accused with a plea bargain. If both remain silent (cooperate), both will serve a lesser sentence of 1 year, the reward for mutual cooperation. Alternatively, if both confess and betray the other (defect) they will each receive 2 years (the punishment for mutual defection). If, however, Smith confesses (defects) whilst Jones remains silent (cooperates) Smith will go free (the temptation payoff) whilst Jones will receive 3 years (the sucker's payoff), and *vice versa*. The dilemma is that, by pursuing rational self-interest through betraying each other, their mutual defection ensures they do worse than if they had both cooperated with each other (by keeping silent). Yet, by switching strategy, each player runs the risk of doing worse still, if the other player opts not to switch. The stability that mutual defection created in such a strategy game is known as a Nash Equilibrium (after the Nobel Prize-winning economist John Nash). It is the only outcome from which each player could only do worse by unilaterally changing strategy.

Life, generally speaking, is not like a single Prisoner's Dilemma game. In reality, humans display a systematic bias towards more cooperative behaviour than is predicted by simple models of rational self-interest. A significant reason for this is that, in everyday situations, people tend to interact more than once. Axelrod, therefore, decided to hold a tournament, inviting game theorists to test various strategies in a game of **iterative** Prisoner's Dilemma, where competitors would play each other many times in succession. Crucially, they were given the opportunity to remember the previous actions of their opponents and change their strategies accordingly. The results he found and the conclusions he drew from them were truly ground-breaking and became the focus of his book *The Evolution of Cooperation*. Most significantly Axelrod showed that cooperative behaviour based solely on reciprocity (I'll scratch your back if you'll

scratch mine) can naturally emerge in an environment of pure self-interested players without any central authority. (Readers may note the similarity here with the wisdom of crowds, which also emerges naturally without top-down interference.) Not cheating yields the best outcomes.

Much like the generalised form of the Prisoner's Dilemma, Axelrod's game had 4 ways of scoring:

- The temptation to defect, $T = 5$ points
- The reward for mutual cooperation, $R = 3$ points
- The punishment for mutual defection, $P = 1$ point
- The sucker's payoff, $S = 0$ points

Many different types of strategy were submitted ranging in complexity from the very simple to the algorithmically very complex¹⁴⁸. One was always to defect, a strategy advocated by standard single-play Prisoner's Dilemma. It can never be taken advantage of (so will never receive the sucker's payoff), but misses the chance to gain a succession of rewards for mutual cooperation, particularly when faced with opponents who will change their strategy to cope with its persistent defection. Another was always to cooperate. This does well with other cooperators (always scoring 3 points with iteration) but performs very badly against defectors. Yet another was a purely random strategy, with cooperation and defection chosen by the toss of a coin. Most of the others were a mix of defection and cooperation according to a complex set of rules. Axelrod found that, over a large but unknown¹⁴⁹ number of iterations, selfish strategies tended to do very poorly in the long run while more cooperative strategies did better, as judged purely by self-interest. The winner was one of the simplest containing just four lines of basic computer code: 'tit for tat'. The strategy cooperates on the first move, and then does whatever its opponent has done on the previous move. Thus, when matched against an 'always defect' strategy, it will always defect after the first move. When matched against an 'always cooperate' strategy, it will always cooperate. This strategy has the benefit of both cooperating with a friendly opponent, getting the full benefits of cooperation, and of defecting when matched against an opponent who defects. When matched against itself, it will always cooperate.

The success of the ‘tit for tat’ strategy appears to hinge on a number of key characteristics. It avoids unnecessary conflict but can be provoked into reaction, and will be tolerant after provocation but is not exploitable. Above all, it demonstrates clarity of behaviour. From these Axelrod formulated several conditions necessary for a strategy to be successful: don’t be first to defect; reciprocate both cooperation and defection; don’t be too clever; and don’t be envious of opponents. On this last point, whilst cooperators can never score more than defectors on a single iteration (3 points versus 5 points), resisting the temptation to cheat will always prove more successful in the long run, provided there are other like-minded players. Essentially, the ‘tit for tat’ strategy changes the mindset. Rather than looking to beat others, it prefers to elicit similarly cooperative behaviour. Instead of trying to be the best, cooperative players try to be the best they can. ‘Tit for tat’ means being firm but fair, forgiving yet retaliatory. It is trust with verification, but punish if deceived.

Axelrod went on to explore the social and evolutionary implications of his tournaments. In particular, he was able to show that a potentially cooperative strategy can gain an initial foothold in an environment which is predominantly non-cooperative, provided there is a sufficient number of invaders willing to try. The size of cluster of cooperative players needed to achieve this is surprisingly small. For games of very long duration (over 200 iterations), Axelrod demonstrated that as few as one in a thousand players following a cooperative strategy was enough to invade a world dominated by defection. Even for games with as few as two iterations, anything over 20% of the interactions taking place between like-minded cooperators was enough for cooperation to emerge. Once established, cooperation quickly takes over as the dominant and most stable strategy and the one most resistant to invasion by others, including defectors.

Axelrod used his tournaments to show a possible mechanism for the evolution of altruistic behaviour from behaviours that are initially purely selfish, by natural selection. Essentially, ‘tit for tat’ forms the basis of reciprocal altruism and the universal moral maxim of the Golden Rule: treat others as one would like others to treat oneself. Earlier in the book, I explored how reciprocity may have evolved in human societies via the risk-reduction hypothesis of food sharing. If Axelrod’s theory is correct, agreeing to share the spoils with a view to spreading the risks associated

with acquiring them would appear to have been an inevitable outcome in the history of human evolution. Socially beneficial outcomes do not happen because of some central authority with foresight and understanding determining them. Rather, they arise simply out of the mutual cooperation of self-interested parties looking to achieve the best possible outcomes in an environment where most others are striving for the same. This is the invisible hand of Adam Smith. The best outcomes are the fairest outcomes. Excessive greed and cheating might pay in the short term, but in the long run they are self-defeating. As Axelrod revealed, where success is achieved by exploiting losers, it will inevitably die out with them.

Manifestly, Axelrod's work is really another interpretation of the wisdom of crowds and by extension the efficient market hypothesis. In a market, cooperation, or compromise, will inevitably lead to the best (fairest) prices. Any information asymmetry, benefiting one side at the expense of another, will ultimately be self-consuming. Indirectly, we have provided another explanation for why fair markets are so difficult to beat. Equally, it exposes the weakness of the moral criticism of gambling. Critics of gambling argue that nothing is created; instead there is merely a self-interested redistribution of wealth. Seen in the light of Axelrod's work, however, the trading of opinions about a mutually agreed price is fundamentally an expression of cooperation. Where two people cooperate and agree on the terms of engagement, what is financially a zero-sum game (with a winner and loser) is transformed into something positive, where both sides win when measured in terms of other qualities like excitement, hope, reward anticipation (and, of course, dopamine). The fact that two people can even agree to cooperate in a manner like this is, to my mind, the ultimate expression that they are willing to reciprocate. Gambling, we might say, is the inevitable behavioural outcome of a species whose individuals, excited by thinking about the future, have evolved to treat each other fairly. The existence of lemons is not evidence to support the idea that the business of gambling (and by extension markets more generally) is wrong, merely that a few of its players are defectors. That there are so few is simply a consequence of the fact that the most stable markets exist where most players are not so. As Michael Shermer says in *The Science of Good and Evil*, “*most people most of the time and in most circumstances... do the right thing for themselves and for others.*”

Whilst cooperation will establish itself naturally, its evolution can be speeded up if players are permitted to communicate. In the single-play Prisoner's Dilemma game, the rational choice is to defect. If, however, you permit the prisoners to communicate with each other beforehand, the best strategy changes to one of cooperation. Communication is not a prerequisite to cooperation, but it helps. Trees in the Amazon rain forest aren't very good at communicating. Hence, they find themselves in an evolutionary arms race to get to the top of the forest canopy in search of sunlight. Essentially, they have to grow taller and taller simply to maintain the same advantage. This is sometimes described as the 'Red Queen Effect', after the character in Lewis Carroll's *Through the Looking Glass*. The red queen finds herself having to run as fast as she can simply to stay in the same place. If only trees could talk. People, of course, **can** talk, and if given the opportunity to do so often make the most of it. The political economist and Nobel Prize winner Elinor Ostrom demonstrated that over-exploitation in a market environment could be avoided where investors were allowed to communicate¹⁵⁰. Her 8 participants were each given a set of tokens and anonymously asked to invest them in one of two markets. The first offered a fixed rate of return whilst the second offered a superior variable rate provided not too many people invested in it. As such, the experiment was a classic game theory example of self-interest versus cooperation. So what happened? When no communication was permitted, self-interest ruled the roost and the players returned just 21% of the maximum possible yield from their investment strategy. Over-exploitation of the second market led to a 'tragedy of the commons'. By contrast, when a single mid-play communication was permitted, the yield increased to 55% of the optimal return, and when repeated communication took place, it jumped to 73%.

The absence of communication is clearly one of the causal factors for the development of lemons markets. The information asymmetry on which they are based arises from sellers preferring not to provide buyers with everything they need to know, and buyers failing to perform sufficient due diligence to find out. It is for this reason that I spent 14 years insisting that transparency for a seller of betting advice is the most important quality he can demonstrate. Being transparent means communicating with your market, making available all information that a buyer should know to help

them make a decision about whether the price you are selling at is a fair one. Tell them the truth and they will reciprocate in kind. Sadly, often the incentive to do so is lacking. The reason: because what Axelrod calls the ‘shadow of the future’ is short.

The success of cooperation in Axelrod’s iterative Prisoner’s Dilemma game hinges on the promise of continual interaction. Players who repeatedly interact with each other learn that trying to take advantage of each other leads to punishment and mutual defection. The longer the shadow of the future, the greater the chance is that cooperation will become established. The internet, however, is a perfect environment for short-lived anonymity, ensuring there is little risk of punishment once the lemons have been sold. Internet tipsters frequently hide their identity when registering their domain. Even when they can be traced, there is little prospect of retaliation, given the global nature of the industry. Typically, the fraudulent seller of betting advice resides in a jurisdiction entirely separate from that of a sucker who’s fallen for the scam. The domain and hosting companies, moreover, are frequently found to be located in a third. Try bringing a law suit against such individuals. Of course, the best form of punishment in such cases is simply to ostracise the defectors and to refuse to play the game with them. To that end, the evolution of cooperation has also ensured that social organisms have naturally acquired excellent cheating detection mechanisms. Remember those meal-sharing vampire bats? During grooming, they are highly adept at spotting which ones have distended bellies after a good meal, and by implication which ones have been unwilling to share. Once detected, defectors are more likely to be punished by ostracism, and in future may go hungry. Humans, too, are very adept at individual recognition and score keeping, essential qualities for detecting cheaters, although large anonymous groups like the internet can make this detection task trickier. Above all, however, one must be sure to use them, and not let our own emotions of greed and denial act as blinkers. Where co-operators seek out cooperators, it’s much harder for cheaters to survive. What’s amazing about the evolution of cooperation is that it even has its own neurochemical signature: oxytocin.

In his book *The Moral Molecule: the new science of what makes us good or evil*, neuroeconomist Paul Zak tells the story of how, as a kid working in a gas station, he was the victim of the ‘Pigeon drop’ con. The mark, sucker

or ‘pigeon’ is persuaded to give up a sum of money in order to secure the rights to a larger sum of money, whilst in reality the scammers make off with his money and the ‘pigeon’ is left with nothing. In this case, someone had found some pearls in the rest room and was handing them in to Zak when the item’s owner telephoned asking whether he might have lost them there. With the news that they had just been found, he said he was on his way back and would give the chap who found it a \$200 reward. Unfortunately, the man who had found them couldn’t wait; he had a job interview to attend and wouldn’t be coming back this way. If Zak could do the honours and return the pearls himself, he’d split the reward. Eager to please, and with \$ signs flashing before his eyes, Zak took \$100 from the cash register and gave it to the man. Needless to say the owner of the pearls, which of course were a cheap imitation, never showed up.

In a sense Jack, too, was a victim of Goldleaf’s pigeon drop. Whilst technically he had purchased some land, James and Philip both new it was worthless and their supposed attempts to help him secure a sale were in reality just part of the scam of encouraging Jack to part with more of his money. Zak says that when he looks back on the incident he believes that it wasn’t greed or any of the other deadly sins that motivated him, but rather trust. Everything the conman said and did appeared to suggest he was placing a large amount of faith in Zak. To return the favour, Zak held a genuine desire to be of assistance. Can we say the same of Jack? Was he returning trust favours in kind? ‘Trust’ was certainly a word he used in my conversations with him. What a pity he was sooner prepared to trust James and Philip than the police who were trying to help him. Cynics, of course, will point out that Zak (and presumably Jack too) was evidently motivated by money as well. If trust was all that Zak cared about, why didn’t he give the man the full \$200? After all, it was he who found the ‘pearls’. Nevertheless, whatever the truth here, the episode obviously stayed with Zak as he embarked on his investigations into why it was that most people behave more generously than traditional economic models predict that they should, even with strangers, and why a few of them still try to cheat.

In 2001, Zak asked some of his students to play the ‘Trust Game’. Like the Prisoner’s Dilemma, it is another classic game theory tool designed to investigate rational self-interest and behavioural departures from it. In the trust game, player A receives an endowment, usually in the form of money,

from the experimenter. They must then choose whether to donate part of that to player B. If they do so the experimenter will increase that amount by a predefined multiple. Player B, in turn, is then asked whether they would like to return part of the increased endowment she has received from player A. The players can communicate, but are otherwise perfect strangers. According to conventional notions of rational behaviour, the game should break down before it has begun. Person B, acting selfishly, has no reason to give any money back and, knowing this, person A shouldn't bother sending any over in the first place. Yet in almost all cases, both players send money to the other. Zak's results were no different. Awarding everyone \$10 simply for showing up and trebling the size of the donation offered by A-players, A- and B-players on average walked away with \$14 and \$17 respectively. The really interesting part, however, was what was going on neurologically. Zak found that levels of a mammalian hormone called oxytocin, more typically associated with childbirth and lactation, strongly correlated with a player's willingness to respond to a sign of trust by giving back real money. In fact, oxytocin levels of B-players who knew they had received money from A-players were 50% higher than when they knew that they had simply been given an endowment randomly, and consequently gave almost twice as much back in return. Moreover, the more money sent by A-players, the greater the oxytocin levels in B-players and subsequently the greater the amount of money given back to A-players. Zak coined oxytocin the 'moral molecule', which became the title of his book.

Evidently, what counts is being trusted. Trust in one person triggers oxytocin in the other, which triggers more trustworthy behaviour, and so on in a virtuous circle that Zak calls the 'Human Oxytocin Mediated Empathy' or HOME circuit. Oxytocin also releases dopamine, which reinforces the feeling of gratification when we treat others well, and serotonin, which gives us a mood lift by reducing anxiety. In doing so, it generates empathy that drives moral behaviour, which inspires trust, which causes the release of more oxytocin, which creates more empathy. Empathy is really the physiologically evolved version of the Golden Rule, the emotional response arising from the associative machinery of memory, laid down by the hormone oxytocin, in much the same way that memories of pleasure trigger the release of dopamine and the emotions of reward anticipation. As Zak says, when we are moved to treat others as we would wish to be treated

ourselves, it is in part because we are “*literally experiencing another person’s pleasure or pain as if it were our own.*” Well, at least most of us do. Human beings are particularly adept at empathising because, it is believed, they possess what are called mirror neurons, neurons that fire in the same way when the person both acts and observes the same action performed by another. They are implicated in what is called the ‘Theory of Mind’, the capacity to understand and infer the intentions, beliefs and desires of others; the word ‘empathy’ literally means ‘feeling inside’. Really all of this is a rather convoluted way of saying that we cooperate because it feels good.

Zak also used oxytocin inhalers to investigate whether artificially increasing concentrations would lead to greater expressions of trust. The results were as unequivocal as in the original trust game: oxytocin-loaded participants displayed much greater levels of trust and generosity than those who used inhalers filled with a placebo. In fact, half of the A-players given oxytocin became so trusting that they donated all of their endowment to their corresponding B-players, more than double the number for those on the placebo. In a follow-up experiment, Zak investigated the effects of oxytocin infusion on the behaviour of players in an ‘Ultimatum Game’. In this game, player A (the proposer) receives a sum of money and proposes how to divide the sum between himself and player B (the responder), who can either accept or reject this proposal. In the first instance, the money is split according to the proposal; in the second neither player receives anything. The standard model of rational self-interest dictates that player B should accept even 1% of the money or less, since that is more than she would have received had she not played the game at all. As with the trust game, however, human beings do not conform to the standard model. Typically, offers of even 30% of the money are usually rejected. Evidently, when faced with such choices, the utility responders derive is based not just on money but also the pleasure of punishment and enforcing a principle of fairness. Similarly, proposers derive utility in maintaining a reputation of virtuousness. Such non-monetary utilities are augmented by oxytocin. Zak found that infusions caused generosity to surge by 80%.

If our evolved HOME circuitry predisposes us for trust rather than full blown rational self-interest, why isn’t everyone virtuous all the time and a few people barely virtuous at all? Zak invokes another hormone; we were

introduced to it earlier: testosterone. Like cortisol, testosterone is also a stress hormone. At times of stress, levels of testosterone and its bioactive metabolite dihydrotestosterone increase, helping to block the binding of oxytocin to its receptor in the brain, thus putting the break on the virtuous circle and moving us from a state of empathy to what Zak call the ‘survival mode’, where self-interest becomes the dominant emotion. Men, of course, have far more testosterone than women. It should come as no surprise, therefore, to learn that in all of Zak’s experiments, women consistently released more oxytocin, and were considerably more generous, trusting and empathetic. Really, this confirms what all women know anyway, but we now have a neurochemical explanation. Not only are men more likely to take risks and gamble, they are also more likely to cheat. Indeed, every character I’ve talked about in this chapter of suckers and sharks has been male. Testosterone, however, through the enhancement of more selfish behaviour, does help to increase the punishment of freeloaders since cheating presents a risk within cooperative groups. Consequently, not only are men more adept at detecting those who cheat, they also derive more pleasure in punishing them.

Zak sees testosterone and oxytocin as a neurological yin and yang partnership, balancing aggression and punishment on the one hand with empathy and cooperation on the other. *“Oxytocin maintains the balance between self and other, trust and distrust, approach and withdrawal.”* It is far too simplistic to say that some of us are good and some of us are evil. As Michael Shermer says (in *Mind of the Market*), we exhibit a dual dispositional nature, trusting and cooperative on the one hand and distrusting, competitive and selfish on the other. Aleksandr Solzhenitsyn sums it up perfectly: *“the line dividing good and evil cuts through the heart of every human being.”* During social interactions we are wired to maintain a balance between trust and scepticism, the nature of which will be dictated by our environment and people in it. Too much self-interest and defection, and we all end up as losers. In that respect, it is surely no coincidence that the countries with the smallest economies in the world also happen to be some of the most corrupt. On the other hand, too much ‘free love’ and some players will start to freeload. One of the benefits of having some testosterone-fuelled punishers in a group is that it helps reinforce cooperation by increasing the cost of cheating. It also turns out that having

a few ‘bastards’ as Zak calls them, those people who are unconditionally non-reciprocators, provides the most stable mix of players. Zak thinks that their existence is no accident but in fact is evolutionarily adaptive. *“Bastards are necessary from an evolutionary standpoint because they keep the physiologic balance between appropriate levels of trust and distrust optimally tuned.”* In other words, ‘bastards’ keep the rest of us, especially the ‘punishers’, on our toes.

All this takes us back to the original question: why are some ‘bastards’ seemingly happy to cheat? Testosterone, and its role in enhancing self-interest, aggressiveness and excessive greed, provided an explanation. Now we can see why. By interfering with and impairing the human oxytocin mediated empathy circuit, testosterone can tip the balance of trust and cooperation in favour of distrust and selfishness. In all of Zak’s experiments with trust there was always about 5% who never gave anything back no matter how much money the other player trusted them with. Counterintuitively, these ‘bastards’ actually had elevated levels of oxytocin. Remember, however, as for dopamine and serotonin, our behaviour does not respond to absolute levels of hormones, but rather relative changes to them. These unconditional non-reciprocators saw no surges in their oxytocin levels; their oxytocin receptors were malfunctioning. Zak calls this condition Oxytocin Deficit Disorder (or ODD): oxytocin is simply not activated when and how it should be. It turns out there are three categories of influence that weaken the oxytocin response: temporary, acquired and genetic.

Temporary ODD may be caused by stress, both acute (adrenaline) and chronic (cortisol). Adrenaline is implicated in situations of fight or flight, where self-preservation becomes the dominant response and empathy towards others inevitably takes a back seat. Cortisol, by contrast, modulates our response to stresses that don’t go away so quickly. Over the longer term, where they become inescapable, elevated levels of cortisol not only lead to more addictive behaviour, repeating goal-directed actions which remain unfulfilled (gambling being one of them), but they make us more selfish and uncooperative as well. To be sure, the two are not mutually exclusive, with greed the common denominator. Indeed, the meaning of the word includes both concepts of excessiveness and selfishness.

Acquired ODD is concerned with more deeply entrenched empathy

impairment. Typically, it results from early emotional trauma, for example abuse and neglect in childhood. Oxytocin receptors not stimulated by love and trust at a time when a brain's neurons are most malleable ensures they fail to develop properly. Zak found that variance in trust game generosity is much larger in players who have been abused early in life. Conceivably, Freud's suggestion that an addictive and deliberately destructive self-interest in gambling acts as a displacement for a childhood resentment of his neglectful parents might have a neurological underpinning in oxytocin deficit.

Most interesting of all, at least in an evolutionary context, are the genetic explanations for ODD. One of these is autism. Autistic individuals have been found to be far less generous in the ultimatum game, with nearly a third offering nothing compared to only 3% in control groups. Some studies have shown that autistics have lower levels of oxytocin, whilst high levels of foetal testosterone have been implicated in impairing the HOME circuit. Indeed, autism has even been called the 'extreme male brain syndrome'. Undoubtedly, boys outnumber girls by 4:1 on the autism spectrum. Classically, autism is characterised by difficulty in communicating, socialising and empathising, with a Theory of Mind lacking that enables people to 'get inside' the emotions of others. Unsurprisingly, autistics are also more likely to accept low offers in the ultimatum game because, as Zak argues, they miss the "*subtleties of give-and-take, which is the essence of productive cooperation.*" Presumably, give-and-take forms the essence of honest trading and gambling, too. You may recall that I believe Jack, the victim of Goldleaf's lemon trading, to be a sufferer of Asperger's syndrome. In unconditionally trusting the sales tactics of James and Philip, was he accepting 'low offers' that most people would quite rightly have refused?

At the extreme end of ODD we find the psychopaths or sociopaths, accounting for roughly 1% of the general population, who don't care about anything except themselves. Whilst they are devoid of empathy they are skilfully adept at pretending otherwise, and are often highly intelligent characters with a contrived social charm that they use to mimic trustworthiness. You won't be surprised to hear that psychopathy is associated with excess testosterone, and that consequently it is three to four times as prevalent in men as in women. It's also about four times as

prevalent in those with senior positions in business and the corporate world. As much as a quarter of all crime and half of serious crime is believed to be attributed to psychopaths. Little research has been done investigating any possible link between psychopathy and fraud since most has focused on violent crime, although carriers of the ‘warrior gene’, the gene variant linked to psychopathy, have been shown to be prone to taking greater financial risks. Nevertheless, many of the characteristics typically associated with psychopathy, such as recklessness, charm, deceitfulness and a lack of guilt are consistent with traits exhibited by those drawn to commit fraud. Jordan Belfort, the convicted penny stock fraudster whose life story was featured in the biopic *The Wolf of Wall Street*, was unquestionably psychopathic. I have little doubt, too, that James Browne was similarly inclined. On several occasions, Jack would insist what a charming young man he was. Choosing to abscond prior to his court hearing would also hardly seem surprising for such a character.

When George Akerlof wrote his paper about a market for lemons, he concluded that the problem of suckers and sharks did not necessarily imply a need for top-down regulatory solutions. In that he was clearly following in the footsteps of Adam Smith. Trust and cooperation, as we’ve learnt, are built from the bottom-up, not by rules and directives imposed by centralised authority but by individuals themselves collectively arriving at the best decisions, in the same way that a wisdom of the crowd delivers the best prices in a market. Lemons markets might well be built out of excessive greed, expressed by both the sellers and buyers, but you can’t force people not to be greedy. The onus, surely, must be on the players themselves to discover that markets in general, and gambling ones more specifically, are places for agreement and consensus which work best when their players play fairly. How to deal with the sharks, then? Essentially, the solutions Akerlof proposed were of a self-regulatory nature, including warranties and guarantees which help to move the risk from the buyer to the seller. Really this is just equivalent to the dictum: trust, but verify. Evolution has given human beings excellent tools to trust and verify. It would seem a pity not to use them effectively.

For some, of course, the whole business of gambling and more generally market economics has been regarded with scepticism. For them, the *de facto* reality of markets means a presence of lemons and exploitation for the

purposes of self-interest by default. Adam Smith's invisible hand is mistrusted because its randomness lacks a causal explanation, whilst wealth concentration is regarded with disdain. Given our evolutionary roots perhaps this is not surprising. The modern winner-takes-all inequality stresses our sense of fairness that evolved to cope with a zero-sum world via reciprocity, cooperation and trust, whilst our pattern recognition engine imparted a low tolerance for uncertainty and unpredictability. For top-down engineers, these difficulties can only be solved by recognising that the corruptibility of man requires a culture of altruism, either secular or religious, to tame him. Yet a self-regulated balancing act between cooperation and competition, trust and scepticism, altruism and selfishness provides the necessary counterargument to the idea that pro-social behaviour can only be imposed via top-down authority. The very criticisms that insist that gambling is wrong are the same ones that would have us believe that all self-interest is bad, and by extension, that markets in general are too. This examination has revealed how completely flawed this thinking is. Markets, like trust, cooperation and, more generally, morality emerge, bottom-up, out of the interaction of players looking to achieve the best for themselves. Evolution has seen to it that for social creatures like human beings this means the best for others as well. Axelrod's experiments showed that cooperation will develop purely out of reciprocal self-interest. Zak's work on oxytocin and trust suggests that this magical by-product has assumed an evolutionary signature all of its own. In the same way that players in markets arrive, as if by magic, at the fairest prices, evolution has arrived at the solution of cooperation as the best possible strategy to assist selfish genes carried by social animals to achieve the goal of survival and reproduction.

Michael Shermer sees parallels between evolution and the market, noting that Sir Charles Darwin is intellectually a direct descendant of Adam Smith, the forefather of modern economics, capitalism and free trade. Much has been made of the incompatibility of his ideas expressed in the *Theory of Moral Sentiments* and *The Wealth of Nations*, specifically how to reconcile the emphasis on benevolence in the former with the emphasis on self-interest in the latter. In fact, there is no incompatibility at all. Smith's concepts of virtuous (or reciprocal) self-interest and mutual sympathy (or empathy) are more or less one and the same which can now be explained by

evolutionary biology. Shermer makes the connection between markets and cooperation explicit.

“Trust and cooperation lead to a viable free market of exchange, and free markets lead to greater trust and cooperation – the very model of a complex adaptive system that learns as it develops.”

In my opinion, honest speculative gambling, free of lemons, represents the most quintessential example of a market built on reciprocated self-interest and trust, where opinions about risk and uncertainty are shared and exchanged. Like trade, where deals are struck between participants rather than matches played between them, gambling, too, involves a deal, a consensual agreement entered into by two or more parties for their mutual benefit. The redistribution of financial rewards might be zero-sum but, as I've previously tried to show, a positive-sum trade in hope, anticipation and psychological control that paradoxically appears so prevalent in games involving uncertainty arguably matters just as much. Indeed, such motivation by uncertainty and risk taking credibly has evolutionary origins. After all, if we did not crave to be right about the future so much, we wouldn't be bothering to gamble at all.

[125](#) Akerlof, G. A., 1970. The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, **84**(3), pp.488-500.

[126](#) The term boiler room refers to the selling of questionable investments by telephone, via unfair and dishonest sales tactics. The term is likely to have originated from the cheap, hastily arranged office space used by practitioners, often just a few desks in the basement or utility room of an existing office building.

[127](#) Armstrong, J.S., 1980. The seer-sucker theory: The value of experts in forecasting. *Technology Review*, **83**, pp.16-24.

[128](#) A list of names that have previously been successfully solicited and which is shared amongst firms happy to engage in fraudulent practices.

[129](#) Roth, G., 2011. *Lost and Found: One Woman’s Story of Losing Her Money and Finding Her Life*. New York: Viking.

[130](#) <http://forum.bettingadvice.com/showthread.php?t=83482> <http://forum.bettingadvice.com/showthread.php?t=85542> <http://forum.bettingadvice.com/showthread.php?t=88126>

[131](#) http://www.verifiedtipsters.com/tipsters_present.php?tipster_username=BettingAdvices
http://www.verifiedtipsters.com/tipsters_present.php?tipster_username=protips

[132 http://www.mybigpartner.com/user/Besttips4ever](http://www.mybigpartner.com/user/Besttips4ever)

[133](#) Ivan, a former co-owner of Tipstersplace.com, was evidently well versed in the marketing of lemons. When engaged in negotiations for a possible sale of his website in 2014, he encouraged the prospective buyer to add fake tipsters to his portfolio as a means of minimising the costs of paying real ones to work for him. Presumably that also makes overall performance look better. In the end it would appear his former partner, Konstantin, bought Ivan's share. Whether Konstantin has followed Ivan's advice is naturally impossible to tell. He says he's never used fake records.

[134 http://www.verifiedtipsters.com/forum/index.php?board=33.0](#)

[135 http://www.verifiedtipsters.com/forum/index.php?topic=2641.0](#)

[136 http://www.verifiedtipsters.com/forum/index.php?topic=2855.0](#)

[137 http://www.verifiedtipsters.com/about_us.php](#)

[138 http://forum.bettingadvice.com/showthread.php?t=89728](#)

[139 http://www.verifiedtipsters.com/tipsters_present.php?tipster_username=jeffhorse](#)

[140 http://www.boxwind.com/site/en/verifiedtipsters.com](#)

[141](#) My offer of £250 to the person who can find it still stands.

[142](#) Sutton, R., 2010. *Does Tough Competition Breed Better Performance, Or Just More Cheating?*

<https://www.psychologytoday.com/blog/work-matters/201006/winner-take-all-incentive-systems-competition-and-cheating-teachers-soccer>. (24th June).

[143](#) Schwieren, C. & Weichselbaumer, D., 2010. Does Competition Enhance Performance or Cheating? A Laboratory Experiment. *Journal of Economic Psychology*, **31**(3), pp.241-253.

[144](#) Robertson, I., 2013. *The Winner Effect: The Science of Success and How to Use It*. London: Bloomsbury.

[145](#) Coates, J. M. & Herbert, J. Endogenous steroids and financial risk taking on a London trading floor. *Proceedings of the National Academy of Sciences of the United States of America* (2008), **105**(16), pp 6167-6172.

[146](#) Barnard, J. W., 2013. Shirking, Opportunism, Self-Delusion and More: The Agency Problem Lives On. *Wake Forest Law Review*, **48**; William & Mary Law School Research Paper No. 02-265.

[147](#) Lammers, J. & Stapel, D.A., 2009. How power influences moral thinking. *Journal of Personality and Social Psychology*, **97**(2), pp.279-289.

[148](#) In fact, Axelrod held two tournaments, the first with 15 entrants and the second with 62.

[149](#) When the number of game iterations is known, the rational strategy is to defect on the last iteration. Consequently, it pays to defect on the penultimate iteration as well and hence the one before that, too, and so on. Cooperation only works where the number of iterations is either unknown or infinite.

[150](#) Ostrom, E., Walker, J. & Gardner, R., 1992. Covenants with and without a Sword: Self-Governance is Possible. *American Political Science Review*, **86**(2), pp 404-417.

THE FOX AND THE HEDGEHOG

In 1953, the philosopher Isaiah Berlin wrote a popular and rather light-hearted composition entitled *The Hedgehog and the Fox*, in which he discussed Russian novelist Leo Tolstoy's interpretation of history. The title is a reference to a passage attributed to the ancient Greek poet Archilochus: "*the fox knows many things, but the hedgehog knows one big thing.*" In his analysis Berlin divided scholars, authors and philosophers into two categories: hedgehogs, who view the world through the lens of a single defining (big) idea (for example, Plato, Marx and Nietzsche), and foxes, who draw on a wide variety of experiences and for whom the world is represented by a plethora of competing (little) ideas (for example, Aristotle and Shakespeare). Isaiah Berlin contended that, whilst Tolstoy aspired to be a hedgehog with singular conviction of thought, he was by nature a fox unable to reject the view that history, as propounded in his book *War and Peace*, is shaped by forces and events that are numerous and fundamentally unknowable.

At this point you may be wondering what on earth all this has to do with gambling. History, after all, is the study of looking back; gambling on the other hand is concerned with predicting the future. The point, however, is this: how one views the evolution of history is inextricably connected to how one understands the business of forecasting the way it is made. Unlike Marx, who believed in deterministic laws that shaped events, Tolstoy abandoned this romantic view of history and instead compared it with calculus: the sum of an infinite amount of small events, feelings and so on whose chaotic mess ensures they cannot be predicted. The Yale Professor of History, John Lewis Gaddis, similarly imagines the process that makes history as a fusion of many alternative realities passing through a funnel to make the present. Such imagery bears obvious similarities to the quantum mechanical interpretation of reality I discussed earlier in the book, where a potentially infinite number of alternative worlds is possible, described by the (probabilistic) wave function, but where only one is observed, the point

at which the wave function collapses. Essentially, these contrasting views of how history is made, and by extension how the future is predicted, are defined by their treatment of causality and uncertainty. It turns out that those who think more probabilistically (the foxes) are generally better at prediction than those who think more deterministically (the hedgehogs). Since good prediction lies at the heart of good gambling, this final chapter will examine some of the characteristics that define it.

Lessons from Political Forecasting

Philip Tetlock, a psychologist and political scientist, spent 20 years from 1983 recording the predictions of 284 experts, including government officials, professors, journalists in the fields of politics, economics and international affairs, covering issues as diverse as US presidential elections, the break-up of the former USSR and independence for Quebec. He published his findings in his book *Expert Political Judgment: How Good Is It? How Can We Know?* From around 82,000 predictions drawn from over 27,000 forecasting questions about the future, he concluded that forecasters were only slightly more accurate than chance, and usually worse than basic extrapolation algorithms, especially over longer prediction periods measured in years. They were overconfident and poor at judging objective probabilities (as defined by retrospective analysis of whether predicted outcomes actually happened). About a quarter of the time, outcomes that forecasters subjectively rated as sure or almost sure things did not occur, whilst about 15% of events considered impossible actually transpired. Such subjectively biased judgement is decidedly reminiscent of Daniel Kahneman's possibility and certainty effects. Clearly, political forecasting is little different to sports or financial equivalents: complex, non-linear and mostly random. Should anyone need reminding, in high luck environments where the links between causes and effects are tenuous, it's easy to confuse experience with expertise.

Despite such underwhelming aggregate performance, Tetlock nevertheless was able to distinguish characteristics that identified someone as being better suited to making more accurate predictions. Intriguingly, it was not what the experts thought that influenced how accurate their

forecasts were – that is to say, their political views – but how they thought it. For Tetlock, differences in thought process bore a striking resemblance to Isaiah Berlin’s hedgehog–fox dichotomy.

“Low scorers look like hedgehogs: thinkers who ‘know one big thing’... who... express considerable confidence that they are already pretty proficient forecasters.... High scorers look like foxes: thinkers who know many small things, are sceptical of grand schemes, see explanation and prediction not as deductive exercises but rather as exercises in flexible ‘ad hocery’ that require stitching together diverse sources of information, and are rather diffident about their own forecasting prowess.”

Furthermore, Tetlock observed differences in the way experts sought to defend and rationalise their mistakes. Whilst foxes tended to express a greater degree of humility about their expertise and an acceptance of errors, hedgehogs were more likely to attribute bad luck to mistakes and were more intolerant of suggestions that they had got things wrong. Their belief system defences tended to fall into three broad categories: it still will happen (or off on timing), it nearly happened (close call) and it would have happened (but for the exogenous shock). Such self-rationalisations are classic examples of hindsight bias which encourages an irrational faith in the inevitability of history and one that can be causally explained retrospectively. The psychology of ‘I was almost right’ can unwittingly encourage an overconfidence that is completely unwarranted. In truth, the forecaster was still wrong. Near misses in gambling sadly have a similar effect, most noticeably in the use of slot machines or the betting of accumulators. Landing all but one of the required outcomes encourages many players to believe that they almost won, when the reality is that a loss is a loss. The symbols displayed on a slot machine or the teams selected for a betting accumulator are completely independent of each other and have nothing to do with how close the player was to winning the jackpot. Unfortunately, near misses motivate players to continue gambling, because they feel they are close to winning.

Forecasters with the biggest news media profiles were particularly bad at prediction and most hedgehog-like in their thinking. Tetlock’s work suggests that there is an inverse relationship between fame and accuracy. Such an observation mirrors Nate Silver’s suggestion that the confidence people express in their predictions may be inversely correlated with their

validity. It also seems remarkably similar to Nassim Taleb's inverse rule of contribution: the higher up the corporate ladder, the lower the evidence of his contribution. Presumably, Taleb is applying such a rule to environments of low or zero-validity. We've seen that betting and investing are two such environments; so, too, it would appear, is political forecasting.

Characteristics of Foxes and Hedgehogs

Building on Nate Silver's useful classification of characteristics displayed by foxes and hedgehogs, I've tried to provide a meaningful summary table.

Foxes	Hedgehogs
Multidisciplinary	Specialised
Adaptable	Unyielding
Self-critical	Stubborn
Tolerant of complexity	Order seeking
Cautious	Confident
Uncertain	Certain
Empirical	Ideological
Probabilistic	Deterministic
Focus on process	Focus on outcomes
Inductive (Bayesian)	Deductive
Bottom-up	Top down
Emergent	Reductionist
Gatherers	Hunters
System 2	System 1

Foxes tend to be more multidisciplinary, incorporating many different interpretations with a view to creating a consensus opinion or a 'wisdom-of-ideas' from all of them. Hedgehogs, on the other hand, prefer to specialise in singular forecasting niches and methodologies. Foxes will be adaptable, trying several forecasting approaches in parallel and giving up ones that

aren't working. Hedgehogs, meanwhile, will go 'all-in' with a chosen strategy, using new data to refine the model. Foxes are more willing than hedgehogs to acknowledge mistakes and less likely to suffer from attribution bias, where successes are attributed internally (skill) and failures are attributed externally (luck). Foxes see the world as complex, with patterns emerging out of an interaction of many variables. Hedgehogs, by contrast, reduce outcomes to a few basic governing rules. Foxes draw conclusions with caution, assigning degrees of provisional (or Bayesian) probability to their likelihood based on bottom-up, inductive, data-driven reasoning, updating those probabilities with new empirical observations. Hedgehogs instead exhibit a confidence grounded in a deterministic cause-and-effect, top-down, ideological view of the world via more deductive reasoning that looks for explanatory closure. Consequently, they are less likely to update their opinions, and are more likely to trivialise evidence that undercuts their preconceptions whilst embracing evidence that reinforces them, the confirmation bias. Foxes are slower, more methodical and willing to engage the rational System 2 in formulating judgements. Hedgehogs are more reliant on gut feeling and intuition, allowing automatic System 1 to find heuristic short cuts to conclusions, and are thus more prone to cognitive bias. Foxes prefer to gather lots of little forecasting successes; hedgehogs hunt for and revel in the glory of jackpot accuracy, inattentive to the reality that these are usually just lucky. Foxes think more about the process of forecasting. This pays dividends when the luck is high and breaks the causal link between skill and results. Hedgehogs, unsurprisingly, focus more on their outcomes. The upshot, in low validity environments, is that foxes generally make better forecasters than hedgehogs.

Unfortunately, human beings have evolved to be hedgehogs. In a world with limited resources and countless threats to survival, it doesn't really pay dividends to over analyse. The fable of the Fox and the Cat, which embodies the same ideas as Archilochus' poetry, explains this succinctly. In the basic story a cat and a fox discuss how many tricks and dodges they have. The fox boasts that he has many; the cat confesses to having only one. When hunters arrive with their dogs the cat quickly climbs a tree, but the fox, suffering from over analysis, self-doubt and decision paralysis, is caught by the hounds. Daniel Kahneman's metaphor of the lions describes

much the same thing. For human beings looking to avoid becoming lunch it is of greater benefit to see patterns grounded in causality, even in a low-probability environment or where that causality is illusory. From an evolutionary perspective, mistakes in such environments arising from slow decision making proved far more costly. Having a hedgehog cognitive style proved to be the winning strategy. That is, after all, why we possess it. As much of the earlier part of the book set out to demonstrate, we are hard-wired to find explanations for things, understanding events in terms of a linear chain of cause and effect, with a view to gaining control over our future that increases our chances of success in our goal-directed behaviours. We are not designed to think probabilistically, but deterministically, interpreting the world through narrative and story-telling rather than statistics, through outcomes rather than processes. As a consequence, and as Kahneman has shown, our subjective estimates don't always match up well with objective reality, ensuring we become victims of numerous fallacies. This cognitive style is perfectly adapted to simple linear environments where the link between cause and effect is clear, but breaks down in a world of non-linear complexity more typically found in prediction markets where most of what happens is just unexplainable random noise, and deciphering a signal is close to an impossible task.

For example, a football match will evolve in one of a potentially infinite number of possible ways. The slightest deviations in starting conditions can result in a completely different reality on account of its complexity as a system, and conceivably even as a consequence of quantum fluctuations as well. This uncertainty clearly has implications for wagers we place on football matches, how we go about researching and choosing them, and finally how we interpret their outcomes. Continuing with this example, a hedgehog might be thinking like this:

Liverpool is in better form than Manchester United, so I'll bet them. They win: great! My judgement was sound and it caused me to win my bet.

A fox, on the other hand, might be thinking like this:

Liverpool is arguably in better form than Manchester United, but I understand that there are many factors that can influence a game that might not have occurred to me. They win. I acknowledge that I picked the right team, but accept that the game could easily have turned out differently. Perhaps my judgement was lucky this time. I should repeat this process many times. Good and bad

luck should then even out over the long term and if I still have a profit left over, that might indicate some of my judgement is sound.

Notice the more thoughtful, deliberate and lengthier thought process of the fox. Of course, to many bettors trained in the art of value hunting, all of this might sound intuitively obvious, and yet it never ceases to amaze me how many punters and tipsters I meet indulge themselves in a self-congratulatory pat on the back when they make the right call on a single match.

Becoming foxier is not easy. It takes conscious cognitive effort (via System 2) that our brain prefers to avoid. In complex prediction markets it means acknowledging that errors and failure are inescapable because so much is unpredictable. Inevitably, that also means that subjective confidence is not to be trusted as an indicator of validity. Doing so, however, is an admission of weakness, not something that comes naturally to human beings primed for ego-defence. Of course, foxes will not have it all their own way. When multivariate thinking is used to build an interpretation of reality, the prediction model it builds can begin to fit itself not only to the underlying signal but to the noise as well. More complex statistical modelling might ostensibly reduce the standard error but won't necessarily improve its validity. The data you analyse represent just one reality of a potentially infinite number of possible realities; perhaps the set of data you are looking at is not necessarily a typical representation of the norm. Consequently, over analysing can lead to poor prediction results. Plainly, a balance needs to be struck.

In his concluding remarks, Philip Tetlock regarded decision making as a trade-off between theory-driven (deduction) and imagination-driven (induction) modes of thinking, between a need for closure and complexity, certainty and uncertainty, confidence and humility. For him good predictive judgement is a balancing act between both cognitive styles, between excessive closed and open mindedness. We might call such people fox-hogs. The bottom line, however, is that the best judges are probably those who think about how they think, rather than those who look only at their results. As I've said to sellers and buyers of betting advice alike on more occasions than I care to remember, just because you won that bet does not mean you caused it to happen. If you fail to think about other possible

explanations for why you made a profit, you may find yourself at a loss to explain why you've stop making one when that time comes.

Foxes versus Hedgehogs: an Example from Sports Prediction

Writing at Scoreboard Journalism in August 2013, Simon Gleave, head of analysis at InfostradaSports.com, compared the performance of a number of sports journalists to a number of computer models in predicting the outcome of the 2013/14 English Premier League^{[151](#)}. A year later, Simon revisited those predictions to see how they had performed and published the illuminating results^{[152](#)}. Whilst both sports journalists and computer models were consistently better (on average) than either random guessing or simply copying the finishing positions from the previous season, the model predictions were superior to those of the journalists. The question is why? For the original research, computer models came with points predictions whereas sports journalists were only asked to predict rankings. Strikingly, it was evident that the majority of the models made fairly conservative estimates for the finishing points totals, with the average for all models for top and bottom being 79 and 33 points respectively, compared to 83 and 29 points for the actual average top and bottom points totals across the previous 18 seasons of the Premier League (the 20-club era). Indeed, one model ranged from only 38 to 71 points. Might this offer a clue as to the way computer models arrive at their predictions?

A useful way of analysing how each finishing rank prediction had performed is to compare them to the actual finishing positions of the teams and then calculate the amount of error within each sample of 20 team predictions. Those familiar with statistics will recognise this as simply the standard error or standard deviation. This is calculated by subtracting the predicted rank from the actual finishing rank for each team, and then calculating the standard deviation of those differences across the 20 predictions. If every prediction made was correct, the standard error would be zero. Conversely, if every forecast made merely predicted a mid table finish (10th) for all teams, the standard error would be about 6. So how did the computer models and journalists fare? Both groups did OK when compared to either mid-table guessing or previous season copying.

However, considerable variation existed amongst both groups. The best models achieved standard errors of just 3.5 whereas the poorest ones could only manage about 4.9. This compares to the best and poorest journalists' errors of 3.5 and 5.3. These figures, however, don't tell the whole story. Whilst the best journalist (Joe Prince-Wright of NBC) was about as good as the best model, the average model performance (4.08) was considerably better than that of the average journalist (4.37).

Whilst this superiority could just have happened by chance, conceivably it might also have something to do with the different ways models and journalists are thinking. A clue as to what might be underpinning that difference can be found by comparing the spread of the errors for models and journalists alike, that is to say, the standard error in the standard errors. Journalists exhibited a wider variability of predictive success than the computer models. In fact, the standard deviation of individual standard errors was 0.43 for the models, whilst it was 0.50 for the journalists. Furthermore, much of the variance in the models' errors arose because of the 3 poorest performers; the rest were fairly closely grouped. Indeed, remove those 3 models and the standard deviation across models' errors drops to 0.20. Do the same for the journalists and it's still as high as 0.40. So why the bigger spread in forecasting accuracy for journalists compared to computer models? Perhaps models were thinking more like foxes whilst the journalists were thinking more like hedgehogs.

Computer models, conceivably, are more likely to consider a wide variety of quantitative data and be more inclined to draw conclusions unbiased by singular and dominant points of view. Sports journalists, by contrast, may be more inclined to make bold predictions about what they think will happen, going out on a limb possibly with a view to garnering a reputation. Jackpot success will look spectacularly profound, dramatic failures spectacularly stupid. Of course, in either case, luck probably plays a major influence, but the result is a wider variation in predictive accuracy. Consider, for example, the prediction from David James (ex England goalkeeper and BT Sport forecaster) that Everton would finish the 2013/14 season in 16th place, one place above relegation. Clearly he must have had his reasons for believing this to be likely. It was a very bold prediction indeed; all the other journalists ranged from between 4th and 9th. In the event, it turned out to be wrong (Everton finished 5th), and in large part

accounted for him overall proving to be the poorest forecaster of all, only marginally better than simply assuming every team would finish 10th .

This is not to argue that all computer models are thinking like foxes and all sports journalists are thinking like hedgehogs. Despite the differences highlighted, there are other general commonalities. Models and journalists alike were all pretty successful in predicting that Manchester United would not win the league (indeed not a single journalist predicted it would happen compared to 2 models). Similarly, they all failed to predict Crystal Palace's 11th place finish with the majority having forecast last place and none better than 18th . And despite the variations, models and journalists, on average, didn't disagree by more than 2 ranking places for every team. Overall, however, the reality, as Philip Tetlock found in his 20 years of research into political forecasting, is that both foxes and hedgehogs alike are not very good, on average, at forecasting uncertain futures, better than random perhaps, but probably not enough to justify the large sums of money spent on hiring these journalists to offer opinions about sporting outcomes in the first place, or indeed modellers spending time predicting the evolution of betting markets more generally if those predictions are meant to offer something better than the market with a view to finding profitable trading opportunities. Whilst I haven't investigated whether any of these predictions could have been used to make a betting profit, it is interesting to note that those defined by the Pinnacle Sports betting market were bettered by only 7 of the other 26 model and journalist forecasts, and only significantly so by 3 of them. Surely this is further evidence, if it were needed, of just how wise and efficient a betting market can be at predicting the future and how difficult it will be to beat, once commission for playing in it is paid. In reality, whilst some forecasters will think more like foxes and some more like hedgehogs, few will be capable of being consistently right in the long term.

What makes an Intelligent Gambler?

In the autumn of 2014, I found myself writing on a betting forum thread with the opening title: *The Keys to Success*. I decided I would post some of my thoughts about what makes an intelligent bettor, and what, by contrast,

makes an unintelligent one. Manifestly, they reflect the dichotomy of cognitive styles that I have examined thus far in this chapter. Their applicability can be extended to other forms of speculative gambling where the nature of the game offers a theoretical profitable expectancy: primarily, poker, investing and trading. Most people might argue against the use of the word ‘gambling’ to describe such activities, which should be reserved solely for games where the odds are always, and by definition, mathematically against you (as at the casino). From my examination of the data in the middle part of the book, however, I trust that you will understand why I have done so. Regardless of the availability of theoretical positive expectancy, few players, it would appear, manage to achieve it. To all intents and purposes, the remainder should simply be considered to be gambling, so gambling is what I call it.

So what makes an intelligent gambler? Most importantly he¹⁵³ is someone who thinks probabilistically and appreciates that most of what happens in gambling is luck, with cause and effect in the uncertain environments within which he operates only very loosely connected. He understands regression to the mean, the law of large numbers, and that sequential wagers represent a memoryless Markov chain where the outcome of the next wager has nothing to do with the previous one. Consequently, he is aware of the fallacies of small samples (generalising expected success from a few wagers), hot hands (a belief in winning streaks) and the maturity of chances (the gambler’s fallacy), and avoids committing them. He understands why losing hurts twice as much as winning is enjoyed but has the discipline not to chase losses as a consequence. He understands that players, against whom he is competing in a relative skills market, define the value of assets and outcomes, not so much third party market facilitators¹⁵⁴. He recognises, furthermore, that collectively this crowd is very often wiser and more efficient at finding ‘true’ value than almost all the individuals who form it, so much so that beating it consistently may be a talent possessed by only the smallest percentage of winners. And he realises that, the better people are at prediction, paradoxically the harder it can become to join such a club. He is modest, not overconfident, accepting that wins can be as much a consequence of good luck as losses are of bad. He recognises that the settlement of an individual wager has little or nothing to say about

its underlying value, and rather perversely would prefer to lose with a positive expectancy than to win with a negative one. Finally, he recognises that success in prediction markets takes a lot of hard work whose application suffers from the law of diminishing returns.

By contrast, unintelligent gamblers think deterministically, believing A has caused B, and that luck has little role to play in this type of positive expectancy gambling. Like Laplace believed many years ago, for him chance is merely the expression of man's ignorance. Thus, provided we know all the facts, winning should be easy. He confidently extrapolates success from a small number of wagers to a more generalised perception that he is sufficiently skilled to return a profit indefinitely. Hot streaks, furthermore, provide a sign that winning breeds more winning. Wagers are won because of things he did to foretell that, whilst losing is generally attributed externally to bad luck, outcomes not happening the way they were meant to or even because of fixes in the market. For him, simply winning is proof enough that the method of arriving at it worked; no need to over analyse the role of chance in any of that. It's the outcome that matters and the statistics of value expectation are largely irrelevant. Following Tetlock he "*would rather risk anointing lucky fools over ignoring wise counsel.*" More formally, such gamblers will sooner reject a true null hypothesis (type 1 error or false positive) than fail to reject a false one (type 2 error or false negative), since they don't really care about the explanations for winning, just the winning itself. Losing sadly fails to change these deep-rooted underlying perceptions, since it's psychologically and emotionally easier for him to deny that he lacks any predictive ability.

Conquering the Unknown

It's time now to draw this book to a close. Researching and writing it has involved a personal journey of nearly two years in the making to understand how and why people choose to gamble. When I first embarked upon it, I had no idea it would take me so long and to such far reaches of human knowledge. It quickly became clear, however, that to do the subject justice I would have to adopt a very multidisciplinary approach to my thinking, to in effect think more like a fox. Our brains may have evolved in an

environment that places a premium on thinking like a hedgehog, to ensure that our primary drives of food, safety and sex are met. Yet all human beings, to a greater or lesser extent, possess a capacity to think like a fox as well; that is to be curious, a capacity that Ian Leslie, author of *Curious: The Desire to Know and Why Your Future Depends on it*, calls the ‘Fourth Drive’. Like the other three, the goal is the same, a sense of control that enhances the chances of survival, but it is unique in requiring a sense of self-awareness. In short, the drive for curiosity reduces to three basic questions: who am I? why am I here? where do I find meaning? Searching for the answers involves conquering the unknown, or as the philosopher, and most foxiest of thinkers, Aristotle called it, ‘the desire to know’.

Before you imagine I have finally drifted off into a world of esoteric nonsense, this is not to argue that people are asking themselves such deep philosophical questions when they place a bet, play a hand, spin a wheel, roll the dice or trade a stock. Evidently, such topics are not on their mind; rather, simpler emotions like fun, excitement, escape, hope and anticipation take centre stage. But arguably these feelings are serving higher order subconscious ones, including a pursuit of success, self-esteem and self-actualisation, which in turn reflect a need to search for meaning in life. To be a successful gambler provides confirmation that he is somebody, who has beaten the system, who has won with wits, someone who can predict the future and consequently conquer the unknown, confirmation that he is in control. Joseph Mazur, in *What’s Luck Got to Do with It?*, captures this idea perfectly.

“[G]ambling behaviour is primarily connected to an intrinsic desire to manipulate luck in order to validate life, to test the forces of uncertainty under a fantasy of knowing something unknowable or to experiment with the new.”

In other words: to be curious.

The wonderful contradiction is that our craving for knowing, meaning and certainty is motivated by the very quality that makes gambling largely unknowable, unconquerable and uncontrollable: randomness. Science is teaching us that uncertainty of rewards, rather than the rewards themselves, is what drives risk takers to continue to take risks. For a curious mind, certainty breeds monotony and dullness. Perhaps this is why human beings, fundamentally curious, have been attracted to gambling for a very long

time, possibly as long as our species has existed. Gambling presents itself as a kind of elixir of knowledge, a means of gaining authority over the unknown, granting us an illusory sense of control, but which ultimately disappoints and has us coming back for more: the addictive power of maybe.

Arguably, it may be sports gamblers who perceive the greatest sense of control because, as Professor Pinhas Dannon of Tel Aviv University explains, “[s]ports gamblers seem to believe themselves the cleverest of all gamblers... think[ing] that with experience and knowledge... they can predict the outcome of a game better than the average person.” But his research¹⁵⁵ has determined that neither betting experience nor knowledge of the details of a soccer game is connected to successful betting outcomes, and that sports gamblers, in the main at least, are operating under an illusion of control. The data I presented in the chapter ‘Monkeys Throwing Darts’ would seem to support these conclusions.

So as Joseph Mazur says, a “*belief in luck turns out to be as natural as religion... a desire to control one’s destiny.*” Small wonder, then, that so much criticism for gambling has come from those with a religious predilection. For me, however, there is a fundamental difference. A belief in God is, in large part, acceptance of absolute truth. The corollary is that there is no need for curiosity. Why seek knowledge that God has not seen fit to present us with? It’s hardly surprising, then, as Ian Leslie argues, that so many of our stories about curiosity – the forbidden fruit and Pandora’s Box for example – are warnings. A belief in luck, however, if understood properly, is an acceptance that ‘truth’ is always uncertain, provisional, statistical and therefore falsifiable, and that there is always something new to learn. Daniel Kahneman has it right: “*An unbiased appreciation of uncertainty is a cornerstone of rationality.*”

Religious determinism has no place for chance and randomness, which are properties to be distrusted, feared and suppressed, in much the same way that social and political determinism and ideology distrust the invisible hand of a market. That is understandable when you think about the hedgehog inside all of us, demanding simple causal explanations for the way things are, with an evolved pattern recognition engine to get that job done. The idea that anything, indeed everything, at its most fundamental

level may not need any cause or explanation at all is possibly the greatest dissonance a human mind could suffer. Try as one might, it is hard to accept the notion of chance as causeless, in precisely the same way that no matter how many times we look at the Müller-Lyer illusion, the lines still look to be different lengths. To relieve this existential angst, I try to think of chance like a magic trick. The beauty of the trick lies not in knowing how it's performed, but on the contrary by preserving its mystery. A belief in certainty and the sense of control that it generates may make us happier, yet it is uncertainty and the addictive power of the unknown that keeps driving us towards achieving it. Knowing something is reassuringly pleasurable, but the mystery of things still unknown and the drive to discover them or figure them out is arguably even more so. At least our dopamine circuits imply this. As Buddha once said, what matters is the journey (the anticipation and the search for meaning), not the destination (the reward or the meaning itself). Perhaps everyone who has ever gambled intrinsically knows this.

Having a healthy appreciation for uncertainty and a cognitive style geared towards probabilistic thinking might conceivably help gamblers most predisposed to excesses to avoid them. Addiction is frequently regarded as a loss of control; the craving response is a continuous but ultimately flawed endeavour to reclaim it. Paradoxically, the reverse – relinquishing some control to get it back – is far more constructive. Becoming comfortable with uncertainty, chance and randomness can teach us to feel safer when not in control. Such a skill might arguably have much wider applications. Life, after all, is a journey full of decision making under conditions of uncertainty. Nassim Taleb warns us against being fooled by randomness but also urges us to live by it, wondering whether being unsure might actually make us happier. As an obsessor of time management, that is something I can most definitely relate to. Demanding control over schedules and with an almost pathological fear of being late, I frequently wonder how less stressful life would be if I could live without a watch or a clock, oblivious to the passing of time. Nobel Prize-winning economist Herbert Simon explains that humans lack the cognitive resources for optimal decision making, where we search for the best outcomes. Since we rarely know the relevant probabilities of all possible outcomes, we are unable to evaluate them with sufficient precision. Instead, he proposes that we 'satisfice' rather than maximise, pursuing a course of action that will

satisfy the minimum requirements necessary to achieve a particular goal. Evidently, ‘satisficing’ has much in common with reciprocity, where you don’t have to beat all the competition to do well for yourself. Taleb is unsure about the direction of causality here: are optimisers unhappy because they are optimising, or trying to optimise because they are unhappy? Either way, he believes that “*randomness seems to operate either as a cure or as Novocain.*” If a healthy attitude to gambling can provide an education in probability, uncertainty and chance to help us become better decision makers, surely that is something to be endorsed, not condemned.

Now that I am at the end, I realise that much of what I have written may well be an expression of my own confirmation bias, the telling of stories that endorse my pre-existing beliefs about gambling. I am, after all, a human being, not an automaton. However, the ‘truths’ I have espoused about gambling and its markets – their randomness, wisdom and magical efficiency which make them so hard to beat – are provisional, not absolutes, waiting to be debated, disputed, challenged and perhaps falsified by other Bayesian thinkers with better data and ideas than mine and those whom I have been privileged to learn from. Indeed, the physicist and mathematician Freeman Dyson wouldn’t even consider such things as ‘truths’ at all, but mysteries continually open to exploration. In the meantime, with a bit of ‘luck’, whilst this book may not help you become a more profitable gambler, it might just make you a wiser one.

¹⁵¹ 26 *Predictions: English Premier League forecasting laid bare*, <https://scoreboardjournalism.wordpress.com/2013/08/21/26-predictions-english-premier-league-forecasting-laid-bare/>

¹⁵² <https://scoreboardjournalism.files.wordpress.com/2014/09/eplpredictions20132014.png>

¹⁵³ Men, after all, make up the vast majority of gamblers, for reasons I’ve previously examined.

¹⁵⁴ To some degree, certain bookmakers might still involve themselves actively in price manipulation, particularly in low-liquidity markets with fewer bettors that expose them to more risk, or protection against sharps looking to exploit the systematic biases of squares.

¹⁵⁵ Huberfeld, R., Gersner, R., Rosenberg, O., Kotler, M. & Dannon, P.N., 2013. Football Gambling Three Arm-Controlled Study: Gamblers, Amateurs and Laypersons. *Psychopathology*, **46**(1), pp.28-33.

BIBLIOGRAPHY

The following books are those which I have referred to during my writing and which provide valuable extra material for those readers looking to explore the subject matters I have examined in further detail. Other journal papers I have used are referenced in the footnotes within the main body of the text.

- Atherton, M.**, 2006. *Gambling: A Story of Triumph and Disaster*. London: Hodder & Stoughton Ltd.
- Axelrod, R.**, 2006. *The Evolution of Cooperation*. Revised edition. New York: Basic Books.
- Bernstein, P.L.**, 1998. *Against the Gods: The Remarkable Story of Risk*. New edition. New York: John Wiley & Sons.
- Brenner, R. & Brenner, G.A.**, 1990. *Gambling and Speculation: A Theory, a History, and a Future of Some Human Decisions*. Cambridge: Cambridge University Press.
- Breuning, L.**, 2012. *Meet Your Happy Chemicals: Dopamine, Endorphin, Oxytocin, Serotonin*. 4th edition. Inner Mammal Institute.
- Burnham, T.**, 2008. *Mean Markets and Lizard Brains: How to Profit from the New Science of Irrationality*. Revised edition. New York: John Wiley & Sons.
- Charles, R.H.**, 1924. *Gambling and Betting: A study dealing with their origin and the relation to morality and religion*. Edinburgh: T. & T. Clark.
- Comley, P.**, 2012. *Monkey with a Pin: Why you may be missing 6% a year from your investment returns*. Ebook published by Peter Comley.
- Croston, G.**, 2012. *The Real Story of Risk: Adventures in a Hazardous World*. New York: Prometheus Books.
- Ellenberg, J.** 2014. *How not to be Wrong: the Hidden Maths of Everyday Life*. London: Penguin Books.
- Epstein, D.**, 2014. *The Sports Gene: Talent, Practice and the Truth About Success*. London: Yellow Jersey Press.

- Epstein, R.A.**, 2009. *The Theory of Gambling and Statistical Logic*. 2nd edition. Burlington: Academic Press.
- Frank, R. & Cook, P.J.**, 2010. *The Winner-Take-All Society: Why the Few at the Top Get so Much more Than the Rest of Us*. London: Virgin Books.
- Isaacs, N.D.**, 2001. *You Bet Your Life: The Burdens of Gambling*. Lexington: University Press of Kentucky.
- Kahneman, D.**, 2011. *Thinking Fast and Slow*. London: Penguin Books.
- Koch, R.**, 2014. *The 80/20 Principle and 92 Other Powerful Laws of Nature: The Science of Success*. New edition. London: Nicholas Brealey Publishing.
- Leslie, I.**, 2014. *Curious: The Desire to Know and Why Your Future Depends on It*. London: Quercus.
- MacKenzie, W.D.**, 1893. *The Ethics of Gambling*. Philadelphia: Henry Altemus.
- Malkiel, B.G.**, 2012. *A Random Walk Down Wall Street: The Time-Tested Strategy for Successful Investing*. 10th revised edition. New York: W. W. Norton & Company. Originally published in 1973.
- Mauboussin, M.**, 2012. *The Success Equation: Untangling Skill and Luck in Business, Sports and Investing*. Boston: Harvard Business Review Press.
- Mazur, J.**, 2010. *What's Luck Got to Do with It? The History, Mathematics and Psychology of the Gambler's Illusion*. Princeton: Princeton University Press.
- Ridley, M.**, 1997. *The Origins of Virtue*. New edition. London: Penguin Books.
- Ruden, R.A.**, 2000. *The Craving Brain*. 2nd edition. New York: Harper.
- Scarne, J.**, 1961. *Scarne's Complete Guide to Gambling*. New York: Simon and Schuster.
- Shermer, M.**, 2004. *The Science of Good and Evil: Why People Cheat, Gossip, Care, Share and Follow the Golden Rule*. New York: Holt Paperbacks.
- Shermer, M.**, 2008. *Mind of the Market: Compassionate Apes, Competitive Humans, and Other Tales from Evolutionary Economics*. New York: Holt Paperbacks.
- Silver, N.**, 2012. *The Signal and the Noise: The Art and Science of Prediction*. London: Penguin Books.

- Smith, A.**, 2014. *The Theory of Moral Sentiments*. CreateSpace Independent Publishing Platform. Originally published in 1759.
- Smith, A.**, 2014. *The Wealth of Nations*. CreateSpace Independent Publishing Platform. Originally published in 1776.
- Surowiecki, J.**, 2005. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few*. New edition. London: Abacus.
- Syed, M.**, 2011. *Bounce: The Myth of Talent and the Power of Practice*. London: Fourth Estate.
- Taleb, N.N.**, 2007. *Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets*. 2nd edition. London: Penguin Books.
- Tetlock, P.E.**, 2005. *Expert Political Judgment*. Princeton: Princeton University Press.
- Veitch, P.**, 2010. *Enemy Number One: The Secrets of the UK's Most Feared Professional Punter*. Newbury: Racing Post Books.
- Zak, P.J.**, 2013. *The Moral Molecule: the new science of what makes us good or evil*. London: Corgi.

About the author



Joseph Buchdahl runs the website Sports-Tipsters.co.uk, independently verifying online sports betting advisory services to provide a measure of quality control and a means of demonstrating transparency, validity and reliability for a sports tipping industry that is otherwise unregulated. He is also the author of *How To Find a Black Cat in a Coal Cellar* and the perennial bestseller *Fixed Odds Sports Betting*, published by High Stakes Publishing

Copyright

This ebook edition first published in 2016
by High Stakes Publishing an imprint of
Oldcastle Books
PO Box 394,
Harpden, AL5 1XJ
highstakespublishing.co.uk

All rights reserved
© Joseph Buchdahl 2016

The right of Joseph Buchdahl to be identified as author of this work has been asserted in accordance with Section 77 of the Copyright, Designs and Patents Act 1988.

This ebook is copyright material and must not be copied, reproduced, transferred, distributed, leased, licensed or publicly performed or used in any way except as specifically permitted in writing by the publishers, as allowed under the terms and conditions under which it was purchased or as strictly permitted by applicable copyright law. Any unauthorised distribution or use of this text may be a direct infringement of the author's and publisher's rights, and those responsible may be liable in law accordingly

ISBN
978-1-84344-858-7 (print)
978-1-84344-859-4 (epub)
978-1-84344-860-0 (kindle)
978-1-84344-861-7 (pdf)

For further information please visit oldcastlebooks.co.uk