Algos & Egos – Understanding Modern Financial Markets

Robert Hillman

Excerpt

 robert.hillman@neuroncapital.com

Note. This is an excerpt from a draft of a book I have been working on (for a long time!) called Algos & Ego. Beneath this excerpt I have included the first few pages.

The book is in draft, has not been proofread and constantly evolving. I am making it available in good faith and encourage constructive criticism!

The Search for Yield and Destabilising Dynamics

Health Warning

This section is hard and should probably only be tackled first by readers with significant knowledge of interest rate swaps and FX options. On the other hand I believe the story offers rich insights that will serve the reader well in understanding how derivatives markets and underlying markets interact – a major theme of the book. It also offers thoughts on how we should think about the role of models in markets and how well-intentioned risk management can often fail. If you are struggling to follow it on your first attempt then I would recommend skipping it and working through Part 2 of the book and then revisiting. My aspiration is that with Part 2 on board the concepts and trains of thought in this section will become a lot clearer.

Searching for Systemic Risk from within the Bank of England in the early Naughties

The first time I personally considered the role of option markets as a source of financial instability was in my last year at the Bank of England. Far from a popular opinion among today's heterodox economists that the Bank was too busy dreaming about rational agents and dynamic stochastic general equilibrium markets, in the early 2000s there were many people across both the Financial Stability ('FS') wing and the 'Markets' area of the Bank that were trying to identify and understand the mechanics and practicalities of complex financial products that had the potential to cause damage. Key figures driving these kinds of inquiries were Paul Tucker and Andy Haldane, the latter going on to become Chief Economist at the Bank and a leading advocate for the agent-based-modelling methodologies that feature heavily in this book.

There was a feeling that to understand these risks required a mix of mathematical geekiness and market knowledge. For years the pointy-heads from the Monetary Instruments and Markets Division ('MIM') took turns to (sometimes a little reluctantly) attend daily meetings in the Markets area, but in 2002 I became the first 'PhD' to permanently move from MIM to the analysis team of the Foreign Exchange Division. From that point there were regular internal secondments of quants and economists down to Markets. Markets was the informal term for three divisions, Foreign Exchange, Gilt-Edged Money Markets, and Risk Management. The first two were essentially the front-office of the central bank. It was from Markets that the Bank attempted to defend the pound in 1992, and from here the Bank sold much of the nation's gold on behalf of the Treasury – much against the Bank's advice.

The Search for Yield

One of the last projects I was involved in epitomised this new type of risk that was appearing on the radar. The underlying fear was that with interest rates low, a 'search for yield' could lead to the creation of investment strategies or products that could turn out to have unpleasant consequences should the background change. One of the places we felt could be particularly vulnerable to these pressures was Japan. Japan had been experiencing deflation since the mid-90s and nominal bond yields had been relentlessly falling for several years. The ten-year Japanese Government Bond yield had dropped below 1% in late 2002. Against this backdrop several banks had created investment products that sought to provide Japanese domestic investors with an optically higher yield. But, inevitably, attractive 'optics' had to be balanced by something unattractive, which we thought in this case involved a risk of amplifying market movements at some point in the future. The Bank also identified that the complex structuring and passing around of cashflows and risks meant there was a risk of contagion across institutions. There was a concern that banks important to the UK domestic economy could find themselves in a sticky position. As it turned out this particular risk did not end up bringing down any systemically important UK financial institutions, but it did end up playing a role several years later during the GFC in early 2008. Risks created five years earlier by banks trying to satisfy retail demand for yield, combined with risks taken much more recently by a small group of hedge funds. In early 2008 these risks collided and created market stress that pre-empted the collapse of several of the funds. I was in one of the funds that was exposed to these risks (but survived) and so had a ring-side seat. One of the funds that was not so fortunate was called JWM Partners. Named after the founder John W. Meriwether, this fund was the most prominent of a bunch of funds that had spun out of the collapse of the infamous Long Term Capital Management fund almost exactly ten years earlier. Nearly all of them were out of business within a year.

Complexity Epitomised

The complexity of the product central bankers was concerned with in 2003 wasn't exactly hidden to be fair, given the name of it was Power Reverse Dual Currency Notes (PRDCs). Not exactly intuitive. From an investor perspective the product appeared attractive as it looked a bit like a long-term bond-like investment that was principal-protected and offered a high initial coupon. To offer a high-coupon in a low yield environment required some extremely creative financial engineering, and that was where it got complicated.

After the initial coupon the subsequent coupons were linked to the dollar-yen exchange rate. In some of the notes the coupon was to be paid out in the foreign currency (often US dollars, but another popular currency was Australian dollars). If the yen depreciated (as it had been doing during the mid-90s) the coupon would rise. But the note had another critical feature. It was callable by the issuer. This meant that the issuer could decide to stop paying coupons and hand the principal back to the investor long before the notional maturity of the bond, which was often twenty or thirty years in the future. In markets' parlance the issuer was 'long an option' – they had some control over their future obligations to the investor. On the flipside the investor was short an option. They had in effect sold something (the call option to the issuer) and in return been paid a premium in the form of a high initial coupon.

The callability limited the possible upside to the investor, but investors had got used to, and expected to continue, simply rolling into a new note as soon as the old note was called. Another possible downside for the investor was having their principal tied up and being stuck with low or even zero coupons for years. This was a risk that investors appeared willing to take, presumably lured by the high initial coupon, the attractive sounding "principal protection", and almost certainly betting on continued yen weakening which would – at worst – lead them to being paid back early. As we will see in the next Part, the willingness to extrapolate based on past experience is endemic in financial markets.

The issuers themselves passed on (for a fee) the headache of risk managing the product to third-parties we will call dealers, often investment banks. These dealers included Japanese banks like Mizuho and Nomura but also western banks like Goldman Sachs, JP Morgan, Citigroup, and Credit Lyonnais[BoE 04]. The end-issuers were often sovereign and supra-national institutions like the European investment Bank. Banks with retail links were

also involved and in the UK Lloyds TSB and Royal Bank of Scotland were among the biggest issuers. This no doubt stimulated the UK central bank's interest.

The issuers needed to use investment banks because implementing a risk management process for these products is fearsomely hard. In the event of the yen appreciating they would find themselves less likely to be able to call the option early, and thus the expected maturity of the note would quickly increase. When they had issued the note their models had told them the chance of the note still being alive in thirty years' time was negligible, and so they would not have hedged that eventuality. But if the yen appreciated sufficiently early on, then suddenly the prospect of the note surviving for decades came into play. In that event there would be a sudden need for a longer-term hedge as the risk horizon facing the issuers extended. It is this kind of nonlinear change in risk profile that actually underpins many structured products. The attractive optics of a nice short-dated investment was designed to lure in investors and it distracted them from a low probability downside. The balance of all these risks was far from intuitively obvious and required some very complicated financial modelling.

False Senses of Security and Model Validation

A few years after I left the Bank I gained a nice personal insight into how quants who were structuring these products approached the modelling. I visited an ex-colleague of mine, Thierry Roncalli, in Paris. He is the quintessential French quant. When I first met him in 1998 he and I were the first two employees of Mark Salmon's brand new Financial Econometrics Research Centre at City University Business School. Thierry arrived with a young family and I helped him find a house in London. I did not realise at first, but it turned out Thierry was already somewhat famous in the quant finance community for having written an extensive library of code in a language called Gauss. Having spent my PhD studying a mixture of econometrics and neural networks, I was completely baffled when he asked me how I treated my 'Greeks'. I think it was at that point I stopped saying I had a PhD in finance.

When I saw Thierry years later in Paris he was head of a model validation team for a giant French investment bank. The team was large, completely French, and very, very quanty. I was interested in what he thought about these potential systematic risks but he explained that while he had some views, modelling the systemic risks was not in his remit. It dawned on me that 'Model validation' was a very introspective exercise. It meant checking that the pricing and risk models that were being used by sales and traders were robust and accurate. But what does accurate mean when you are trying to price something that is incredibly complex, a function of multiple essentially random variables, and whose future performance and possible market impact would be almost entirely unpredictable? In truth Model Validation meant checking that the models were true to any academic foundation should it have existed, but much more importantly that they gave consistent answers over time, and across different variations of traded products. To test the latter you would need to confront the models with actual market prices. But herein lies a conundrum. The market prices themselves would have been in large part spat out by the same, or very similar models. This kind of self-referentialism is an issue that will crop up in other places in this book and makes the study of finance (and economics) so fascinating, hard, and different to many other sciences. It was great catching up with Thierry, but I left feeling a slight unease about the amount of confidence organisations placed in things that sound clever and reassuring like Model Validation. Was anyone thinking about the broader risks?

Patient Zero

Managing the risk of these notes was clearly difficult. As both the yen-dollar exchange rate, and or relative interest rates between the US and Japan varied, the risk would change. In some states of the world the dealer would be seen to profit, in other to lose money, but just like our Spanish peseta put seller, dealers don't want profits or losses to appear unpredictability in the future. Both profits and losses generate risk, and dealers typically want to keep a 'flat book'.

Most models suggested that to offset the risk of future payments in the longer term one helpful step would be for the dealers to buy long-dated yen call options. To be clear they needed to purchase something now that

would offer them the opportunity to buy yen at favourable rates in the future if the yen had appreciated strongly. Unfortunately, but not surprisingly, there are not many natural sellers of twenty or thirty year maturity FX options. As we will see in later chapters, one solution to trying to obtain a long-term hedge is to continuously take a series of short-term hedges. For the FX risk the dealer could look to buy one-year yen call options, and then as the year was nearly up roll into another one-year ahead option and so on. This imbalance in demand and supply was reflected in yen option prices: low-strike (i.e. for levels when the yen was relatively strong versus the US dollar) yen calls were significantly more expensive than yen puts, mirroring the option skew that had become evident in US equity index options after 1987. The other much simpler way to try and hedge the risk of the yen appreciating is via replicating the option exposure by trading the underlying. In this case the process is easy to describe. If you are concerned about the yen appreciating then as the yen starts appreciating simply buy some yen. If it falls back, sell it. But if it continues to appreciate, buy some more. As we will see in Part 2, this continuous process of trading pro-cyclically like this creates a pay-off that resembles the shape of a call option. But the potential problem with this kind of process is that if enough people are doing it simultaneously, it stands to amplify movements in the exchange rate – buying begetting more buying, making the whole problem worse.

Another risk that dealers faced if the yen appreciated was that they would be much less likely to call the note early. Very quickly the expected maturity of the note could shift from a couple of years out to thirty years or more. In one sense this might seem great. The issuer gets to sit on the investor's principal for a long time, but on the other hand it dramatically alters the hedge required. Now all of a sudden there could be a decent chance that the dealer may want to terminate the note at some unpredictable point in the distant future. They still 'own' this optionality but previously had not valued it much, but with a much longer expected duration this eventuality now created a long-term risk. The risk would come into play if in the distant future the yen exchange rate depreciated (rather than in the near term as the investors and issuers had originally anticipated). The dealers' models indicated that the major drivers of the likely long-term value of the exchange rate are the relative interest rate differentials. In the short-term, the spot exchange rate volatility (as it reacts to news or data releases) is the dominant driver of exchange rate uncertainty because relative interest rate differentials are not likely to move very much. But looking decades ahead shifts in interest rates are much more likely to dominate. Other things being equal, the yen would depreciate far into the future if long term interest rates in Japan rose relative to the US. This meant that suddenly there would be requirement for dealers to hedge against long term interest rate movements. Specifically, they would want to take an exposure in Japanese interest rates that would create profits for them in Japanese rates fell relative to current expected rates, and losses if rates rose. They therefore entered into long-dated interest rate swap agreements with banks where the dealers received a fixed rate payment for 30 years and agreed to pay the floating short-term rate for 30 years. Interest rate swaps, even long-dated ones, are much more liquid markets than FX options, but nonetheless the sudden rehedging by dealers that could follow a sharp yen appreciation would generate a sudden demand to receive a fixed interest rate. When the demand for something suddenly increases so does the price, in this case it meant that it put downward pressure on long-dated yields in Japan.

Fast Forward to March 2008 – Hedge Funds and The Box Trade

It wasn't until 2008 when the yen did start appreciating rapidly against the US dollar that the stress scenario the Bank of England had been considering some five years earlier began to look like it might come into play. As an article in the FT in April 2008 described the situation: "It's been quite extraordinary what has happened," confesses one leading bond salesman at a major Japanese institution. "The market dislocations we have seen are unprecedented in the career of most traders." Sound familiar?

What had happened was that the scenario the Bank had anticipated had indeed started playing out, but it occurred entangled with a different dynamic whose source lay with a relatively small number of hedge funds. Two of these funds were in fact direct descendants of LTCM, JWM Partners and Platinum Grove Asset Management (which included Myron Scholes one of the high-flying academic names most associated with LTCM). Others have been described by Ludwig Chincarini (whose book provides a fascinating insiders account of

this period featuring lots of interviews with key players) as Copycat funds: Smith Breeden, Endeavour, III, and Parkcentral Global. Endeavour had spun out of Salomon's London arbitrage team in much the same way LTCM had. Endeavour, another similarly focused fund Rubicon, and the fund I worked for, London Diversified Fund Management were situated on adjacent floors in a small office building on Mount Street in Mayfair, London. Each fund was individually secretive, but it was no secret that we all shared the same sales teams at the large investment banks and brokers. Occasionally we shared meetings when visiting economists or strategists from investment banks popped by.

These hedge funds had noticed that there was something of a cyclicality in the spread between long dated swap spreads. This is the spread between government backed bonds, and interest-rate swaps as issued by commercial banks. It was measured as the yield on a swap minus the yield on a government bond with a similar term. As commercial banks were (an assumption soon to become less failsafe) deemed to be riskier than governments, this spread was always positive. You needed to be paid more to engage in a similar transaction with a commercial bank than with a government. Inspection of the historical swap-spread chart indicated it could fluctuate but would never go negative. And simple statistical models could be estimated on the spread time series that would offer a numerical quantification of wide it could go, and how regularly it might be expected to fluctuate around its long-term average. Armed with this quantitative analysis, when it was towards the bottom end of the range hedge funds were tempted to place bets that it would revert back towards, and probably beyond its historical average. To implement the hedge fund would buy long-term Japanese Government Bonds (JGBs) and enter into a long-dated swap whereby the fund paid a fixed rate and received a floating rate. This type of trade is called relative-value because it takes a view on the relative pricing of different assets (in this case government bonds and swaps). It is also an example of a quantitative trade. It could be 'discovered', sized, and managed by the application of statistical analysis.

There was also another aspect of the trade that appealed to hedge funds. At certain times of the year, especially fiscal year end (March) life insurance companies were typically looking to buy long-dated bonds. Their potential obligations are obviously very long term and naturally they would look to match those liabilities with assets (bonds) with similar durations. Hedge funds thought they understood this supply-and-demand dynamic and that it meant that at certain times the yields on long-dated bonds could be temporarily driven by demand and supply imbalances. This is a perfect opportunity for a hedge fund, with no such forced obligations, to step into the market and take advantage of the temporary mispricing. During 2007 the long-dated swap spread was falling to the bottom of its range and this sucked hedge funds in.

To implement the trade as just described required taking risk in two products, long-dated JGBs and swaps. But being hedge funds, they thought they could be smarter. At the same time as putting on the risk in the 20-year section of the yield curves, they put on the reverse trade on shorter-dated instruments. They sold JGBs with a 7-year maturity and entered into 7-year swaps that would profit from falling rates (they received a fixed rate and paid a floating rate). One reason for doing both the shorter dated and the longer dated spreads was that the spread of the spreads looked even more attractive from a historical and statistical perspective. If you drew the two yield curves (the swap curve being above the government curve) and you put a dot on each point of the four points on the curves you had a position and joined the dots you had a kind of box shape, hence the trade was called the box trade. This box spread was even more mean-reverting, and looked like a no-brainer as long as the fund could hold onto the exposure for long enough. Another implication of this statistical observation was that the measured risk in the box trade was less than in the long-dated leg alone. In practice this means that a fund could take more positions onto its books and end up lowering its (measured) risk. Clearly there is something wrong with this idea, and it is a concept we will return too many times.

This was a classic hedge fund trade. It had something for everyone. The quants predisposed to using algos and statistics could see the trade's potential. The risk management team (if it existed) would be pleased their sophisticated modelling was able to conjure up risk numbers with great numerical precision. And the egos in the room felt smart that they were taking advantage of a temporary structural demand and supply imbalance that

could be rationalised intellectually and reinforce a marketing message that hedge funds could deliver strong returns because they could take opportunities others could not.

Risks Collide

In early March 2008 the scenario the Bank of England had identified began to play out. As the global credit crisis deepened, several US banks were looking to be on the ropes, most noticeably Bear Stearns. The yen started appreciating during the first few months and this kicked off need for sudden re-hedging by the PRDC dealers. The problem was that as the predicted downwards pressure on long-term swap rates began to be realised it led to losses for the hedge funds who had been betting on those long-term rates rising relative to JGBs. At the same time, as concerns over credit rose US Libor rates rose, pulling up shorter term swap rates the world over. This hurt the other part of the hedge funds box trades. One thing that LTCM had learned was that it was probably better to react to losses sooner rather than later, and some of the hedge funds began trying to exit their positions as soon as losses started building. Many hedge funds had risk management processes in place that forced them to cut risk in response to losses. For example, an individual trader within such a fund might be forced to cut their risk by half if they lost 5%, and by half again if they lost a further 5%. This discrete risk reduction system is sometimes called a traffic-light system. Similar automatic risk reduction processes would likely have been implemented at a fund level as well. These measures were put in place with the best of intentions and were often a central plank of the marketing pitch. But as will begin to sound like a broken record in this book, often what seems sensible at an individual level in normal times, may be very suboptimal in practice is everyone is trying to do the same thing in abnormal times.

In order to unwind their positions the hedge funds had to engage in offsetting trades. Their box trade had them paying a long-dated (20 year or more) fixed rate and receiving floating rates, betting long term rates would rise. To unwind this they need to go the other way and enter into agreements to receive a fixed rate and pay a floating rate. But this only added to the downward pressure on swap yields being generated by the PRDC hedgers. Within a few days the long-term swap spread moved sharply negative (remember all the historical analysis had indicated this had never happened before). Meanwhile in exchange rate land the exact scenario the Bank had identified played out. As the yen had started appreciating, PRDC hedging caused more yen buying. Both yen appreciation and movements in yield curves became strongly amplified. It also became quickly apparent that the same box trade was held by a handful of hedge funds now under obvious strain, and just like had occurred ten years earlier this knowledge would have only served to worsen their plight. As each fund tried to stop out of their positions simultaneously liquidity evaporated and prices made those extraordinary moves that our trader quoted in the FT reported. The impact on the hedge funds was in many cases to all intents and purposes terminal. It was reported Endeavour lost 27% in a single day in March 2008. As the year wore on more and more relative value trades went wrong amidst deteriorating market conditions with counterparties pulling back from providing liquidity, if they did not collapse themselves. By the end of the year losses were huge. Of the funds that were caught in the Japanese box trade in March 2008, JWMP was down 42%; Platinum Grove 50%; Endeavour and III 54%; and both Smith Breeden and Parkcentral were down 100%. John Meriwether shut his fund down for the second time in 2009 (Source: Ludwig Chincarini, The Crisis of Crowding, Bloomberg Press 2012).

Key Takeaways

In the end the anticipated PRDC hedging risk did not bring down any systemically important institutions, but it certainly played a role in adding to stresses that were already building in the financial system. It did wipe out several hedge funds, but these were clearly not, unlike as had been feared with LTCM in 1998, systemically critical. Many of the original retail investors, which included both individuals and municipals also found themselves stuck with an illiquid investment product that because of its potential duration could feasibly outlive them. These losses were not enough to cause systemic hardship, but enough to generate concerns over investor protection and the marketing of the products.

For me the story has two messages. The first is about how risks created for one client, or some situation, can create a complex web of interconnected risks that dynamically change over time. This web of connectedness can spread out over time and space, across institutions and people. In some cases the longevity of these risks can be strikingly long. As we will see later in the book, for a variety of reasons more and more investors are exposed to, and care about very long-term risks to their investments.

The second message is that apparently individually sensible risk management practices may do than worse than fail to protect capital. They can actually do the exact opposite of what they are intended and increase losses not reduce them. The principle reason they fail is because when designing risk management practices institutions struggle to anticipate what the market environment will be like at the time at which they are likely to implement their risk protocols. They invariably overestimate the success with which they can reduce risk. They appear to pay little attention to what others might be doing at the same time. There are two immediate effects that we can see time and time again when these events play out. Firstly, movements in markets can be amplified. Volatility becomes larger than expected. Secondly, we see markets suddenly become highly correlated, beyond what had been witnessed in the past. In this case we saw the yen rising rapidly and the yields on long-dated interest rate products also falling together. These sudden rises in correlation between previously fairly disconnected markets can cause a lot of problems. Later in the book we will explore ways in which it might be possible to improve this situation.

Opening Sections to the book

Why

In September 2017 the "central bank's banker" the Bank for International Settlements warned that against a backdrop of low inflation and interest rates, there may be "amplification mechanisms" that "lie in wait for unwary fixed income investors". Such mechanisms threaten to cause interest rates to overreact to news, and potentially lead to systemic risks through the creation of financial instability. Shortly after, the International Monetary Fund shined a spotlight onto certain types of algorithmic trading strategies that have become popular among the world's largest asset managers. These were not high-frequency strategies of the Flash Crash kind that had already gained public consciousness. These were much lower frequency and seemingly sensible approaches to adjusting portfolio allocations over time with the aim of improving the investment realisations of millions of investors who entrust their savings to professional investors. In 2018, David Harding a leading fund manager and pioneer of algorithmic trading strategies designed for institutional investors, complained that the market was becoming too crowded. He had been forced to rapidly adapt his long standing and highly profitable strategy for fear of subjecting his investors to an inevitable long and painful death and he also acknowledged the risk that his own trading could contribute to widespread financial volatility. Other less manifestly quantitative, but no less successful traders, had also been complaining that markets were becoming too difficult to trade but Harding's warning came from an insider, marking a new stage in the perception if not the reality that there could be unpleasant unintended consequences from the application of individually sensible risk management practices.

Why is there so much controversy about apparently innocent questions like the impact of how investors choose to invest or manage their risk? One of the fiercest battlegrounds today is around a simple choice investors face. Do you entrust your money to an active manager who will try to pick and choose where to invest it, or do you give it to a tracker fund that will employ an algorithm that mimics – no more or no less – than the performance of a market-cap weighted index, commonly known as the 'market'. There are some who see this battle as marking the final maturing of a market too long the scourge of snake-oil salesmen peddling expensive yet valueless funds. But others see it as providing a catalyst for the next financial crisis, this time led by innovation in financial engineering of retail products. These questions are not just of interest to academics, professional investors and managers. Two other themes, the distrust of experts and the rise of populism are lurking in the

background. The atmospheric conditions for financial market volatility to trigger a radical shift in societal and governmental attitudes towards financial markets seem highly charged.

This book is about how we might think about these issues scientifically. How can we approach the question of whether further growth in algorithmic trading strategies may impact market dynamics? What kinds of information should investors think about when faced with choosing between, or even sticking with, investment products that appear to have stopped working? History is of limited use. Many of today's products have only existed a few years, and as the saying goes - it's not your grandfather's market. Regulator interest has been piqued, but how are they addressing these issues?

How

This book explores why and how computer simulation methods like agent-based-modelling may provide answers to questions facing players in today's financial markets. Computer simulation is a type of generative science that embraces notions of complexity and sits between empirical data science and theoretical modelling. I will make parallels with meteorology where scientists face similar challenges – measuring and assessing the cause and implications of extreme weather without sufficient and relevant historical data to support a purely data-driven approach. I offer examples of how simulation can – today – help us address real world problems like policy making trade-offs, enable real-time risk forecasting, and help portfolio allocation problems facing investors looking to manage their financial well-being through retirement. My approach will be relentlessly practical. For too long I have heard romantic visions of why agent-based-modelling is our only hope to address the flawed orthodoxy of economics. But what is the result? Today in 2019 there are virtually no firms that specialise in agent-based-approaches. I would wager that if you surveyed wealth managers or advisors, through traders, to risk-managers, hardly any would be able to tell you what agent-based modelling is, and next to none would be able to give an example of where they themselves use it or have indeed even heard of it being used. So – within economics and finance at least - the great agenda, the loudly trumpeted paradigm shift towards complexity science has not yet occurred. I think I know why, and that narrative will emerge through the book. But by the end of the book, my aim is to have presented a thorough description of what computer simulation entails, why it is useful, how we can implement it to help us get closer to better investment outcomes.

A large part of my story is that markets have evolved to look like a computer simulation rather than the other way around. In this sense the early, yet somewhat disappointing, aspiration of modelling human behaviour using analogies of evolutionary dynamics and computational learning techniques was not daft, it was merely ahead of its time. To get across this part of the story across I will describe how algorithms came to play a role in determining how modern markets operate and how they work. I intend to give as many case studies and real-world examples as possible, drawing on my own experience as a quant and trader in hedge funds and earlier as a quantitative economist in a central bank. My aim is to be ruthlessly honest, telling both sides of the story. I am a trader and thus a risk-manager and I am strongly motivated to pursue lines of thinking that have positive economic consequences for me.