

Introduction

1.1 EXTREME EVENTS

This book is about how best to construct investment portfolios if *a priori* it is reasonable to assume that markets *might* exhibit *fat-tailed behaviour*. It is designed to appeal to a wide variety of practitioners, students, researchers and general readers who may be interested in extreme events or portfolio construction theory either in isolation or in combination. It achieves this aim by

- (a) Exploring extreme events, why they might arise in a financial context and how we might best analyse them.
- (b) Separately exploring all the main topics in portfolio construction theory applicable even in the absence of fat tails. A special case of any more general approach capable of effectively handling extreme events is the situation where the extent of fat-tailed behaviour is too small to be discernible.
- (c) Blending points (a) and (b) together to identify portfolio construction methodologies better able to cater for possible fat-tailed behaviour in the underlying assets or liabilities.

Given its wide intended audience, the book covers these topics both from a more qualitative perspective (particularly in the earlier and later chapters) and from a more quantitative (i.e., mathematical) perspective (particularly in the middle chapters). Where possible, this book has been segmented so that valuable insights can be gained without necessarily having to read the whole text. Conversely, in the author's opinion, valuable insights arise throughout the book, including the parts that are more mathematical in nature. More general readers are therefore encouraged not to skip over these parts completely, although they do not need to worry too much about following all the details.

By *fat-tailed behaviour* we mean that the distribution of future returns is expected to involve more extreme events than might be expected to occur were returns to follow the (multivariate) (log-) Normal distributions often assumed to apply to markets in basic portfolio construction texts.¹

Most practitioners believe that most markets are 'fat-tailed' given this terminology. There is a wide body of empirical academic literature that supports this stance, based on analysis of past market behaviour. There is also a growing body of academic theory, including some involving behavioural finance, explaining why fat-tailed behaviour seems so widespread. So, we might also characterise this book as exploring how best to construct investment portfolios *in the real world*.

Of course, practitioners and academics alike are not themselves immune from behavioural biases. It is one thing to agree to pay lip service to the notion that market behaviour can be

¹By 'multivariate' we mean that the returns on different return series have a joint distribution, the characterisation of which includes not only how individual return series in isolation might behave, but also how they behave when considered in tandem, see Chapter 3. By '(log-) Normal' we mean that the natural logarithm of $1 + r$ is Normally distributed, where the return, r , is expressed in fractional form, see Section 2.3.1.

fat-tailed, but quite another to take this into account in how portfolios are actually constructed. Following the dot.com boom and bust in the late 1990s and early 2000s, markets settled into a period of unusually low volatility. Strategies that benefited from stable economic conditions, e.g., ones that followed so-called ‘carry’ trades or strategies relying on continuing ready access to liquidity, proved successful, for a time. The 2007–09 credit crisis, however, painfully reminded the complacent that markets and economies more generally can and do experience extreme events.

1.2 THE PORTFOLIO CONSTRUCTION PROBLEM

We do not live in a world in which we have perfect foresight. Instead, portfolio construction always involves striking a balance between risk and reward, i.e., the risk that the views implicit in our portfolio construction will prove erroneous versus the rewards that will accrue if our views prove correct. Everyone involved in the management of portfolios, whether of assets or of liabilities, faces a *portfolio construction problem*. How do we *best* balance risk and return? Indeed, what do we mean by ‘best’?

Given the lack of perfect foresight that all mortals face, it is not reasonable to expect a book like this to set out guaranteed ways of profiting from investment conditions come what may. Instead, it seeks to achieve the lesser but more realistic goal of exploring the following:

- (a) core elements of portfolio construction;
- (b) mathematical tools that can be used to assist with the portfolio construction problem, and their strengths and weaknesses;
- (c) ways of refining these tools to cater better for fat-tailed market behaviour;
- (d) mindsets best able to cope well with extreme events, and the pitfalls that can occur if we do not adopt these mindsets.

1.3 COPING WITH REALLY EXTREME EVENTS

Lack of perfect foresight is not just limited to a lack of knowledge about exactly what the future holds. Typically in an investment context we also do not know how uncertain the future will be. Using statistical terminology, we do not even know the precise form of the probability distribution characterising how the future might evolve.

The differentiation between ‘risk’ and ‘uncertainty’ is a topic that several popular writers have sought to explore in recent times, e.g., Taleb (2004, 2007). In this context ‘risk’ is usually taken to mean some *measurable* assessment of the spread of possible future outcomes, with ‘uncertainty’ then taken to mean lack of knowledge, even (or particularly) concerning the size of this spread.

In this book, we take up this baton particularly in Chapters 8 and 9. Holding such an insight in mind is, I think, an important contributor to successful portfolio construction. In particular, it reminds us that really extreme events seem to have a nasty habit of occurring more often than we might like. Put statistically, if there is a 1 in 10^{10} (1 in 10 billion) chance of an event occurring given some model we have adopted, and there is a 1 in 10^6 (1 in a million) chance that our model is fundamentally wrong, then any really extreme events are far more likely

to be due to our model being wrong than representing random (if unlikely) draws from our original model.²

Yet such insights can also be overplayed. The portfolio construction problem does not go away merely because the future is uncertain. Given a portfolio of assets, someone, ultimately, needs to choose how to invest these assets. Although it is arguably very sensible for them to bear in mind intrinsic limitations on what might be knowable about the future, they also need some framework for choosing between different ways of structuring the portfolio.

This framework might be qualitatively formulated, perhaps as someone's 'gut feel'. Alternatively, it might be quantitatively structured, based on a more mathematical analysis of the problem at hand. It is not really the purpose of this book to argue between these two approaches. Indeed, we shall see later that the outcome of essentially any qualitative judgemental process can be reformulated as if it were coming from a mathematical model (and arguably vice versa).

Perhaps the answer is to hold onto wealth lightly. All of us are mortal. The more religious among us, myself included, might warm to this philosophy. But again, such an answer primarily characterises a mindset to adopt, rather than providing specific analytical tools that we can apply to the problem at hand.

1.4 RISK BUDGETING

Some practitioners point to the merits of risk budgeting. This involves identifying the total risk that we are prepared to run, identifying its decomposition between different parts of the investment process and altering this decomposition to maximise expected value-added for a given level of risk. It is a concept that has wide applicability and is difficult to fault. What business does not plan its evolution via forecasts, budgets and the like? Indeed, put like this risk budgeting can be seen to be common sense.

Again, though, we have here principally a *language* that we can use to describe how to apply investment principles. Risk budgeting principally inhabits the 'mindset' sphere rather than constituting an explicit practical toolset directly applicable to the problem at hand. This should not surprise us. Although sensible businesses clearly do use budgeting techniques to good effect, budgeting per se does not guarantee success. So it is with *risk* budgeting.³

However, language is the medium through which we exchange ideas and so cannot be ignored. Throughout this book, we aim to explain emerging ideas using terms that can be traced back to risk budgeting concepts. This helps clarify the main aspects of the methodology under discussion. It also helps us understand what assumptions need to be made for the relevant methodology to be valid.

² More precisely, in this situation we need the probability of occurrence of the event that we are considering to be much higher (on average) than 1 in 10,000 in the 1 in a million circumstances when our underlying model is assumed to prove to be fundamentally wrong. This, however, is typically what is implied by use of the term 'fundamentally wrong'. For example, suppose that our model presupposes that possible outcomes are Normally distributed with zero mean and standard deviation of 1. Then the likelihood of an outcome worse than c. -6.4 is 1 in 10 billion. However, suppose that there is actually a one in a million chance that our model is 'fundamentally wrong' and that the standard deviation is not 1 but 10. Roughly 26% of outcomes when the standard deviation is 10 will involve an outcome worse than c. -6.4. So, in this instance an event this extreme is roughly 40 times as likely to be a result of our original model being 'fundamentally wrong' as it is to be a fluke draw from the original model.

³ Likewise, no portfolio construction technique is able to *guarantee* success.

1.5 ELEMENTS DESIGNED TO MAXIMISE BENEFIT TO READERS

As explained in Section 1.1, this book aims to appeal to a wide variety of audiences. To do this, I have, as with my earlier book on *Market Consistency*, sought a suitable balance between mathematical depth and readability, to avoid some readers being overly daunted by unduly complicated mathematics. The book focuses on core principles and on illuminating them where appropriate with suitably pitched mathematics. Readers wanting a more detailed articulation of the underlying mathematics are directed towards the portfolio construction pages of the www.nematrian.com website, referred to throughout this book as Kemp (2010).

To maximise the benefit that both practitioners and students can gain from this book, I include two sections at the end of each main chapter that provide:

- (a) Comments specifically focusing on the practitioner perspective. To navigate successfully around markets typically requires an enquiring yet somewhat sceptical mindset, questioning whether the perceived benefits put forward for some particular technique really are as strong as some might argue. So, these sections either focus on the ways that practitioners might be able to apply insights set out earlier in the relevant chapter in their day-to-day work, or highlight some of the practical strengths and weaknesses of techniques that might be missed in a purely theoretical discussion of their attributes.
- (b) A discussion of some of the more important implementation challenges that practitioners may face when trying to apply the techniques introduced in that chapter. Where the same challenge arises more than once, I generally discuss the topic at the first available opportunity, unless consideration of the challenge naturally fits better in a later chapter.

The book also includes an Appendix containing some exercises for use by students and lecturers. Each main chapter of the book has associated exercises that further illustrate the topics discussed in that chapter. The exercises are reproduced with kind permission from Nematrian Limited. Hints and model solutions are available on the www.nematrian.com website, as are any analytical tools needed to solve the exercises.

Throughout the book, I draw out principles (i.e., guidance, mainly for practitioners) that have relatively universal application. Within the text these principles are indented and shown in bold, and are referenced by P1, P2, etc.

1.6 BOOK STRUCTURE

The main title of this book is *Extreme Events*. It therefore seems appropriate to focus first, in Chapters 2 and 3, on fat tails and extreme events. We explore some of the ways in which fat-tailed behaviour can be analysed and the existence or otherwise of extreme events confirmed or rejected. We differentiate between analysis of fat tails in single return series in Chapter 2 and analysis of fat tails in joint (i.e., multiple) return series in Chapter 3. The shift from ‘one’ to ‘more than one’ significantly extends the nature of the problem.

Before moving on to portfolio construction per se, we consider in Chapter 4 some ways in which we can identify what seems to be driving market behaviour. Without some underlying model of market behaviour, it is essentially impossible to assess the merits of different possible approaches to portfolio construction (or risk modelling). We consider tools such as principal components analysis and independent components analysis, and we highlight their links with other statistical and econometric tools such as multivariate regression.

In Chapters 5–7 we turn our attention to the portfolio construction problem.

Chapter 5 summarises the basic elements of portfolio construction, both from a quantitative and from a qualitative (i.e., ‘fundamental’) perspective, if fat tails are *not* present. At a suitably high level, both perspectives can be viewed as equivalent, apart perhaps from the mindset involved. In Chapter 5 we also explore some of the basic mathematical tools that commentators have developed to analyse the portfolio construction problem from a quantitative perspective. The focus here (and in Chapter 6) is on *mean-variance* portfolio optimisation (more specifically, mean-variance optimisation assuming *time stationarity*). We consider its application both in a single-period and in a multi-period world.

In Chapter 6 we highlight the sensitivity of the results of portfolio construction analyses to the input assumptions, and the tendency of portfolio optimisers to maximise ‘model error’ rather than ‘risk-return trade-off’. We explore ways of making the results more *robust* to errors affecting these input assumptions. The academic literature typically assumes that input assumptions are estimated in part from past data. We might argue that asset allocation is a forward-looking discipline, and that the assumptions we input into portfolio construction algorithms should properly reflect our views about what might happen in the future (rather than about what has happened in the past). However, some reference to past data nearly always arises in such analyses. We pay particular attention to Bayesian approaches in which we have some prior (‘intrinsic’) views about the answers or input parameters that might be ‘reasonable’ for the problem and we give partial weight to these alongside partial weight to external (often past) data. The best-known example of this is probably the Black-Litterman approach. Some Bayesian approaches can also be viewed as justifying heuristic⁴ techniques that can be applied to the portfolio construction problem. This again highlights the high-level equivalence that exists between quantitative and qualitative approaches to portfolio construction.

In Chapter 6 we also introduce ‘market consistent’ portfolio construction, in which we derive input assumptions not from past data but from market implied data. Such an approach is ‘robust’ in the sense that the input assumptions are in theory not subject to the same sorts of estimation errors as ones derived from historical behaviour. We also explore tools that practitioners less keen on Bayesian approaches have developed to tackle estimation error, particularly resampled portfolio optimisation. We show that they are less divorced from Bayesian approaches than might appear at first sight.

In Chapter 7 we identify how to incorporate fat tails into portfolio construction theory. We start by exploring what happens when we relax the assumption of time stationarity, by introducing the concept of *regime shifting*. This involves assuming that the world is, at any point in time, in one of several possible states, characterised by different distributions of returns on the different assets and liabilities under consideration. The mixing of distributions introduced in such a model naturally leads to fat-tailed behaviour. We then extend these ideas to encompass more general ways of incorporating fat-tailed behaviour. We focus mainly but not exclusively on situations where the regime is characterised not by a single Normal distribution but by a distributional mixture of Normal distributions (because this type of model is sufficiently general that it can approximate other ways in which fat tails might arise). We also explore approaches that involve continuously varying parameterisations of the different regimes and focus on behaviour in continuous rather than discrete time.

⁴ In this context, a ‘heuristic’ technique is one that is akin to a rule of thumb that is not principally proposed on the basis of some formal mathematical justification but more because the approach is relatively convenient to implement.

Chapters 2 to 7 are largely concerned with a probability-theoretic view of portfolio construction. In them, we identify, somehow or other, a distributional form to which we believe future outcomes will adhere (or more than one in the case of regime shifting). At least in principle, this involves specifying a likelihood of occurrence for any given future scenario. However, the belief that we can in all cases actually identify such likelihoods arguably involves an overly rosy view about our ability to understand the future. More to the point, regulators and other bodies who specify how financial services entities should operate may want to ensure that entities do not put too many of their eggs into what may prove an uncertain basket.

In recent years this has led to an increased focus on *stress testing*, which we explore in Chapter 8. Stress testing, in this context, generally involves placing less emphasis on likelihood and more emphasis on magnitude (if large and adverse) and on what might make the scenario adverse. We can view ‘reverse stress testing’ and ‘testing to destruction’ as being at one extreme of this trend. In the latter, we hypothesise a scenario adverse enough to wipe out the business model of the firm in question (or the equivalent if we are considering a portfolio) irrespective of how likely or unlikely it might be to come to pass. We then focus on what might cause such a scenario to arise and whether there are any strategies that we can adopt that might mitigate these risks.

Chapter 9 extends the ideas implicit in Chapter 8 to consider ‘really extreme’ events. It is more heuristic and ‘mindset-orientated’ than the rest of the book. This is inevitable. Such events will almost certainly be so rare that there will be little if any directly relevant data on them. Market implied portfolio construction techniques potentially take on added importance here. Merely because events are rare does not mean that others are not exposed to them too. The views of others, distilled into the market prices of such risks, may be subject to behavioural biases, but may still help us in our search to handle extreme events better.

Finally, in Chapter 10 we collate and summarise in one place all the principles highlighted elsewhere in the book.