

Introduction

1.1 MARKET CONSISTENCY

This book is about *market consistency*, a term that we use throughout this book as a catch-all for the activity of taking account of ‘what the market has to say’ in financial practice. We will explore, from first principles, when it is (and when it isn’t) most appropriate to listen to what the market is saying. We will also explore how in practice we might extract information from ‘the market’ (and also what ‘the market’ is) in those circumstances where market consistency has merit.

We can characterise the incorporation of market consistency in a piece of financial analysis as involving the creation of a suitable model or idealised abstraction of how something works, which is then calibrated using appropriate market derived information. We focus in this book on models applicable to three core areas of financial practice, namely:

- (a) *Valuation methodologies*, i.e. the placing of values or prices on positions in financial instruments or other sorts of assets or liabilities;
- (b) *Risk management processes*, i.e. the assessment and management of the (financial) risks that are created by holding such positions; and
- (c) *Portfolio construction techniques*, i.e. the selection of which sorts of such positions or risks it is most desirable to hold (or avoid), and in what quantities.

We do so because these three disciplines are closely allied, both in theory and in practice. We do not value something in complete isolation. Instead, there must be reasons for doing so. One key reason is to understand better the characteristics of the positions we currently have. But what is the point of gaining this understanding? Surely, it is to be better placed to understand the potential behaviour that the positions might exhibit in the future. We want to understand the *risks* and possible *rewards* attaching to them, i.e. how they might behave in adverse and favourable circumstances.¹ The natural next step is then to consider and take decisions about how best to manage the risks and potential rewards we face. This in turn naturally leads to the question of which exposures we should adopt in the first place.

¹ We focus in this book principally on activities that concentrate on the present and the future. Another reason for valuing things refers more to the past, namely, to enable owners of an entity to judge how well its managers have performed as custodians of the owners’ assets. Using movements in ‘fair’ (i.e. market or market consistent) values for this purpose is a well-established technique in the context of performance measurement of investment portfolios (or unitised funds), see e.g. Kemp (2005). However, misuse of such techniques within bank management and trader remuneration structures can create systemic instabilities, as Turner (2009) points out when discussing the causes of the 2007–2009 credit crunch. He notes that such remuneration structures can encourage irrational exuberance and lead to excessive rewards being paid for financial innovation that delivers no fundamental economic benefit but merely exploits opacity and asymmetry of information and knowledge. Such rewards may then eventually turn out merely to have been based on ‘illusory’ market value gains.

1.2 THE PRIMACY OF THE 'MARKET'

Why might we care about what the market has to say? The world of finance is rarely far away from anyone these days. For some, life is a day-to-day struggle to make ends meet. Others, more fortunate, may have surplus funds deposited with banks or invested in the multitude of financial products now emanating from the world's financial centres. But even they will often have been in debt at some stage in their lives. Perhaps this will have been to finance the purchase of a car or house, or, for the more entrepreneurially minded, to support a new business venture. The same is true on a larger scale for companies, charities, even entire countries. For better or worse, money, as a means of storing and transferring value, has proved to be one of humankind's more important inventions. Indeed, it has been like this for many centuries, at least for the wealthier end of society. Julius Caesar built up huge debts (like several other Roman politicians of his day) and then amassed huge wealth on his way to seizing supreme power in the Roman Republic. Banking provided the wealth that enabled the Medici patronage of the Renaissance.

The last few decades have arguably seen a spurt in financial innovation. There has been huge growth in derivatives markets and in the range and sophistication of financial products and instruments now available to individuals, corporations and financial entities. In part, this reflects the technological innovation, economic growth and capital accumulation that large parts of the world have seen since the Second World War. It also, in my opinion, reflects the particular focus given during this period by financial theory and practice to the concept of *the market*. By this we mean some possibly hypothetical construct in which whatever we are interested in can be bought or sold without (too much) difficulty. Economic theory has always argued that properly functioning marketplaces are important for effective competition and hence efficient allocation of resources across an economy. The core innovation over the last few decades has been to apply this more general economic insight to finance itself.

We have seen, for example, changes in the underlying business models that many banks have adopted, away from a 'borrow and hold' business model (in which a bank would raise money from its depositors and then itself lend the proceeds, or some multiple of them, to some of its other customers) and towards an 'originate and sell' business model (in which the bank's assets are repackaged and sold to other capital markets participants). Some reversal of this trend is also now apparent as more straightened economic times loom ahead.

The 'sell' element of such a business model ultimately involves transfer of risk exposures to third parties. It is of course facilitated if there are ready markets in such exposures. It places particular importance on the (market) prices at which transfers of these exposures can take place. Many banks and bank-like entities have historically been involved to some extent in market making activities. Increasingly, they have become 'market makers' of their own core cash flow streams.

Other players in the financial markets, such as pension funds and insurance companies, have perhaps been less affected by this fundamental change in financial mindsets. However, they too do not operate in a vacuum. The mere possibility of transferring blocks of pension benefits or policyholder entitlements reminds (or ought to remind) them that they too live not only surrounded by markets but also, in some sense, *within* them. For better or for worse, a focus on 'what the market has to say' is likely to be here to stay.

This does not mean that enthusiasm for 'the market' will not wax and wane over time. As I wrote this book, a number of pillars of the financial services industry were revealing eye-watering sized losses. These were arguably a consequence of their overexuberant hope that markets would continue to operate in ways to which they had previously become accustomed. Some of these firms subsequently defaulted. Others had to be rescued in mammoth government

sponsored bailouts as governments endeavoured to bring stability to their financial systems. Enthusiasm for listening to what the market has to say often diminishes when what it is saying is unpalatable!

1.3 CALIBRATING TO THE ‘MARKET’

Merely observing that there may be useful information extractable from the ‘market’ doesn’t actually help us extract this information. Nor does it help us work out when such information is of most use or how to use it in ways that do not exacerbate systemic, i.e. economy-wide, risks. This is particularly true when the ‘market’ does not have all of the characteristics of the perfectly behaved construct of economic theory. In this book we explore how market consistency can be applied in the world in which we actually live, where markets are imperfect.

As mentioned above we focus in this book on three interrelated disciplines, namely valuation, risk management and portfolio construction. In each case, we focus on practical ways of incorporating greater market consistency, whilst simultaneously providing a systematic treatment of the underlying theory. For example, when dealing with valuation, we explore the theory and regulatory drivers currently favouring greater use of *marking-to-market* (as well as describing some of the current countervailing drivers) and we explore the differences, or often the lack thereof, between mark-to-market and mark-to-model valuations. However, there are also sections offering more practical guidance on how to determine market consistent valuations for assets or liabilities where markets are limited, illiquid or even almost non-existent.

In the field of risk management, market consistency can mean different things to different people. In my view, we are likely in due course to see a paradigm shift towards greater use of market implied risk measures. The principles involved are explored in some detail, because of the important ramifications this would have for practitioners in this field. However, I accept that this view has yet to achieve wide acceptance (not least because of the practical challenges involved and the unpalatable answers it might generate). The book therefore also provides a full treatment of the more limited types of market consistency that are incorporated in current risk measurement and management paradigms. These might best be described as attempting to quantify (probabilistically or otherwise) the ‘real world’ likelihood of some risk materialising over a given timeframe.

The application of market consistency to portfolio construction is simultaneously both core and peripheral. Decisions about what to buy or sell should take account of how much they might cost to buy or sell. However, a world in which the ‘market’ is taken to be the sole arbiter of knowledge is not one that can be fully reconciled with the concept of active management. By this we mean taking views about when the market is right and when it is wrong and acting accordingly. The need here is to understand and take due notice of the market but not to let it be the sole input into your decision-making process.

1.4 STRUCTURE OF THE BOOK

For practical reasons, this book is, in the main, segmented between the interrelated disciplines described above. Most larger financial organisations segment their business by activity. The day-to-day working life of readers within the financial services industry will therefore tend to have a bias towards one of these three disciplines. For example, asset management and investment banking businesses are often subdivided into three parts, a *back office*, a *middle office* and a *front office*. Typically, individuals working in the back office have day-to-day

responsibility for the processes used to administer and value relatively simple instruments. Those working in the front office are responsible for deciding what positions to buy or sell. Those working in the middle office may provide a bridge between back and front offices. It is also becoming the norm for risk management to be explicitly differentiated from the front (and back) office and thus to fall within the remit of the middle office function. Similar types of role distinctions do also apply to insurers and pension funds, but in many cases the activities in question are outsourced or given different names.

The structure adopted by this book is as follows.

Chapters 2 to 6 cover in the main core material applicable to all of the disciplines being considered. Chapter 2 focuses on when market consistency is and isn't appropriate, extrapolating from the properties of money to establish what sorts of properties we might expect *monetary values* to exhibit. From these properties it deduces when it is most and least relevant to be market consistent. It also explores at a high level some of the main drivers currently for and against market consistent approaches, and how this can influence what in practice is actually meant by 'market consistency'. Chapter 3 focuses more on valuation activities, and on how in practice different standards setters and commentators interpret market consistency and other similar terms. Chapter 4 provides a primer on derivative pricing theory. It cannot be claimed that it does so without reference to *any* mathematics. However, hopefully even those less welcoming of complex mathematical arguments will feel that they have gained some useful insights after reading it. In this context 'derivative pricing theory' is really a catch-all for virtually all of the financial theory that underpins the rest of the book (apart from the theory relating to the interaction of risk and return, which we introduce principally in Chapter 12). Chapter 5 explores a particularly important issue in practice when applying market consistent principles to less liquid instruments. This is how to understand and identify a suitable 'risk-free' interest rate or yield curve. For a book that is focused in part on how to handle less liquid markets it is also natural to include in this part of the book a chapter specifically on liquidity theory, i.e. Chapter 6.

Chapters 7 to 10 consider market consistency in the context of risk management. Chapter 7 covers the fundamental theory, such as the description of different sorts of risk measures and how they are typically calculated. Chapter 8 focuses on capital adequacy. It provides examples of how current regulatory frameworks try to identify the appropriate minimum capital that an organisation should hold to protect itself against risks that might lead to insolvency. Chapter 9 explores how to apply market consistency in what might be called the current risk management paradigm. It focuses principally on how market valuations might vary in the future, and hence how the risks being expressed in these positions might be best managed. Chapter 10 considers how these approaches would need further refinement if we want them to adopt a *fully* market consistent paradigm. This involves applying market consistency not merely to the valuations used within the risk assessment but also to the probability distributions ascribed to the future movement of these valuations. Reasons for adopting this paradigm are also covered in this chapter.

One claim some commentators made during the 2007–09 credit crisis was that inappropriate use of marking-to-market techniques can create a lack of confidence in the financial soundness of banks in stressed market conditions, which can then become a self-fulfilling prophecy. It behoves a book on market consistency to consider carefully this logic. We explore in Chapter 11 ways in which as a society we may best protect against system-wide concerns whilst not diluting other benefits that may come from a greater focus on 'what the market has to say'.

Chapter 12 focuses on portfolio construction techniques. As noted earlier, the application of market consistency to portfolio construction is less direct than for valuation and risk management, but arguably no less important. There are also several analogies that can be drawn

out between market consistency as applied to portfolio construction and market consistency as applied to other financial disciplines.

Chapter 13 draws together many of the strands developed in earlier chapters. It provides case studies exploring how to incorporate market consistency in various types of computations relating to different types of assets and liabilities. It also explores questions like what to do when all available market observable prices relate to relatively illiquid instruments with relatively large or uncertain bid–offer spreads. There would be less need for this book if markets were always ‘perfect’.²

Finally, Chapter 14 summarises and repeats in one place all of the market consistency principles highlighted elsewhere in the book.

Throughout the book we draw out principles, i.e. guidance to *practitioners*, that have relatively universal application, independent of any particular regulatory or current market practice drivers. Within the text these principles are indented and shown in bold, and are referenced by P1, P2, etc.

Each chapter contains at least one such principle, along with many other insights. Thus readers only interested in certain aspects of market consistency should still find this book worthwhile reading, even if they limit their attention just to those chapters particularly relevant to their own specialisms. Conversely, there are valuable insights throughout the book, including in the more mathematical sections (even for readers who don’t wish to follow the mathematics in detail). I would therefore encourage all readers to consider exploring parts of the book that they might have assumed were only tangentially relevant to their own specialisms, because of the greater depth of understanding that this might bring.

1.5 TERMINOLOGY

In parts of the book focused on valuation, we use as essentially interchangeable terms such as market consistency and ‘marking-to-market’. Another term with much the same meaning is the accounting concept of *fair value*. When financial services regulators use the term ‘realistic valuation’ they also normally have a similar concept in mind. We define the market consistent value of an asset or liability to mean:

- (a) Its market value, if it is readily traded on a market at the point in time that the valuation is struck; or
- (b) A reasoned³ best estimate of what its market value would have been had such a market then existed, in all other situations.

² When practitioners talk about markets not being ‘perfect’ they typically have several different concepts in mind simultaneously. These include ‘incompleteness’ (i.e. markets not having as complete a range of instruments as we might like), ‘market frictions’ (i.e. markets being subject to transaction costs, etc.) and ‘inconsistency’ (i.e. the same exposures being priced at a particular point in time inconsistently in different parts of the market). There is sometimes a tendency to equate ‘inconsistency’ with ‘irrationality’, but as we will see in Section 2.10, markets do not need to be behaving rationally for market consistent principles to be applicable. This is fortunate since most practitioners seem to believe that markets do at times exhibit irrational behaviour.

³ By ‘reasoned’ we here mean that the valuer has carefully thought through the sorts of principles set out later in this book, rather than (at the other extreme) merely plucking a number out of thin air. It would have been nice in this book to have been able to define a prescriptive approach that practitioners could follow in all circumstances to come up with a single ‘right’ answer. However, this just isn’t possible in practice. As we shall see, market consistency, when applied to less liquid assets and liabilities, inevitably involves some subjectivity. The best that we can hope for is to place appropriate limits or constraints on this subjectivity and on the mindsets adopted by those who necessarily have to exercise this subjectivity.

Such a definition is similar to the more standard accounting definition of fair value as ‘the value at which an arm’s-length transaction involving willing, knowledgeable counterparties would take place’. However, explicit inclusion of the word ‘market’ within the terminology has the advantage of highlighting that for a non-traded asset or liability we are not wanting to focus on the valuer’s *own* intrinsic assessment of its value. Rather we are interested in modelling how some actual or hypothetical market would be expected to value the asset or liability. By implication, we also demand that any such model should, if possible, be calibrated back to market prices of instruments that are more readily traded.⁴

More generally, choosing to use the term ‘market consistency’ has the advantage of not unwittingly guiding the reader towards an overly valuation centric focus, thus downplaying other disciplines to which market consistency may be applicable. Focusing too much on terms like marking-to-market (or indeed any other phrase involving the word ‘mark’) runs this risk because a ‘mark’ is typically associated with the price we place on an instrument in our books. As we have already noted (and we will stress subsequently), we do not value things in isolation. Valuations ultimately gain their wider meaning and context from the purposes to which they are put.

Given that this book covers risk management, it is also appropriate to include in this section terminology to help categorise the main sorts of risks faced by financial services entities. We shall explore in more detail later what we (and others) might mean by ‘risk’. However, at this juncture, a helpful subdivision often used in practice is the following. It should be noted that it is not always easy (or even useful) to identify clear boundaries between some of these different types of risk.

- (a) *Market risk*, i.e. the risk of loss due to adverse market movements. More generally, we might focus on adverse market movements affecting the entity’s asset/liability position, although this might be called *asset-liability risk*. Market movements in this context would typically include movements in equity values and in interest rates.
- (b) *Credit risk*, i.e. the risk that the creditworthiness of a name or counterparty to which an entity is exposed declines, causing the entity loss. At one extreme would be actual default of the counterparty. A subtlety here is whether credit risk should be deemed to include only *default risk* (i.e. some intrinsic assessment now of the risk that the counterparty or issuer might default in the future) or whether it should also include *ratings migration risk* or *spread risk*. The *spread* on a bond-like or cash-like instrument is the difference in the redemption yield available on the instrument versus the corresponding yield available on some standard reference instrument. For example, people refer to ‘spread’ versus government bond yields, as the difference between the yield on the instrument in question and the yield on government bonds of equivalent duration, type and currency.⁵ However, spread could be measured versus Libor, see Section 5.2, or some other interest rate or yield measure, etc. The market price of a bond subject to default risk is influenced by likelihood of future default. We might attempt to proxy this by some statistic based on the credit rating that a credit rating agency or an internal credit ratings team ascribes to the instrument. However, the *market* price, and hence spread, will also be influenced by

⁴ Another difference is that market consistency does not need to be limited to the special case where both sides are ‘willing’ participants in the trade. In some situations of practical importance (e.g. assessment of the amount of capital needed to withstand stressed conditions), such an assumption may be inappropriate.

⁵ Yields (and hence spreads) can be quoted in a variety of ways, e.g. annualised, semi-annualised, etc. These are generically called ‘annualisation conventions’. For further details see e.g. Kemp (2009).

the market's expectation of how likelihood of default might change over time. Even if one ignores market prices and focuses on some perceived 'intrinsic' likelihood of default derived from credit ratings, these can also change through time (even if the instrument has not defaulted). A rating ascribed to a particular instrument can migrate up or down. At issue is whether spread risk is a form of credit risk (i.e. defining credit risk as risks associated with 'credit' instruments) or whether it is a form of market risk (i.e. defining market risk as anything relating to movements in market prices whatever the instrument type).

- (c) *Liquidity risk*, which the UK's Financial Services Authority (FSA) defines as the risk that a firm, although balance sheet-solvent, cannot maintain or generate sufficient cash resources to meet its payment obligations in full as they fall due, or can only do so at materially disadvantageous terms, see FSA (2007). Some view a part of the spread payable on non-default-free instruments as relating to their liquidity characteristics, again highlighting the difficulties in rigidly demarcating between different types of risk.
- (d) *Insurance risk*, i.e. risks specific to insurance companies, typically relating to the uncertain outcome of insurance contingencies. These would typically include life contingencies, i.e. risks linked to mortality, morbidity or longevity.⁶ They would also include property/casualty and other sorts of non-life insurance risk. Non-life insurance is called 'general' insurance in some jurisdictions. It is not always easy to differentiate what risks are 'insurance-related' and what are not, other than by falling back onto the practical but partly circular definition that insurance risks are ones that are carried by insurance companies.
- (e) *Operational risk*, i.e. the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events. In the six-way classification of risk described in this section a wide range of risks would be deemed to fall into this category, including *legal risk* and, possibly, *strategic risk* and *reputational risk*.⁷
- (f) *Group risk*, i.e. the additional risk to a particular legal entity caused by it being within a larger group structure. For example, resources may be diverted from the entity in question to other group companies if the latter companies suffer a large loss, which can have adverse knock-on effects which would not have arisen had the entity been stand-alone.

For convenience, we also follow the convention, adopted by many other writers in this field, of using the term *firm* to encompass not just bodies with an explicit corporate form and purpose but also other entities that operate within the financial services arena, such as pension funds. Where the context demands, we clarify the specific type of 'firm' or 'entity' on which our attention is focused.

⁶ 'Mortality' and 'longevity' risk are different facets of a single underlying risk, namely, that of uncertainty of when people might die. Mortality risk typically refers to the situation where the insurer is at risk if the person dies sooner than expected whilst longevity risk typically refers to the situation where the insurer is at risk if the person lives longer than expected (e.g. an annuity continues in payment longer than expected). Morbidity refers to ill-health rather than death.

⁷ Strategic risk and reputational risk are excluded from the definition of Operational Risk in Basel II.

