Cambridge University Press
978-1-107-16585-4 — Bond Pricing and Yield Curve Modeling
Riccardo Rebonato
Excerpt
More Information

PART I

THE FOUNDATIONS

CHAPTER 1

What This Book Is About

> It is not my intention to detain the reader by expatiating on the variety, or the importance of the subject, which I have undertaken to treat; since the merit of the choice would serve to render the weakness of the execution still more apparent, and still less excusable. But […] it will perhaps be expected that I should explain, in a few words, the nature and limits of my general plan.
>
> Edward Gibbon, *The Decline and Fall of the Roman Empire*[1]

1.1 MY GOAL IN WRITING THIS BOOK

In this book I intend to look at yield-curve modelling from a 'structural' perspective.[2] I use the adjective *structural* in a very specific sense, to refer to those models which are created with the goal of *explaining* (as opposed to *describing*) the yield curve. What does 'explaining' mean? In the context of this book, I mean accounting for the observed yields by combining the expectations investors form about future rates (and, more generally, the economy) and the compensation they require to bear the risk inherent with holding default-free bonds. (As we shall see later, there is a third 'building block', ie, convexity.)

This provides one level of explanation, but one could go deeper. So, for instance, the degree of compensation investors require in order to bear 'interest-rate risk' could be derived ('explained') in more fundamental terms from the strategy undertaken by a rational, risk-averse investor who is faced with a set of investment opportunities and wants to maximize her utility from consumption

[1] From the Prologue.

[2] A note on terminology. In the term-structure literature the adjective 'structural' is often applied to those models that are based on a specification of the economy – a specification that may go all the way down to preferences, utility maximization and equilibrium. I use the term 'equilibrium models' to refer to these descriptions. We shall only dip our toes in these topics in Chapter 15. For those readers who already understand the meaning of the expression, structural models in this book are those that straddle the $\mathbb{P}$- (real-world) and $\mathbb{Q}$- (risk-neutral) measures. If this does not make much sense at the moment, all will be revealed.

in a multiperiod economy. I will sketch with a broad brush the main lines of this fundamental derivation, but will not pursue this line of argument in great detail. The compensation exacted by investors for bearing market risk (the 'market price of risk') will instead be empirically related (say, via regressions) either to combinations of past and present bond prices and yields, or to past history and present values of macroeconomic variables.

Another way to look at what I try to do in this book is to say that I *describe* the market price of risk in order to *explain* the yield curve. If one took a more 'fundamental' approach, one could try to *explain* the market price of risk as well, but would still have to *describe* something more basic, say, the utility function. Sooner or later, all scientific treatments hit against this hard descriptive core; even theoretical physics is not immune to the curse, or blessing, of having to describe. See, in this respect, the Section 7 of this chapter.

In keeping with the quote that opens this chapter, I will not dwell on why yield curve modelling is important – after all, if the reader were not convinced of this, she probably would not be reading these words. Still, one may well ask, 'Why write a book on *structural* yield-curve modelling?' The answer is that since the mid-2000s there have been exciting developments in the theoretical and empirical understanding of the yield curve dynamics and of risk premia. The 'old' picture with which many of us grew up is now recognized to be in important respects qualitatively wrong. To go from the old to the new class of models requires a rather substantial piece of surgery, not a face-lift, but it is well worth the effort.

Unfortunately, the existing literature on these exciting new topics is somewhat specialized and uses elegant but, to the uninitiated, rather opaque and forbidding-sounding concepts (such as the state-price deflator or the stochastic discount factor). Gone is the simplicity with which even a relative newcomer could pick up Vasicek's paper and, with a good afternoon's work, understand what it was about.

It is therefore my intention to 'translate' and introduce these exciting new developments using the simplest mathematical tools that allow me to handle correctly (but not rigorously) the material at hand. In doing so, I will always trade off a pound of mathematical rigour for an ounce of intuition.

I will also try to explain the vocabulary of the 'new language', and rederive in the simplest possible way the old (and probably familiar) no-arbitrage results using the modern tools. This will both deepen the reader's understanding and enable her to read the current literature.

In addition to expectations and risk premia, there is a third important determinant to the shape of the yield curve, namely 'convexity'. In Part V explain in detail what convexity is, and why it is, in some sense, unique. (In a nutshell, to extract the risk premium you just have to be patient and will be 'fairly' rewarded for your patience; to earn convexity, you have to work very, very hard.) For the moment, the important point is that in the treatment I present in this book these three building blocks (expectations, risk premia, and convexity), together with the principle of no-arbitrage, explain all that there is to know about the yield curve.[3]

1.2 WHAT MY ACCOUNT LEAVES OUT

Is it true that, once we account for expectations, risk premia and convexity, there is really nothing else to the dynamics of credit-risk-free yield curves, at least at the level of description that we have chosen? Of course it isn't. To understand what is left out some historical perspective may help.

The current modelling approach places the Expectation Hypothesis at its centre. This does not mean that 'only expectations matter', but that the only (or the main) deviations from expectations come from risk premia (and the neglected relation, convexity). As Fontaine and Garcia (2015) state '[w]hat distinguishes modern literature is the emphasis on interest rate risk as the leading (or sole) determinant of the risk premium.'[4] As a result 'sources of risk premium other than interest rate risk found a refuge in undergraduate textbooks while the academic agenda leapt forward, developing an array of sophisticated yet tractable no-arbitrage models.'[5]

So what is left behind by the expectations–risk premia–convexity triad?

To begin with, I devote little attention to liquidity, which can become very important, especially in periods of market distress.[6] However, in most market conditions the securities I deal with in this book – US Treasury bonds, German Bunds, UK gilts – are among the most liquid instruments available to investors. Liquidity, one can therefore argue, should be relatively unimportant in a reasonable hierarchy of important factors.[7] If the reader is interested in liquidity-specific issues (such as the pricing of on-the-run versus off-the-run Treasury bonds), the approach of Fontaine and Garcia (2008) discussed in some detail

[3] As noted earlier, I will mention briefly the links between my building blocks and more fundamental macroeconomic and monetary-economics concepts (see Chapters 3 and 15), but I will do so simply to give the reader a qualitative understanding of the form a more fundamental approach to yield curve modelling would take.

[4] p. 463. [5] ibid., pp. 463–464.

[6] In Chapter 18 I present a general pricing methodology that will allow the reader to build her own affine model, DIY-style. Using this toolkit, there is nothing to stop the reader from introducing a factor called 'liquidity', equip it with the necessary parameter paraphernalia (reversion speed, reversion level, volatility, etc) and plug it in the multipurpose affine framework that I develop in Chapter 18. By construction, her 'fits' will be at least as good, and probably better, than before she introduced the 'liquidity' factor. However, it is not easy to find a 'principled' way to assign the correct explanatory contribution to this factor: are we really modelling liquidity, or have we just over-parametrized our model?

[7] Of the models that we explore in Part VII, two deal with liquidity. One is the D'Amico, Kim and Wei (2010) approach, which deals with nominal and *real* rates, explicitly models liquidity – and the authors make the point that the inclusion of this factor is important in order to have a correct estimation of the model parameters and a convincing description of inflation expectations. Dollar-denominated inflation-linked bonds were, especially in the early years after their introduction, far less liquid that their nominal Treasury counterparts, and a strong case can therefore be made for an explicit modelling of liquidity.

Cambridge University Press
978-1-107-16585-4 — Bond Pricing and Yield Curve Modeling
Riccardo Rebonato
Excerpt
More Information

in Chapter 32, is very useful.[8] When it comes to government bonds, however, it must be kept in mind that a bond-specific maturity factor presents a serious challenge for traditional (and frictionless) no-arbitrage models, which are built on the assumption that all bonds are created exactly equal, once their return and risk characteristics are properly taken into account.[9]

The other main possible missing ingredient from the description presented in this book is market segmentation – the idea that classes of investors, such as pension funds, might have preferred 'habitats' (maturity ranges) where they 'like to' invest. According to proponents of segmentation, by so doing, these investors create an imbalance of supply and demand that arbitrageurs either do not manage to eliminate, or do eliminate, but by taking risk, for which compensation – and hence risk premium – is exacted.[10] According to researchers such as Vayanos and Vila (2009), the compensation for the risky activities of pseudo-arbitrageurs then leaves a detectable signature in the shape of a risk-premiuma contribution to various yields. Readers interested in the topic of segmentation are referred to Vayanos and Vila (2009) for a theoretical treatment along these lines, and Chen et al. (2014) for an empirical discussion of the maturity preference exhibited by insurance firms.

These topics, and other sources of imperfections such as the zero bound of rates, are well treated in Fontaine and Garcia (2015) – the title of their chapter ('Recent Advances in Old Fixed Income Topics: Liquidity, Learning, and the Lower Bound') gives a good flavour of what the reader can find in their work. As mentioned, we look at liquidity in Chapter 32, and we deal with the zero bound in Chapter 19. We do not deal with market segmentation, and only cursorily with learning-related issues; see, however, the opening sections of Chapter 28.

1.3 AFFINE MODELS

Let's therefore assume that we are happy with our identification of the three building blocks (expectations, risk premia and convexity) and of the glue

I also deal with liquidity in Chapter 32, which is devoted to the Diebold and Rudebusch approach. The treatment is based on the insight by Fontaine and Garcia (2008), and can be applied to other liquidity-unaware models as well.

[8] 'On-the-run' bonds are freshly-minted, newly-issued Treasury bonds. They enjoy special liquidity, and therefore yield several basis points less (are more expensive) than earlier-issued ('off-the-run') Treasury bonds of similar maturity. This on-the-run/off-the-run spread can become significantly larger in periods of market distress, when liquidity becomes very sought after.

[9] As Fontaine and Garcia (2008) write, 'a structural specification of the liquidity premium raises important challenges. The on-the, run-premium is a real arbitrage opportunity unless we explicitly consider the cost of shorting the more expensive bond, or, alternatively, the benefits accruing to a bondholder from a lower repo rate. These features are absent from the current crop of term-structure model' (pp. 9–10).

[10] As Fontaine and Garcia (2015) point out, liquidity and segmentation need not be looked at as totally different sources of friction or inefficiency because '[t]he clientele demand for new and old bonds is similar in spirit to the view that investors have "preferred habitats"' and '[t]he clientele demand may be scattered across bond maturities, but it can also be scattered across the illiquidity spectrum' (p. 472).

(noarbitrage) that holds them together. What we need next is a way to combine these ingredients in a coherent and logically consistent manner. This is what a model does, and this is why a large part of this book is devoted to discussing models of the yield curve. *Which* models, though?

Because of their unsurpassed intuitional appeal and their analytical tractability, I deal mainly with a popular class of structural models – the affine class.[11] In order to give a transparent understanding of how these models weave together these three building blocks to determine the shape of the yield curve, I will start my discussion from the simplest incarnation of affine models – the Vasicek (1977) model.[12]

The Vasicek model is unparalleled for the intuitive understanding it affords, and it is for this reason that I introduce it, perhaps unwisely, very early in the book – even, that is, before dealing with the theoretical underpinnings of term-structure modelling. Quite simply, I want the reader to have a vivid, if, at this point, probably imprecise, picture of what we will be talking about more precisely and more abstractly in the later parts of the book, when more complex, and more opaque, models come to the fore.

In general, I strongly encourage the reader who feels her intuition beginning to fail her when looking at the more complex models to adopt ruthlessly the strategy of *reductio ad Vasicek*, ie, to ask herself, 'What is the equivalent of this concept/formalism/result in the Vasicek model?' She is encouraged to do so, not because the Vasicek model is perfect, but because it lays bare with great clarity the mechanics and intuition behind more complex affine models.

For all the virtues of the Vasicek model, recent empirical evidence suggests that the explanation of risk premia Vasicek-family models afford is *qualitatively* wrong. Since the risk premium constitutes the explanatory bridge between expectations and observed prices, and since the Vasicek approach is the progenitor of all the more recent affine models, this does not seem to bode well for affine structural approaches to term-structure modelling.

Luckily, the same empirical evidence also suggests how the first-generation, Vasicek-like, affine models can be modified and enriched. I therefore present in Part VI of this book what we now know about term premia, and in Part VII how these empirical findings can be incorporated in the new-generation affine models.

[11] See, for instance, Dai and Singleton (2000) for a systematic calssification of affine models, and Duffee (2002) for a discussion of *essentially* affine models – loosely speaking, models which remain affine both in the real-world and in the pricing measures. Good reviews of affine models can be found in Bloder (2001), who also deals with Kalman filter estimation methods, and Piazzesi (2010). Extensions to stochastic affine-volatility models are found in Longstaff and Schwartz (1992) and Balduzzi et al. (1996).

[12] I must make very clear from the start that I will deal in this book with *Gaussian* affine models, which are far simpler than the square-root models of the Cox–Ingersoll–Ross (1985a, b) family. Admittedly, Gaussian affine models do allow for negative rates, but recent experience suggests that this should be considered more of a virtue than a blemish. (At the time of this writing, Germany just issued short-dated government bonds with a negative yield.)

Speaking of affine models means that we require a special type of relationship between yields and the state variables. But how should we choose these variables? As we shall discuss towards the end of the book, from a very abstract point of view, and as long as some quantities are exactly recovered by the different models, the choice of variables makes very little difference. In practice, however, this choice informs the statistical estimation techniques used in the calibration, the degree of 'structure' on the dynamics of the state variables (via the condition of no-arbitrage), the parsimony of the model and the user's ability to understand and interpret the model. Section 1.5 of this introductory chapter makes these statements more precise. First, however, we want to look a bit more carefully at the various types of yield curve models, so that the reader can clearly see what we are going to deal with and what we will not touch upon. Probably, the reader should not throw away her book receipt before reading the next section.

1.4 A SIMPLE TAXONOMY

There are many different types of term-structure models. They are different in part because they have been created with different purposes in mind and in part because they look at the same problem from different angles. A reasonable taxonomy may look as follows.

1. *Statistical models* aim to *describe* how the yield curve moves. Their main workhorses here are the Vector Auto-Regressive (VAR) models, which are often employed to forecast interest rates and to estimate the risk premium as the difference between the forward and the forecasted rates. This task sounds easy, but, as I discuss later in the book, the quasi-unit-root nature of the level of rates (and many more statistical pitfalls) makes estimations based purely on time-series analysis arduous, and the associated 'error bars' embarrassingly large. See, eg, the discussion in Cochrane and Piazzesi (2008).[13]

   In the attempt to improve on this state of affairs, no-arbitrage structural models, which add *cross-sectional* information to the time-series data, come to the fore. *In this book we shall take a cursory and instrumental look at statistical models, mainly to glean statistical information about one important ingredient of our structural models, ie, the market price of risk.*

   The important thing to stress is that statistical models fit observed market yield curves well and have good predictive power but lack a strong theoretical foundation, because, by themselves, they cannot guarantee absence of arbitrage among the predicted yields. Their strengths and weaknesses are therefore complementary to those of the no-arbitrage models discussed in the text that follows: these are theoretically sound, but sometimes poor at fitting the market yield

[13] See, in particular, the discussion of their Panel 1 on p. 2 of their paper.

covariance structure and the observed yield curves, and worse at predicting their evolution. See, in this respect, the discussion in Diebold and Rudebusch (2013)[14] and Section 1 in Chapter 32.

One of the underlying themes developed in this book is the attempt to marry the predictive and fitting virtues of statistical models with the theoretical solidity of the no-arbitrage models. Chapters 32, 33 and 34 should be read in this light.

2. *Structural no-arbitrage models* (of which the Vasicek (1977) and Cox–Ingersoll–Ross (1985a, b) are the earliest and best-known textbook examples) make assumptions about how a handful of important driving factors behave; they ensure that the no-arbitrage condition is satisfied; and they derive how the three components that drive the yield curve (expectations, risk premia and convexity) should affect the shape of the yield curve. The no-arbitrage conditions ensure that the derived prices of bonds do not offer free lunches. As I explain in footnote 1, I speak of structural no-arbitrage models when they straddle the physical (real-world, $\mathbb{P}$) and risk-neutral ($\mathbb{Q}$) measures – as opposed to restricted no-arbitrage models that are formulated only in the $\mathbb{Q}$ measure.

   The distinction is important for at least two reasons. First, if we want to understand how bond prices are formed based on expectations and risk aversion, we cannot look at just one measure: market prices are compatible with an infinity of different combinations of expectations and market prices of risk.

   The second reason is subtler. It is well known that if we only look at the risk-neutral ($\mathbb{Q}$) measure three factors (as we shall see, the first three principal components) explain the movements in prices extremely well. However, if we also want to explain excess returns (risk premia) we may have to use more variables (perhaps up to five, according to Cochrane and Piazzesi (2005, 2008), Adrian, Crump and Moench (2013) and Hellerstein (2011)).[15] The message here is that variables virtually irrelevant in one measure may become important when the two measures are linked. More about this later. *Structural no-arbitrage models constitute the class of models this book is about*.

3. *'Snapshot' models* (such as the Nelson–Siegel (1987) model, or the many splines models of which Fisher, Nychka and Zervos's (1995) is probably the best known) are cross-sectional devices to *interpolate* prices or yields of bonds that we cannot observe, given a set of prices or yields that we *can* observe.[16] They also produce as a by-product the model yields of the bonds we *do* observe. If supplemented with

[14] p. 76.

[15] See in this respect the discussion on p. 140 of Cochrane and Piazzesi (2005) and on p. 3 of Hellerstein (2011).

[16] For two early, but still valid, evaluations of yield-curve estimation models, see Bliss (1997) and Anderson et al. (1996).

the ubiquitous but somewhat ad hoc assumption that the residuals (the differences between the model and the market prices) are mean reverting, these models give practitioners suggestions about whether a given observed bond yield (hence, price) is 'out of line' with a reasonable smooth interpolation of where it should lie.[17] Liquidity corrections such as those discussed in Fontaine and Garcia (2008) can be very important in these 'cheap/dear' analyses.

Apart from the smoothness-based assessment of the relative cheapness or dearness of different bonds, snapshot models are extremely important for structural affine models because they assume the existence of a continuum of discount bonds. So the output of snapshot models (a snapshot discount function) is the input to structural models.

In general, there is no deep meaning to the parameters of fitted snapshot models. However, some recent developments have given a time-series, dynamic interpretation to their parameters, and married them with Principal Component Analysis. (See, eg, (Diebold and Rudebusch, 2013).) So, these latest developments combine features of structural, statistical and snapshot models. We shall revisit this approach later in the chapter.

4. *Derivatives models* (eg, the Heat–Jarrow–Morton (1992), the Brace–Gatarek–Musiela (1997), the Hull and White (1990), the Black–Derman–Toy (1990), …) are based on *relative* pricing and on the enforcement of no-arbitrage. Because of this, they strongly rely on first-order cancellation of errors (between the derivative they are designed to price and the hedging instruments used to build the riskless or minimum-variance portfolio; see the discussion in Nawalha and Rebonato (2011)). Therefore they do not strive to provide a particularly realistic description of the underlying economic reality. After the first generation (Vasicek (1977), Cox et al. (1985a,b), derivatives models squarely set up camp in the risk-neutral $\mathbb{Q}$ measure, and affect a disdainful lack of interest for risk premia. I do not deal with this class of models in this book.

1.5 THE CHOICE OF VARIABLES*

1.5.1 Latent versus Observable Variables

As mentioned previously, an important theme that recurs throughout the book is that the choice of the type of state variable is a very important, and often

[17] Snapshot models are also important because all structural models use as their building blocks discount bonds, which are not traded in the market but which make mathematical analysis (immensely) easier. The output of snapshot models (the discount curve) is therefore the input to structural models.