# 1

# Facts, Factors, and Questions

In this chapter we introduce some important conceptual, descriptive, and theoretical considerations regarding nominal government bond yield curves. Conceptually, just what is it that are we trying to measure? How can we best understand many bond yields at many maturities over many years? Descriptively, how do yield curves tend to behave? Can we obtain simple yet accurate dynamic characterizations and forecasts? Theoretically, what governs and restricts yield curve shape and evolution? Can we relate yield curves to macroeconomic fundamentals and central bank behavior?

These multifaceted questions are difficult yet very important. Accordingly, a huge and similarly multifaceted literature attempts to address them. Numerous currents and cross-currents, statistical and economic, flow through the literature. There is no simple linear thought progression, self-contained with each step following logically from that before. Instead the literature is more of a tangled web; hence our intention is not to produce a "balanced" survey of yield curve modeling, as it is not clear whether that would be helpful or even what it would mean. On the contrary, in this book we slice through the literature in a calculated way, assembling and elaborating on a very particular approach to yield curve modeling. Our approach is simple yet rigorous, simultaneously in close touch with modern statistical and financial economic thinking, and effective in a variety of situations. But we are getting ahead of ourselves. First we must lay the groundwork.

## 1.1 Three Interest Rate Curves

Here we fix ideas, establish notation, and elaborate on key concepts by recalling three key theoretical bond market constructs and the relationships among them: the discount curve, the forward rate curve, and the yield curve. Let $P(\tau)$ denote the price of a $\tau$-period discount bond, that is, the present value of \$1 receivable $\tau$ periods ahead. If $y(\tau)$ is its continuously compounded yield to maturity, then by definition

$$P(\tau) = e^{-\tau y(\tau)}. \tag{1.1}$$

Hence the discount curve and yield curve are immediately and fundamentally related. Knowledge of the discount function lets one calculate the yield curve.

The discount curve and the forward rate curve are similarly fundamentally related. In particular, the forward rate curve is defined as

$$f(\tau) = \frac{-P'(\tau)}{P(\tau)}. \tag{1.2}$$

Thus, just as knowledge of the discount function lets one calculate the yield curve, so too does knowledge of the discount function let one calculate the forward rate curve.

Equations (1.1) and (1.2) then imply a relationship between the yield curve and forward rate curve,

$$y(\tau) = \frac{1}{\tau}\int_0^\tau f(u)du. \tag{1.3}$$

In particular, the zero-coupon yield is an equally weighted average of forward rates.

The upshot for our purposes is that, because knowledge of any one of $P(\tau)$, $y(\tau)$, and $f(\tau)$ implies knowledge of the other two, the three are effectively interchangeable. Hence with no loss of generality one can choose to work with $P(\tau)$, $y(\tau)$, or $f(\tau)$. In this book, following much of both academic and industry practice, we work with the yield curve, $y(\tau)$. But again, the choice is inconsequential in theory.

Complications arise in practice, however, because although we observe prices of traded bonds with various amounts of time to maturity, we do not directly observe yields, let alone the zero-coupon yields at fixed standardized maturities (e.g., six-month, ten-year, ...), with which we work throughout. Hence we now provide some background on yield construction.

## 1.2 Zero-Coupon Yields

In practice, yield curves are not observed. Instead, they must be estimated from observed bond prices. Two historically popular approaches to constructing yields proceed by fitting a smooth discount curve and then converting to yields at the relevant maturities using formulas (1.2) and (1.3).

The first discount curve approach to yield curve construction is due to McCulloch (1971, 1975), who models the discount curve using polynomial splines.[1] The fitted discount curve, however, diverges at long maturities due to the polynomial structure, and the corresponding yield curve inherits that unfortunate property. Hence

---

[1] See also McCulloch and Kwon (1993).

such curves can provide poor fits to yields that flatten out with maturity, as emphasized by Shea (1984).

An improved discount curve approach to yield curve construction is due to Vasicek and Fong (1982), who model the discount curve using exponential splines. Their clever use of a negative transformation of maturity, rather than maturity itself, ensures that forward rates and zero-coupon yields converge to a fixed limit as maturity increases. Hence the Vasicek-Fong approach may be more successful at fitting yield curves with flat long ends.

Notwithstanding the progress of Vasicek and Fong (1982), discount curve approaches remain potentially problematic, as the implied forward rates are not necessarily positive. An alternative and popular approach to yield construction is due to Fama and Bliss (1987), who construct yields not from an estimated discount curve, but rather from estimated forward rates at the observed maturities. Their method sequentially constructs the forward rates necessary to price successively longer-maturity bonds. Those forward rates are often called "unsmoothed Fama-Bliss" forward rates, and they are transformed to unsmoothed Fama-Bliss yields by appropriate averaging, using formula (1.3). The unsmoothed Fama-Bliss yields exactly price the included bonds. Unsmoothed Fama-Bliss yields are often the "raw" yields to which researchers fit empirical yield curves, such as members of the Nelson-Siegel family, about which we have much to say throughout this book. Such fitting effectively smooths the unsmoothed Fama-Bliss yields.

## 1.3 Yield Curve Facts

At any time, dozens of different yields may be observed, corresponding to different bond maturities. But yield curves evolve dynamically; hence they have not only a cross-sectional, but also a temporal, dimension.[2] In this section we address the obvious descriptive question: How do yields tend to behave across different maturities and over time?

The situation at hand is in a sense very simple—modeling and forecasting a time series—but in another sense rather more complex and interesting, as the series to be modeled is in fact a series of *curves*.[3] In Figure 1.1 we show the resulting three-dimensional surface for the United States, with yields shown as a function of maturity, over time. The figure reveals a key yield curve fact: yield curves move a lot, shifting among different shapes: increasing at increasing or decreasing rates, decreasing at increasing or decreasing rates, flat, U-shaped, and so on.

Table 1.1 presents descriptive statistics for yields at various maturities. Several well-known and important yield curve facts emerge. First, time-averaged yields (the "average yield curve") increase with maturity; that is, term premia appear to exist, perhaps due to risk aversion, liquidity preferences, or preferred habitats. Second, yield volatilities decrease with maturity, presumably because long rates involve averages of expected future short rates. Third, yields are highly persistent, as evidenced not only by the very large 1-month autocorrelations but also by the sizable 12-month autocorrelations.

[2] We will be interested in dynamic modeling and forecasting of yield curves, so the temporal dimension is as important as the variation across bond maturity.

[3] The statistical literature on functional regression deals with sets of curves and is therefore somewhat related to our concerns. See, for example, Ramsay and Silverman (2005) and Ramsay et al. (2009). But the functional regression literature typically does not address dynamics, let alone the many special nuances of yield curve modeling. Hence we are led to rather different approaches.

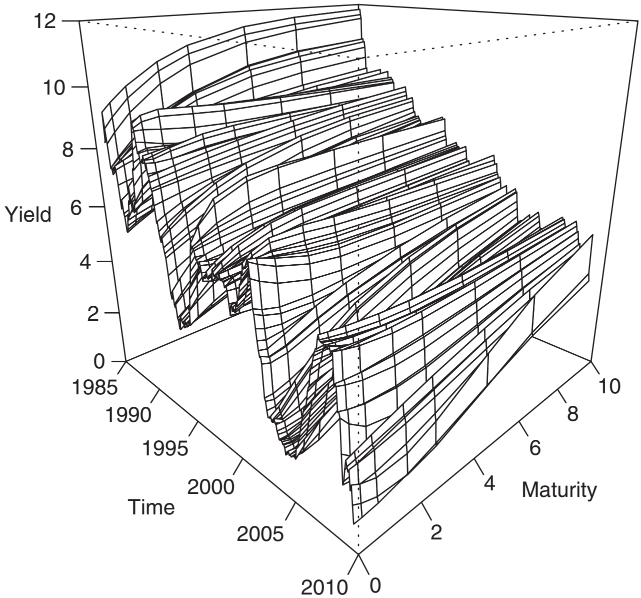


**Figure 1.1.** Bond Yields in Three Dimensions. We plot end-of-month U.S. Treasury bill and bond yields at maturities ranging from 6 months to 10 years. Data are from the Board of Governors of the Federal Reserve System, based on Gürkaynak et al. (2007). The sample period is January 1985 through December 2008.

Table 1.2 shows the same descriptive statistics for yield *spreads* relative to the 10-year bond. Yield spread dynamics contrast rather sharply with those of yield levels; in particular, spreads are noticeably less volatile and less persistent. As with yields, the 1-month spread autocorrelations are very large, but they decay more quickly, so that the 12-month spread autocorrelations are noticeably smaller than those for yields. Indeed many strategies for active bond trading (sometimes successful and sometimes not!) are based on spread reversion.

**Table 1.1.** Bond Yield Statistics

| Maturity (months) | $\bar{y}$ | $\hat{\sigma}_y$ | $\hat{\rho}_y(1)$ | $\hat{\rho}_y(12)$ |
|---|---|---|---|---|
| 6 | 4.9 | 2.1 | 0.98 | 0.64 |
| 12 | 5.1 | 2.1 | 0.98 | 0.65 |
| 24 | 5.3 | 2.1 | 0.97 | 0.65 |
| 36 | 5.6 | 2.0 | 0.97 | 0.65 |
| 60 | 5.9 | 1.9 | 0.97 | 0.66 |
| 120 | 6.5 | 1.8 | 0.97 | 0.68 |

*Notes*: We present descriptive statistics for end-of-month yields at various maturities. We show sample mean, sample standard deviation, and first- and twelfth-order sample autocorrelations. Data are from the Board of Governors of the Federal Reserve System. The sample period is January 1985 through December 2008.

## 1.4 Yield Curve Factors

Multivariate models are required for sets of bond yields. An obvious model is a vector autoregression or some close relative. But unrestricted vector autoregressions are profligate parameterizations, wasteful of degrees of freedom. Fortunately, it turns out that financial asset returns typically conform to a certain type of *restricted* vector autoregression, displaying *factor structure*. Factor structure is said to be operative in situations where one sees a high-dimensional object (e.g., a large set of bond yields), but where that high-dimensional object is driven by an underlying lower-dimensional set of objects, or "factors." Thus beneath a high-dimensional seemingly complicated set of observations lies a much simpler reality.

Indeed factor structure is ubiquitous in financial markets, financial economic theory, macroeconomic funda-

**Table 1.2.** Yield Spread Statistics

| Maturity (months) | $\bar{s}$ | $\hat{\sigma}_s$ | $\hat{\rho}_s(1)$ | $\hat{\rho}_s(12)$ |
|---|---|---|---|---|
| 6 | −1.6 | 1.3 | 0.98 | 0.44 |
| 12 | −1.4 | 1.1 | 0.98 | 0.46 |
| 24 | −1.1 | 0.9 | 0.97 | 0.48 |
| 36 | −0.9 | 0.7 | 0.97 | 0.47 |
| 60 | −0.6 | 0.4 | 0.96 | 0.44 |
| 120 | NA | NA | NA | NA |

*Notes*: We present descriptive statistics for end-of-month yield spreads (relative to the 10-year bond) at various maturities. We show sample mean, sample standard deviation, and first- and twelfth-order sample autocorrelations. Data are from the Board of Governors of the Federal Reserve System, based on Gürkaynak et al. (2007). The sample period is January 1985 through December 2008.

mentals, and macroeconomic theory. Campbell et al. (1997), for example, discuss aspects of empirical factor structure in financial markets and theoretical factor structure in financial economic models.[4] Similarly, Aruoba and Diebold (2010) discuss empirical factor structure in macroeconomic fundamentals, and Diebold and Rudebusch (1996) discuss theoretical factor structure in macroeconomic models.

In particular, factor structure provides a fine description of the term structure of bond yields.[5] Most early studies involving mostly long rates implicitly adopt a single-factor world view (e.g., Macaulay (1938)), where the factor is the level (e.g., a long rate). Similarly, early

[4] Interestingly, asset pricing in the factor framework is closely related to asset pricing in the pricing kernel framework, as discussed in Chapter 11 of Singleton (2006).

[5] For now we do not distinguish between government and corporate bond yields. We will consider credit risk spreads later.

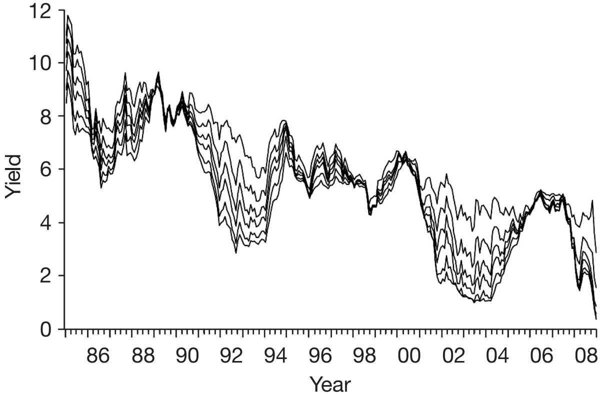


**Figure 1.2.** Bond Yields in Two Dimensions. We plot end-of-month U.S. Treasury bill and bond yields at maturities of 6, 12, 24, 36, 60, and 120 months. Data are from the Board of Governors of the Federal Reserve System, based on Gürkaynak et al. (2007). The sample period is January 1985 through December 2008.

arbitrage-free models like Vasicek (1977) involve only a single factor. But single-factor structure severely limits the scope for interesting term structure dynamics, which rings hollow in terms of both introspection and observation.

In Figure 1.2 we show a time-series plot of a standard set of bond yields. Clearly they *do* tend to move noticeably together, but at the same time, it's clear that more than just a common level factor is operative. In the real world, term structure data—and, correspondingly, modern empirical term structure models—involve *multiple* factors. This classic recognition traces to Litterman and Scheinkman (1991), Willner (1996), and Bliss (1997), and it is echoed repeatedly in the literature. Joslin et al. (2010), for example, note:

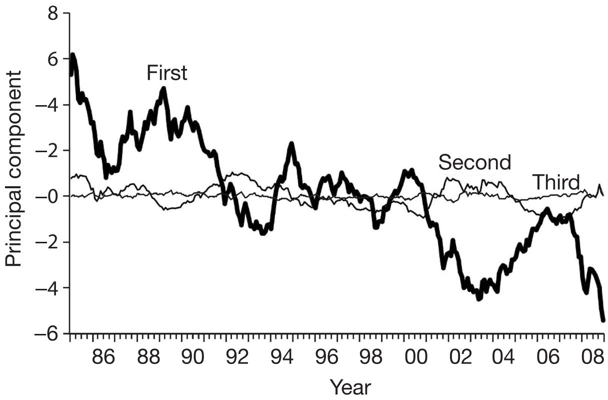


**Figure 1.3.** Bond Yield Principal Components. We show the first, second, and third principal components of bond yields in dark, medium, and light shading, respectively.

> The cross-correlations of bond yields are well decribed by a low-dimensional factor model in the sense that the first three principal components of bond yields . . . explain well over 95 percent of their variation. . . . Very similar three-factor representations emerge from arbitrage-free, dynamic term structure models . . . for a wide range of maturities.

Typically three factors, or principal components, are all that one needs to explain most yield variation. In our data set the first three principal components explain almost 100 percent of the variation in bond yields; we show them in Figure 1.3 and provide descriptive statistics in Table 1.3.

The first factor is borderline nonstationary. It drifts downward over much of the sample period, as inflation was reduced relative to its high level in the early 1980s. The first factor is the most variable but also the most predictable, due to its very high persistence. The second factor is also highly persistent and displays a clear

**Table 1.3.** Yield Principal Components Statistics

| PC | $\hat{\sigma}$ | $\hat{\rho}(1)$ | $\hat{\rho}(12)$ | $R^2$ |
|---|---|---|---|---|
| First | 2.35 | 0.97 | 0.67 | 0.98 |
| Second | 0.52 | 0.97 | 0.45 | 0.95 |
| Third | 0.10 | 0.83 | 0.15 | 0.70 |

*Notes*: We present descriptive statistics for the first three principal components of end-of-month U.S. government bill and bond yields at maturities of 6, 12, 24, 36, 60, and 120 months. We show principal component sample standard deviation, first- and twelfth-order principal component sample autocorrelations, and the predictive $R^2$ (see Diebold and Kilian (2001)) from an $AR(p)$ approximating model with $p$ selected using the Schwartz criterion. Data are from the Board of Governors of the Federal Reserve System, based on Gürkaynak et al. (2007). The sample period is January 1985 through December 2008.

business cycle rhythm. The second factor is less variable, less persistent, and less predictable than the level factor. The third factor is the least variable, least persistent, and least predictable.

In Figure 1.4 we plot the three principal components (factors) against standard empirical yield curve level, slope, and curvature measures (the 10-year yield, the 10Y-6M spread, and a 6M+10Y-2*5Y butterfly spread, respectively). The figure reveals that the three bond yield factors effectively *are* level, slope, and curvature. This is important, because it implies that the different factors likely have different and specific macroeconomic determinants. Inflation, for example, is clearly related to the yield curve level, and the stage of the business cycle is relevant for the slope. It is also noteworthy that the yield factors are effectively orthogonal due to their exceptionally close links to the principal components, which are orthogonal by construction.

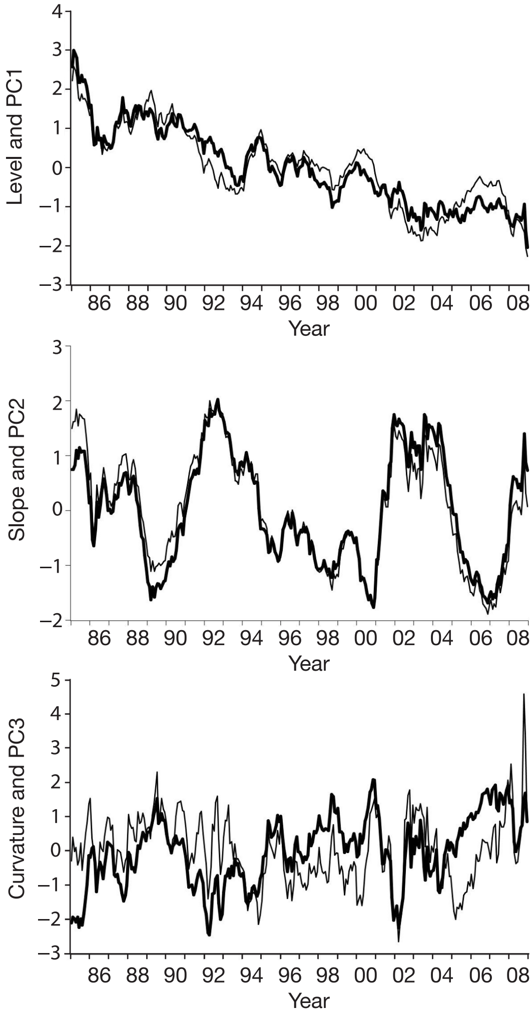


**Figure 1.4.** Empirical Level, Slope, and Curvature and First Three Principal Components of Bond Yields. We show the standardized empirical level, slope, and curvature with dark lines, and the first three standardized principal components (denoted PC1, PC2, and PC3) with lighter lines.

The disproportionate amount of yield variation associated with the common *level* factor, together with its high persistence, explains the broad sweep of earlier-discussed facts, in particular the high persistence of yields and the greatly reduced persistence of yield spreads (because the common level factor vanishes from the spreads). Reality is of course a bit more complicated, as slope and curvature factors are also operative, but the effects of the level factor dominate.

A factor structure for yields with a highly persistent level factor is constrained by economic theory. Economic theory strongly suggests that nominal bond yields should not have unit roots, because the yields are bounded below by zero, whereas unit-root processes have random walk components and therefore will eventually cross zero almost surely. Nevertheless, the unit root may be a good approximation so long as yields are not too close to zero, as noted by Dungey et al. (2000), Giese (2008), and Jardet et al. (2010), among others.[6] Work in that tradition, most notably Dungey et al. (2000), finds not only integration but also clear *co*integration, and the common unit roots associated with cointegration imply factor structure.

## 1.5 Yield Curve Questions

Thus far we have laid the groundwork for subsequent chapters, touching on aspects of yield definition, data construction, and descriptive statistical properties of yields and yield factors. We have emphasized the high

[6] Alternatively, more sophisticated models, such as the "square-root process" of Cox et al. (1985), can allow for unit-root dynamics while still enforcing yield nonnegativity by requiring that the conditional variance of yields approach zero as yields approach zero.

persistence of yields, the lesser persistence of yield spreads, and, related, the good empirical approximation afforded by a low-dimensional three-factor structure with highly persistent level and slope factors. Here we roam more widely, in part looking backward, expanding on themes already introduced, and in part looking forward, foreshadowing additional themes that feature prominently in what follows.

### 1.5.1 Why Use Factor Models for Yields?

The first problem faced in term structure modeling is how to summarize the price information at any point in time for the large number of nominal bonds that are traded. Dynamic factor models prove appealing for three key reasons.

First, as emphasized already, factor structure generally provides a highly accurate empirical description of yield curve data. Because only a small number of systematic risks appear to underlie the pricing of the myriad of tradable financial assets, nearly all bond price information can be summarized with just a few constructed variables or factors. Therefore, yield curve models almost invariably employ a structure that consists of a small set of factors and the associated factor loadings that relate yields of different maturities to those factors.

Second, factor models prove tremendously appealing for statistical reasons. They provide a valuable compression of information, effectively collapsing an intractable high-dimensional modeling situation into a tractable low-dimensional situation. This would be small consolation if the yield data were not well-approximated with factor structure, but again, they are. Hence we're in a most fortunate situation. We *need* low-dimensional factor structure for statistical tractability, and, mercifully, the data actually *have* factor structure.

Related, factor structure is consistent with the "parsimony principle," which we interpret here as the broad insight that imposing restrictions implicitly associated with simple models—even false restrictions that may degrade in-sample fit—often helps to avoid data mining and, related, to produce good out-of-sample forecasts.[7] For example, an unrestricted vector autoregression provides a very general linear model of yields typically with good in-sample fit, but the large number of estimated coefficients may reduce its value for out-of-sample forecasting.[8]

Last, and not at all least, financial economic *theory* suggests, and routinely invokes, factor structure. We see thousands of financial assets in the markets, but for a variety of reasons we view the risk premiums that separate their expected returns as driven by a much smaller number of components, or *risk factors*. In the equity sphere, for example, the celebrated capital asset pricing model (CAPM) is a single-factor model. Various extensions (e.g., Fama and French (1992)) invoke a few additional factors but remain intentionally very low-dimensional, almost always with fewer than five factors. Yield curve factor models are a natural bond market parallel.

### 1.5.2 How Should Bond Yield Factors and Factor Loadings Be Constructed?

The literature contains a variety of methods for constructing bond yield factors and factor loadings. One

[7] See Diebold (2007) for additional discussion.

[8] Parsimony, however, is not the only consideration for determining the number of factors needed; the demands of the precise application are of course also relevant. For example, although just a few factors may account for almost all dynamic yield variation and optimize forecast accuracy, more factors may be needed to fit with great accuracy the cross section of yields at a point in time, say, for pricing derivatives.

approach places structure only on the estimated factors, leaving loadings free. For example, the factors could be the first few principal components, which are restricted to be mutually orthogonal, while the loadings are left unrestricted. Alternatively, the factors could be observed bond portfolios, such as a long-short for slope or a butterfly for curvature.

A second approach, conversely, places structure only on the loadings, leaving factors free. The classic example, which has long been popular among market and central bank practitioners, is the so-called Nelson-Siegel curve, introduced in Nelson and Siegel (1987). As shown by Diebold and Li (2006), a suitably dynamized version of Nelson-Siegel is effectively a dynamic three-factor model of level, slope, and curvature. However, the Nelson-Siegel factors are unobserved, or latent, whereas the associated loadings are restricted by a functional form that imposes smoothness of loadings across maturities, positivity of implied forward rates, and a discount curve that approaches zero with maturity.

A third approach, the no-arbitrage dynamic latent factor model, which is the model of choice in finance, restricts both factors and factor loadings. The most common subclass of such models, affine models in the tradition of Duffie and Kan (1996), postulates linear or affine dynamics for the latent factors and derives the associated restrictions on factor loadings that ensure absence of arbitrage.

### 1.5.3 Is Imposition of "No-Arbitrage" Useful?

The assumption of no-arbitrage ensures that, after accounting for risk, the dynamic evolution of yields over time is consistent with the cross-sectional shape of the yield curve at any point in time. This consistency condition is likely to hold, given the existence of deep and well-organized bond markets. Hence one might argue that the real markets are at least approximately arbitrage-free, so that a good yield curve model must display freedom from arbitrage.

But all models are false, and subtleties arise once the inevitability of model misspecification is acknowledged. Freedom from arbitrage is essentially an internal consistency condition. But a misspecified model may be internally consistent (free from arbitrage) yet have little relationship to the real world, and hence forecast poorly, for example. Moreover, imposition of no-arbitrage on a misspecified model may actually *degrade* empirical performance.

Conversely, a model may admit arbitrage yet provide a good approximation to a much more complicated reality, and hence forecast well. Moreover, if reality is arbitrage-free, and if a model provides a very good description of reality, then imposition of no-arbitrage would presumably have little effect. That is, an accurate model would be at least approximately arbitrage-free, even if freedom from arbitrage were not explicitly imposed.

Simultaneously, a large literature suggests that coaxing or "shrinking" forecasts in various directions (e.g., reflecting prior views) may improve performance, effectively by producing large reductions in error variance at the cost of only small increases in bias. An obvious benchmark shrinkage direction is toward absence of arbitrage. The key point, however, is that shrinkage methods don't force absence of arbitrage; rather, they coax things *toward* absence of arbitrage.

If we are generally interested in the questions posed in this subsection's title, we are also specifically interested in answering them in the dynamic Nelson-Siegel context.

A first question is whether our dynamic Nelson-Siegel (DNS) model can be made free from arbitrage. A second question, assuming that DNS can be made arbitrage-free, is whether the associated restrictions on the physical yield dynamics improve forecasting performance.

### 1.5.4 How Should Term Premiums Be Specified?

With risk-neutral investors, yields are equal to the average value of expected future short rates (up to Jensen's inequality terms), and there are no expected excess returns on bonds. In this setting, the expectations hypothesis, which is still a mainstay of much casual and formal macroeconomic analysis, is valid. However, it seems likely that bonds, which provide an uncertain return, are owned by the same risk-averse investors who also demand a large equity premium as compensation for holding risky stocks. Furthermore, as suggested by many statistical tests in the literature, the risk premiums on nominal bonds appear to vary over time, which suggests time-varying risk, time-varying risk aversion, or both (e.g., Campbell and Shiller (1991), Cochrane and Piazzesi (2005)).[9]

In the finance literature, the two basic approaches to modeling time-varying term premiums are time-varying quantities of risk and time-varying "prices of risk" (which translate a unit of factor volatility into a term premium). The large literature on stochastic volatility takes the former approach, allowing the variability of yield factors to change over time. In contrast, the canonical Gaussian affine no-arbitrage finance representation (e.g., Ang

[9] However, Diebold et al. (2006b) suggest that the importance of the statistical deviations from the expectations hypothesis may depend on the application.

and Piazzesi (2003)) takes the latter approach, specifying time-varying prices of risk.[10]

### 1.5.5 How Are Yield Factors and Macroeconomic Variables Related?

The modeling of interest rates has long been a prime example of the disconnect between the macro and finance literatures. In the canonical finance model, the short-term interest rate is a linear function of a few unobserved factors. Movements in long-term yields are importantly determined by changes in risk premiums, which also depend on those latent factors. In contrast, in the macro literature, the short-term interest rate is set by the central bank according to its macroeconomic stabilization goals—such as reducing deviations of inflation and output from the central bank's targets. Furthermore, the macro literature commonly views long-term yields as largely determined by expectations of future short-term interest rates, which in turn depend on expectations of the macro variables; that is, possible changes in risk premiums are often ignored, and the expectations hypothesis of the term structure is employed.

Surprisingly, the disparate finance and macro modeling strategies have long been maintained, largely in isolation of each other. Of course, differences between the finance and macro perspectives reflect, in part, different questions, methods, and avenues of exploration. However, the lack of interchange or overlap between the two research

[10] Some recent literature takes an intermediate approach. In a structural dynamic stochastic general equilibrium (DSGE) model, Rudebusch and Swanson (2012) show that technology-type shocks can endogenously generate time-varying prices of risk—namely, conditional heteroskedasticity in the stochastic discount factor—without relying on conditional heteroskedasticity in the driving shocks.

literatures that occurred in the past is striking. Notably, both the DNS and affine no-arbitrage dynamic latent factor models provide useful statistical descriptions of the yield curve, but in their original, most basic, forms they offer little insight into the nature of the underlying economic forces that drive its movements.

Hence, to illuminate the fundamental determinants of interest rates, researchers have begun to incorporate macroeconomic variables into the DNS and affine no-arbitrage dynamic latent factor yield curve models. For example, Diebold et al. (2006b) provide a macroeconomic interpretation of the DNS representation by combining it with vector-autoregressive dynamics for the macroeconomy. Their maximum-likelihood estimation approach extracts 3 latent factors (essentially level, slope, and curvature) from a set of 17 yields on U.S. Treasury securities and simultaneously relates these factors to 3 observable macroeconomic variables (specifically, real activity, inflation, and a monetary policy instrument). By examining the correlations between the DNS yield factors and macroeconomic variables, they find that the level factor is highly correlated with inflation and the slope factor is highly correlated with real activity. The curvature factor appears unrelated to any of the main macroeconomic variables.

The role of macroeconomic variables in a no-arbitrage affine model is explored in several papers. In Ang and Piazzesi (2003), the macroeconomic factors are measures of inflation and real activity, and the joint dynamics of macro factors and additional latent factors are captured by vector autoregressions.[11] They find that output

[11] To avoid relying on specific macro series, Ang and Piazzesi construct their measures of real activity and inflation as the first principal component of a large set of candidate macroeconomic series.

shocks have a significant impact on intermediate yields and curvature, while inflation surprises have large effects on the level of the entire yield curve.

For estimation tractability, Ang and Piazzesi allow only for unidirectional dynamics in their arbitrage-free model; specifically, macro variables help determine yields but not the reverse. In contrast, Diebold et al. (2006b) consider a bidirectional characterization of dynamic macro–yield interactions. They find that the causality from the macroeconomy to yields is indeed significantly stronger than in the reverse direction, but that interactions in both directions can be important. Ang et al. (2007) also allow for bidirectional macro-finance links but impose the no-arbitrage restriction as well, which poses a severe estimation challenge. They find that the amount of yield variation that can be attributed to macro factors depends on whether the system allows for bidirectional linkages. When the interactions are constrained to be unidirectional (from macro to yield factors), macro factors can explain only a small portion of the variance of long yields. In contrast, when interactions are allowed to be bidirectional, the system attributes over half of the variance of long yields to macro factors. Similar results in a more robust setting are reported in Bibkov and Chernov (2010).

Finally, Rudebusch and Wu (2008) provide an example of a macro-finance specification that employs more macroeconomic structure and includes both rational expectations and inertial elements. They obtain a good fit to the data with a model that combines an affine no-arbitrage dynamic specification for yields and a small fairly standard macro model, which consists of a monetary policy reaction function, an output Euler equation, and an inflation equation. In their model, the level factor reflects market participants' views about the underlying or medium-term inflation target of the central bank, and the slope factor captures the cyclical response of the central bank, which manipulates the short rate to fulfill its dual mandate to stabilize the real economy and keep inflation close to target. In addition, shocks to the level factor feed back to the real economy through an ex-ante real interest rate.

## 1.6 Onward

In the chapters that follow, we address the issues and questions raised here, and many others. We introduce DNS in chapter 2, we make it arbitrage-free in chapter 3, and we explore a variety of variations and extensions in chapter 4. In chapter 5 we provide in-depth treatment of aspects of the interplay between the yield curve and the macroeconomy. In chapter 6 we highlight aspects of the current frontier, attempting to separate wheat from chaff, pointing the way toward additional progress.