No. 9
March 2010

StaffPAPERS

FEDERAL RESERVE BANK OF DALLAS

The Difficult Art of Eliciting Long-Run Inflation Expectations from Government Bond Prices

Carlos E. J. M. Zarazaga

No. 9, March 2010

StaffPAPERS Federal Reserve Bank of Dallas

The Difficult Art of Eliciting Long-Run Inflation Expectations from Government Bond Prices

Carlos E. J. M. Zarazaga

Senior Research Economist and Advisor

Abstract

Central banks are always concerned with keeping long-run inflation expectations well anchored at some implicit or explicit low target inflation rate. To that end, they are constantly on the lookout for indicators that can gauge those expectations accurately. One such indicator frequently reported in the specialized financial press and by central banks around the world is constructed with the forward rates technique, which exploits price differentials between government bonds of various maturities. This article examines the theory behind those indicators and assesses the extent to which they can be trusted in practice.

JEL codes: E31, E43, E58, G12

Keywords: Inflation expectations, term structure, expectations hypothesis, time-varying risk premia

---

The views expressed herein are exclusively those of the author and do not necessarily reflect those of the Federal Reserve Bank of Dallas or Federal Reserve System. The author wishes to thank Ferre DeGraeve, Jim Dolmas, Evan Koenig, Roland Meeks, Mark Wynne and an anonymous referee for useful comments on an earlier draft of this article as well as Nicole Ball for diligent research assistance. I am also grateful to Kathy Thacker for detailed editorial comments.

Monetary authorities all over the globe must be constantly alert to the impact their decisions have on inflation expectations in the long run. That is particularly true at times when monetary policy significantly departs from a pattern well understood by the markets. It is at those times that central bankers, aware of the time inconsistency problem first pointed out by Kydland and Prescott (1977), may be duly concerned with the possibility that an unusually accommodative monetary policy stance will be interpreted not as a temporary deviation from business as usual, perhaps warranted by extenuating economic or financial circumstances, but as a switch to a higher-inflation regime.

In fact, Clarida, Galí, and Gertler (2000) have argued that the main reason for the escalation of inflation rates in the U.S. between the mid-1960s and the beginning of the 1980s was that policymakers at the time didn't pay enough attention to then-rising inflation expectations and therefore failed to set policy instruments at the level required by the Federal Reserve's price stability mandate. After the appointment of Paul Volcker as Fed chairman, that changed and U.S. monetary policy took a more proactive stance, with the Fed decisively raising nominal short-term interest rates in response to higher expected inflation. Having learned from past mistakes, central bankers are always on the lookout for reliable long-run inflation-expectations indicators to guide their policy decisions.

One such popular indicator often reported by central banks and financial institutions is derived from the so-called forward rates implied by government bond yields (Figure 1).

Figure 1: Long-Run Inflation-Expectations Indicator (as Implied by Forward Rates)

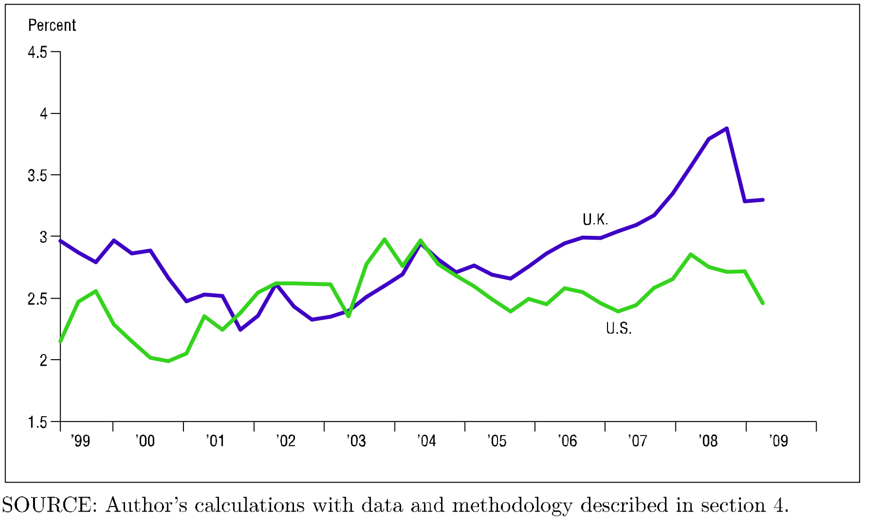

SOURCE: Author's calculations with data and methodology described in section 4.

According to the plots, long-run inflation expectations steadily rose in the U.K. between the fourth quarter of 2001 and third quarter of 2008, only to abruptly turn down almost 0.6 percentage point in the last quarter of 2008, when the local economy started to feel the effects of an unprecedented global economic crisis that threatened major banks and financial institutions in that country as well as on the other side of the Atlantic. In the U.S., the same indicator had been zigzagging in no definite direction, occasionally displaying large variations in a relatively short period of time. Examples include the 0.4 percentage-point increase that took place between the second and third quarter of 2003, and the 0.3 percentage-point decline observed shortly after the collapse of Lehman Brothers, between the fourth quarter of 2008 and first quarter of 2009.

Can inflation expectations move up and down so much in a few quarters in two countries with long histories of relatively low and stable inflation? Could it be that the indicator in Figure 1 is measuring not what is intended—fluctuations in long-run inflation expectations—but rather fluctuations in something else? The question is the subject of this article.

To provide a fair account of the answers that have been given to that question, the article will explore the possibility that the "something else" behind the fluctuations in this popular inflation-expectations indicator is time-varying risk premia, one of the many components that determine the path of interest rates over time. These premia are the channels through which the risk profile of different types of assets is reflected in their prices and returns. Loosely speaking, investors will be willing to hold in their portfolios an asset with an uncertain payoff only if it pays a higher rate of return—that is, a risk premium—over and above that paid by an asset with the same average but sure payoff. Whether or not those risk premia are constant or move over time turns out to have potentially important implications for the ability of asset-pricing models to match the observed volatility of specific asset prices or of long-run inflation-expectations indicators like the one in Figure 1.

In particular, the speculation that time-varying risk premia may be responsible for most of the fluctuations exhibited by inflation-expectations indicators constructed from government bond prices had been confined mostly to academic circles until former Federal Reserve Chairman Alan Greenspan mentioned that possibility in his monetary policy testimony before Congress in July 2005 (Greenspan, 2005a).[1] He pointed to time-varying risk premia as the main suspect for what he had referred to in his February 2005 testimony (Greenspan, 2005b) as a "conundrum"—the allegedly puzzling insensitivity that long-term nominal interest rates were exhibiting then to the monetary policy tightening cycle under way, in sharp contrast with previous experiences during such cycles.[2]

Although Greenspan's remarks elevated time-varying risk premia to the category of plausible explanation for the conundrum, the quantitative importance of those premia in accounting for the dynamics of government bond prices and associated interest rates at different maturities remains a matter of controversy. Cochrane (2007), for example, pointed out alternative explanations for the conundrum that don't involve time-varying risk premia.

In any case, this continues to be an active area of research with impor-

---

[1] See, for example, Fama (1990) and Mishkin (1990).

[2] See Rudebusch, Sack, and Swanson (2007) and Backus and Wright (2007) for more rigorous accounts of the conundrum.

tant implications for the correct interpretation of inflation-expectations indicators constructed with the forward rates technique. In particular, owing to the empirically pervasive presence of time-varying risk premia, such widely followed indicators may provide false signals and suggest that long-run inflation expectations have become unanchored when, in fact, they may remain firmly anchored at the implicit or explicit inflation rate targeted by the monetary authority. The precise source of this potential confusion can be appreciated only with a thorough understanding of the theoretical underpinnings of that class of indicators. To that end, this article derives forward-rates-based inflation-expectations indicators from basic asset-pricing principles and examines their quantitative performance under empirically plausible time-varying risk premia, always within the tractable analytical framework offered by the so-called affine factor models widely used in the financial literature.

The following is a more detailed road map to the rest of this document: Section 1 presents a nontechnical overview of the approach behind the construction of long-run inflation-expectations indicators and discusses the important role that a theory known as the expectations hypothesis plays in their interpretation. Section 2 identifies the relevant assets for the subject of this article—zero-coupon government bonds—and addresses the problem of valuing them with the same basic asset-pricing principles used for any security. This discussion leads almost naturally to the theoretical construct of affine factor models of the term structure of interest rates, which are widely used in the literature for their analytical and computational tractability. Section 3 studies how that class of models can conveniently break down government bond prices into different components—those capturing long-term inflation expectations and those capturing risk premia—and the implications of different assumptions about the nature of the latter for the reliability of inflation-expectations indicators constructed from forward rates. The constant risk premia case is examined first with a simple, easy-to-follow one-factor model. The time-varying risk premia case is studied subsequently with a more sophisticated, not-so-easy-to-follow three-factor model that can distinguish between real and nominal variables and, therefore, rationalize the nominal–real forward-rate-spread formula that is behind the calculations used to produce popular inflation-expectations indicators like the one in Figure 1. Section 4 assesses the extent to which such indicators measure what is intended—the evolution of long-run inflation expectations over time.

1. THE YIELD CURVE AND THE EXPECTATIONS HYPOTHESIS

As hinted at in the introduction, the main inputs in the construction of long-run inflation-expectations indicators are government bond yields and their term structure, known as the *yield curve.* The yield curve is a plot of the rates of return (nominal or real, as the case may be) on zero-coupon government bonds of different maturities (that is, bonds that do not pay an interest stream or coupon), like the one in Figure 2.

Figure 2: U.S. Nominal Bond Yield Curves

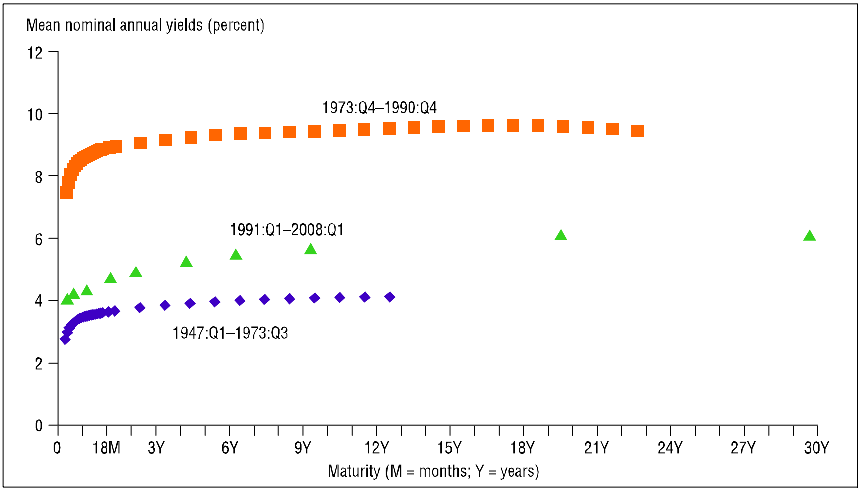

SOURCE: Author's calculations based on raw data from McCulloch and Kwon (1993) for the period 1947:Q1–1990:Q4, available at www.econ.ohio-state.edu/jhm/ts/mcckwon/mccull.htm, and from the Federal Reserve Board for the period 1991:Q1–2008:Q1.

Understanding the forces that drive bond yields may also require knowing their time-series properties; that is, the behavior of a given maturity throughout time, as seen in Figure 3. Notice that slicing the graph with a cut parallel to the horizontal axis going away from the reader produces the cross-section, or yield curve, plotted for the three periods in Figure 2. Two empirical features stand out in Figure 3: 1) The yield curve seems to be upward sloping on average, and 2) yields exhibit fairly large volatility at all maturities.

Figure 3: U.S. Nominal Yields: Maturities of 0–18M, 12M, 30M and 2–13Y

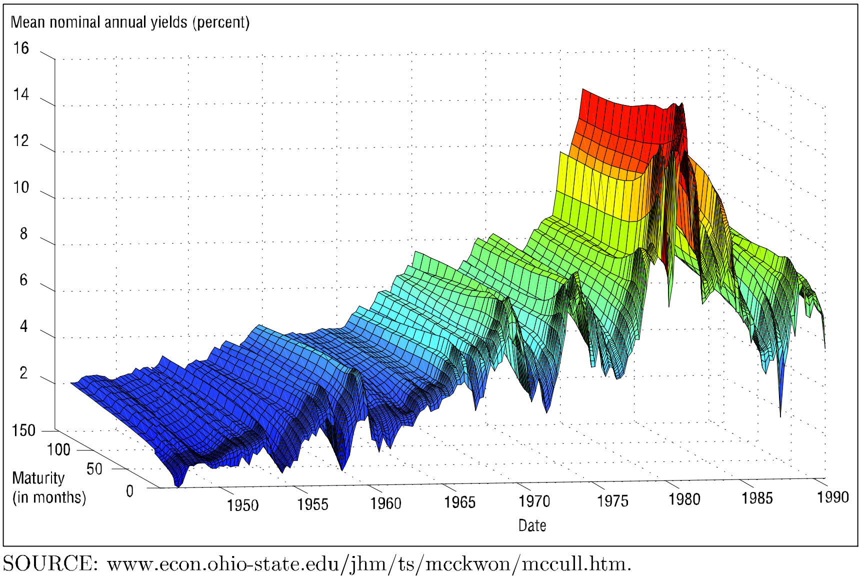

SOURCE: www.econ.ohio-state.edu/jhm/ts/mcckwon/mccull.htm.

The significance of the cross-section and time-series properties of government bond yields is that they might contain valuable information regarding long-run inflation expectations. Nonexperts might find helpful at this point an overview of the difficult art of extracting that information with the class of indicators that is the focus of this article—those that exploit price differentials between government bonds of different maturities.

The basic idea behind those indicators is that, under certain conditions discussed later, the typical nominal interest rate of the term structure can be broken down into different components, as heuristically captured by the following expression:

$$\begin{aligned} 1 + i_t^{(n)} &= \text{long-run expected inflation} + \text{long-run real interest rate} + \\ &\quad \text{inflation short-term dynamics}_t^{(n)} + \\ &\quad \text{real interest rate short-term dynamics}_t^{(n)} + \\ &\quad \text{inflation risk premium}_t^{(n)} + \text{real interest rate risk premium}_t^{(n)}, \quad (1) \end{aligned}$$

where $i_t^{(n)}$ denotes the nominal interest rate contracted at time t on a nominal financial claim with an expiration date n periods ahead.

Expression (1) is a highly stylized version of the rigorous one that will be derived in the section titled "A One-Factor Model of the Yield Curve" on page 14, which highlights that the nominal interest rates on government bonds, or nominal yields, studied in this article are typically made up of six components:

- one that captures the inflation rate expected in the long run,
- one that captures the real interest rate expected in the long run,
- two that capture the short-term dynamics of the two variables mentioned above (possibly different for each maturity), and
- two more in the last line that correspond to inflation and real interest rate risk premia (also possibly different for each maturity) whose exact nature and interpretation are discussed at length in this paper.

This representation of nominal government bond yields readily suggests the general approach behind the different methods that have been proposed to elicit long-run inflation expectations from the yield curve: to infer in every period the components other than the first in (1) from the distinct bits of information about them contained in each of the interest rates in the yield curve and then subtract those components from some appropriately chosen nominal yield or derivative of it to obtain the long-run inflation-expectations component as a residual. This is exactly what the forward rates inflation-expectations indicator plotted in Figure 1 tries to accomplish in a clever way that will be discussed in "The Nominal–Real Forward Rate Spread and Its Components" on page 32.

However, as is often the case with concepts that appear good on paper, the actual implementation of the residual method just outlined for the calculation of inflation-expectations indicators must overcome a number of practical difficulties. A serious complication that is the focus of this article is whether the risk premia components of nominal interest rates vary over time.

Until recently, the expectation hypothesis had ruled out time-varying risk premia as a potential source of the fluctuations typically exhibited by inflation-expectation indicators. The hypothesis, once widely accepted and formally discussed in "Risk Premia Components and the Expectations Hypothesis" on page 20, maintains that the risk premia in the last line of (1) are some constant, albeit possibly different for each maturity. Given that those components of the interest rate don't move over time, the argument goes, they could not account for relatively large variations in inflation-expectations indicators such as those plotted in Figure 1.

However, as pointed out by Backus and Wright (2007), a large number of empirical studies have confirmed the early finding by Macaulay (1938) that the expectations hypothesis doesn't offer a realistic description of interest rate dynamics in general and of government bond yields in particular, inclining the balance in favor of the view that the risk premia components of the yield curve exhibit substantial variations over time.[3] As discussed in this article, the consequences of that alternative assumption can be quite dramatic for the reliability of inflation-expectations indicators because they can give the impression, as the one in Figure 1 does, that long-run inflation expectations shifted quite a bit over a certain period when, in fact, they may have been firmly anchored all the time at the long-run inflation rate targeted by the monetary authority.

Before getting any deeper into the subject, however, it is necessary to introduce some notation and concepts involved in pricing government bonds, the main input in the construction of many of the inflation-expectations indicators proposed in the literature.

2. GOVERNMENT BOND PRICES AND YIELDS

Notation Preliminaries

As mentioned in the previous section, the construction of inflation-expectations indicators exploits heavily the information contained in the term structure of government bond yields; that is, in the price differential between bonds that differ just in their maturity. In practice, however, government bonds differ also in the amount and timing of interest payments. To get around this problem and make government bonds issued in different conditions comparable to each other, their prices are typically calculated in terms of their zero-coupon equivalent.

A zero-coupon bond is a special kind of security, one that pays no interest or dividends until expiration, at which time investors receive just the face value of the security. For example, a thirty-years-to-maturity zero-coupon government bond with a \$1 face value will pay investors \$1 thirty years later and nothing in between (therein the zero-coupon terminology). Since under most circumstances comparable investment options carry a strictly positive nominal interest rate, investors will not buy zero-coupon bonds unless they sell for less than their face value; that is, at a discount.

[3] See Stambaugh (1988), Cook and Hahn (1990), Campbell and Shiller (1991), Rudebusch (1993), Cochrane and Piazzesi (2005), and references in footnote 1.

To make this point more concrete, suppose that an investor knows with perfect certainty that annual interest rates will be 5 percent for the next thirty years. This investor will be willing to buy the thirty-year zero-coupon government bond at the price $\$\frac{1}{(1+\frac{5}{100})^{30}}$; that is, at a price of about 23 cents.[4]

Thus, the price of a thirty-year zero-coupon government bond is linked to its yield by the identity

$$P_t^{(30)} \equiv \frac{1}{[Y_t^{(30)}]^{30}}, \tag{2}$$

where $P_t^{(30)}$ is the price of the bond and $Y_t^{(30)}$ the underlying gross annual rate of return, or yield (equal to 1.05 in the example above). The superscripts in parenthesis should not be interpreted as mathematical exponentiation but simply as tags identifying the time to maturity of a bond originally issued at time t.

The above identity implies that the annual yield of an n period zero-coupon government bond can be calculated from its price by the formula

$$Y_t^{(n)} \equiv \frac{1}{[P_t^{(n)}]^{\frac{1}{n}}}. \tag{3}$$

Because formula (2) holds for any government bond price, it is an identity, not a theory of the ultimate determinants of government bond prices. Such a theory is the topic of the next section.

The Pricing Kernel and the Valuation of Government Bonds

Given that government bonds are securities, it seems only natural that any theory of government bond prices should start by imposing the requirement that such bonds are valued with the same pricing formula used to price any security:

$$P_t^{(i)} = E_t\left[m_{t+1}(x_{t+1}^{(i)} + P_{t+1}^{(i)})\right], \tag{4}$$

where $P_t^{(i)}$ is the price of any given security or asset i at time t, $x_{t+1}^{(i)}$ is the stochastic stream of dividends or interest payments that the security or asset will generate at $t+1$, m_{t+1} is a stochastic discount factor that converts $t+1$ payoffs into time t-equivalent values, and E_t is the conditional expectations operator that "instructs" investors to take the average over all possible joint realizations of the terms within brackets in period $t+1$, conditional on events observed up to time t. The presence of the word

[4] Notice that if the investor pays 23 cents for that bond today, he will be paid \$1 when the bond matures thirty years later, which implies that the original investment will have grown by the factor $[(1+\frac{5}{100})]^{30}$, exactly the same as if he had invested the 23 cents in a certificate of deposit renewed annually at the same 5 percent annual interest rate.

factor may be a source of ambiguities later in this article, so the stochastic discount factor m_{t+1} will be referred to hereafter as the pricing kernel.[5]

Because zero-coupon bonds do not pay interest, x_{t+1} in formula (4) is always zero, and the zero-coupon government bond pricing formula reduces therefore to:

$$P_t^{(n)} = E_t \left[m_{t+1} . P_{t+1}^{(n-1)} \right]. \tag{5}$$

Again, the superscript on P_t does not denote an exponentiation operator but is simply a label that keeps track of the years left to maturity of a bond issued in a given period t. Thus, the notation $P_{t+1}^{(n-1)}$ on the right-hand side of (5) denotes the price at time $t+1$ of a zero-coupon n-period bond originally issued at time t, which will mature and pay its \$1 face value $n-1$ periods later. Alternatively, the notation identifies a newly issued bond at time $t+1$ with maturity $(n-1)$ periods later.

It is intuitively appealing to postulate that the price of those two equivalent bonds should be the same. Otherwise, investors would buy only the cheaper version of two bonds that pay the same amount at the same time, \$1 at time $n-1$. As rigorously demonstrated in finance theory, the equality of two payoff-equivalent securities or bonds is guaranteed by imposing the no-arbitrage condition, which is satisfied if and only if the pricing kernel m_{t+1} is strictly positive in all contingencies; that is, $m_{t+1}(s) > 0$ for all s, where s identifies all possible contingencies that can materialize after time t. Following standard practice in the literature, this no-arbitrage condition will be imposed throughout the article by working with strictly positive pricing kernels.

Although the imposition of the no-arbitrage condition turns the asset-pricing formula (5) into a theory, in the sense that it would be falsifiable if it turned out that observed prices do give rise to arbitrage opportunities, it is not yet a very useful model of government bond prices because the pricing kernel m_{t+1} is not observable. While it is fairly easy to get historical records of government bond prices, there are no statistics reporting the sequence of values for m_{t+1} underlying those prices.

Rendering the theoretical government bond pricing equation (5) operational requires linking the unobservable pricing kernel m_{t+1} to observables such as the payoffs and prices of the securities under study, and even to macroeconomic variables such as aggregate consumption. Historically, this task has been accomplished with two different approaches: 1) a reverse-engineering approach, which tries to infer from the observed prices of the class of securities under study the pricing kernel underlying their valuations, and 2) a straight-engineering approach, which infers from theoret-

[5] The motivation for this alternative term widely used in the literature is that, according to equation (4), all asset prices grow, as it were, out of the same seed or kernel m_{t+1}. The term kernel has been long used in mathematics in the context of integral operators. The qualifier *pricing* has been added in finance to indicate the fact that the basic asset-pricing formula (4) is a particular incarnation of an integral equation (as captured by the expectations operator E), solved by finding the function or kernel m_{t+1} that recovers in the left-hand side, for each asset under study, the same price function for that asset plugged into the right-hand side of the equation with the time subscript appropriately shifted.

ical considerations and households' preferences over different commodity bundles (which can include leisure time) the pricing kernel that ought to be used to price all assets, including government bonds.

The next section illustrates with a concrete example how the reverse-engineering approach backs out the pricing kernel underlying government bond prices. A discussion of the alternative, straight-engineering approach is provided in Appendix A.

Reverse Engineering the Pricing Kernel from Government Bond Yields. As the previous section suggests, the key to pricing government bonds correctly, in the sense that the theoretical government bond prices obtained with formula (5) have some resemblance to those actually observed, is to get the pricing kernel right. That is precisely the principle behind the exercise in the reverse-engineering tradition carried out next.

Suppose the data suggest that the logarithm of the one-period government bond yield is normally distributed conditional on information available at time t (in other words, that the one-period yield is conditionally lognormally distributed) and that it evolves over time according to the first-order autoregressive process:

$$\ln Y_{t+1}^{(1)} = (1-\rho)(\delta - \frac{1}{2}\sigma_\varepsilon^2) + \rho \; \ln Y_t^{(1)} + (\rho + \theta)\; \varepsilon_{t+1}, \tag{6}$$

where $0 < \rho < 1$, and ε_t is an independent and identically distributed normal random variable with mean 0 and variance σ_ε^2.

One property of this process is that it is mean-reverting. That is, the one-period yield tends to revert to its long-run average value, $\delta - \frac{1}{2}\sigma_\varepsilon^2$, after it has been knocked off that value by a shock, ε_t.[6]

It seems intuitive to conjecture that government bond prices and yields inherit some of their properties from those of the pricing kernel underlying the valuation of those bonds and, therefore, to guess that the evolution of the pricing kernel through time will be described by a stochastic process with an autoregressive structure similar to that of the one-period government bond yield:

$$\ln m_{t+1} = (1 - c_2)c_1 + c_2 \ln m_t + c_3\mu_{t+1},$$

where c_1, c_2, and c_3 are unknown deterministic coefficients and μ_{t+1} is a stochastic process assumed to render $\ln m_{t+1}$ a normally distributed random variable conditional on information available at time t, on the reasonable conjecture that $\ln Y_{t+1}^{(1)}$ inherits that distribution from the pricing kernel. Rearranging terms results in the expression for the pricing kernel that will be used in what follows:

$$\ln m_{t+1} = c_1 + c_2(\ln m_t - c_1) + c_3\mu_{t+1}. \tag{7}$$

[6] One complication that will not be discussed in detail here is that the components ρ and θ of the coefficient on the residual cannot be identified only from the time-series properties of government bonds. It is also necessary to use the information contained in the cross section of government bond prices. For details, see Backus and Zin (1994).

Notice that this specification of the pricing kernel automatically satisfies the no-arbitrage condition because the logarithm is a well-defined function only for strictly positive arguments. The challenge is to infer the specific distribution of the unforecastable shock μ_{t+1} and the specific values of the coefficients in the pricing kernel (7) from the observed statistical properties of one-period government bond yields.

To that effect, notice that in terms of the notation introduced with the government bond pricing formula (5), $P_{t+1}^{(0)}$ is the price at period $t+1$ of a bond issued at period $t+1$ that will pay its face value of \$1 in that period. Such a bond is as good as cash and therefore commands a price equal to its face value; that is, $P_{t+1}^{(0)} = 1$. Taking into account this observation in the bond pricing formula (5) results in the following valuation of a one-period government bond issued at time t:

$$
\begin{aligned}
P_t^{(1)} \quad &\equiv \quad \frac{1}{Y_t^{(1)}} = E_t\left[m_{t+1}.P_{t+1}^{(0)}\right] = E_t[m_{t+1}\,.\,1] = \\
&e^{c_1 + c_2(\ln m_t \; - \; c_1) + c_3 E_t[\mu_{t+1}] \; + \frac{1}{2} E_t[c_3\mu_{t+1} - E_t[c_3\mu_{t+1}]]^2}. \qquad (8)
\end{aligned}
$$

The last equality follows from the fact that if a random variable X is conditionally lognormally distributed—as m_{t+1} is conjectured to be—its conditional expected value is given by $E_t[X] = e^{E_t[\ln X] + \frac{1}{2} Var_t[\ln X]}$, where $Var_t[\ln X]$ is a shortcut for $E_t\left[\ln X - E_t[\ln X]\right]^2$, which is the variance of $\ln X$, conditional on events observed up to time t.[7]

The next step in the process of inferring the specific values of the unknown coefficients in formula (8) is to take logarithms of both sides of it and shift the resulting expression forward one period to obtain:

$$
\begin{aligned}
\ln Y_{t+1}^{(1)} \quad &= \quad -c_1 - c_2(\ln m_{t+1} \; - \; c_1) - c_3 E_{t+1}[\mu_{t+2}] \; - \\
&\frac{1}{2} c_3^2 E_{t+1}\left[\mu_{t+2} - E_{t+1}[\mu_{t+2}]\right]^2. \qquad (9)
\end{aligned}
$$

Taking logarithms of equation (8); solving the resulting expression for $c_1 + c_2(\ln m_t \; - \; c_1)$; plugging that solution into the right-hand side of (7); using the result in lieu of $\ln m_{t+1}$ in (9); and conjecturing that the stochastic process μ_{t+1} inherits the constant variance of ε_{t+1} in (6)—$E_{t+1}\left[\mu_{t+2} - E_{t+1}[\mu_{t+2}]\right]^2 = E_t\left[\mu_{t+1} - E_t[\mu_{t+1}]\right]^2 = \sigma_\mu^2$—results, after

[7] The formula for the price of the one-period government bond $P_t^{(1)} = E_t[m_{t+1}]$ conveys the intuition for the condition $m_{t+1} > 0$ to rule out arbitrage opportunities mentioned earlier: If the pricing kernel were not strictly positive in all possible contingencies but 0 in some of them and strictly negative in the others, its expected value over all events potentially observable after time t, $E_t[m_{t+1}]$, would be strictly negative, with the implication that the price at time t of a bond paying \$1 for sure one period later would be negative. That is, investors would not only get paid for “buying” such a bond (this is what a negative price means) but also would receive \$1 at its expiration the following period. It is implausible to think that such an arbitrage opportunity to make money at no cost would last for long.

some algebra, in the following expression:

$$\ln Y_{t+1}^{(1)} = (1-c_2)(-c_1 - \frac{1}{2}c_3^2\sigma_\mu^2) + c_2 \ln Y_t^{(1)} - c_3\left\{ c_2\left[\mu_{t+1} - E_t[\mu_{t+1}]\right] + E_{t+1}[\mu_{t+2}]\right\}. \quad (10)$$

Comparing this expression with (6) reveals four equalities that the unknown entities in (10) must satisfy to reproduce the observed AR(1) process:

$$c_2 = \rho; \quad (11)$$

$$c_1 = -\delta; \quad (12)$$

$$c_3^2\sigma_\mu^2 = \sigma_\varepsilon^2; \quad (13)$$

$$-c_3\left\{\rho\left[\mu_{t+1} - E_t[\mu_{t+1}]\right] + E_{t+1}[\mu_{t+2}]\right\} = (\rho+\theta)\,\varepsilon_{t+1}.$$

The last equation can be rewritten:

$$\mu_{t+1} = E_t[\mu_{t+1}] - \left[\frac{\rho+\theta}{c_3\rho}\varepsilon_{t+1} + \frac{1}{\rho}E_{t+1}[\mu_{t+2}]\right], \quad (14)$$

which implies

$$c_3^2\sigma_\mu^2 = c_3^2 E_t\left[\mu_{t+1} - E_t[\mu_{t+1}]\right]^2 = \frac{1}{\rho^2}E_t\left[\rho\varepsilon_{t+1} + \theta\varepsilon_{t+1} + c_3 E_{t+1}[\mu_{t+2}]\right]^2 = \sigma_\varepsilon^2.$$

The last equality holds if

$$E_{t+1}[\mu_{t+2}] = -\frac{\theta}{c_3}\varepsilon_{t+1}. \quad (15)$$

This result implies $E_t[\mu_{t+1}] = -\frac{\theta}{c_3}\varepsilon_t$ and, therefore, that (14) can be rewritten:

$$\mu_{t+1} = \frac{1}{c_3}(-\theta\varepsilon_t - \varepsilon_{t+1}). \quad (16)$$

Replacing the equalities given by expressions (11), (12), and (16) in (7) represents the pricing kernel in terms of the known parameter values and stochastic variable that describe the dynamics of the one-period government bond yield according to equation (6):

$$\ln m_{t+1} = -\delta + \rho(\ln m_t + \delta) - \theta\varepsilon_t - \varepsilon_{t+1}.$$

It is a straightforward exercise to verify that this pricing kernel does generate one-period government yields with the time-series properties characterized by equation (6): Simply repeat with it the same steps that led up to equation (10).[8]

Affine Factor Models of the Pricing Kernel. To gain insight into the thinking that led to the formulation of affine factor models of the pricing kernel, solve equation (9) for $\ln m_{t+1}$ and replace the unknown coefficients and stochastic process μ_t with the right-hand side of expressions (11), (12), (13), and (15) to obtain:

$$\ln m_{t+1} = \frac{\delta}{\rho} - \delta - \frac{1}{\rho}\ln Y_{t+1}^{(1)} + \frac{1}{\rho}\theta\ \varepsilon_{t+1} - \frac{1}{\rho}\frac{1}{2}\sigma_\varepsilon^2.$$

Replacing $\ln Y_{t+1}^{(1)}$ in this expression with the right-hand side of (6) results in the following equivalent representation of the pricing kernel:

$$\ln m_{t+1} = -\frac{1}{2}\sigma_\varepsilon^2 - \gamma \ln Y_t^{(1)} - \lambda\varepsilon_{t+1}. \qquad (17)$$

Note that the right-hand side of the expression above has an affine structure; that is, it consists of a constant, plus terms that are linear in variables (or factors) and in stochastic unforecastable shocks. In the case of this example, the constant is $-\frac{1}{2}\sigma_\varepsilon^2$ and the only factor happens to be the one-period yield, $\ln Y_t^{(1)}$, with coefficient γ, typically known in the literature as factor loading because it "loads" its value onto the factor that drives the pricing kernel. The parameter λ scales the unforecastable innovation (in the sense that $E_t(\varepsilon_{t+1}Y_t) = 0$) and is referred to in some contexts as the price of risk, for reasons that will become clear shortly. Although in this particular example, the parameters γ and λ happen to be equal to 1, their values in the general case will be different, so for purposes of exposition it will be convenient to keep those symbols explicit in expression (17).

Together with the lognormality of the relevant stochastic variables, affine pricing kernels like (17) impart that structure to the prices of the securities they value. As a result, securities prices can be computed almost effortlessly with mechanical linear recursive formulas.[9]

This computational advantage, along with their analytical tractability, explains the popularity of affine factor models of the pricing kernel and the

[8] Notice that an AR(1) process for the one-period government bond yield induces an autoregressive moving-average representation—an ARMA(1,1)—of the pricing kernel, with a long-run mean of $-\delta$. Of course, more complicated time-series patterns for one-period government bonds induce more complex time-series properties in the pricing kernel as well. See Backus and Zin (1994) for a thorough treatment of this subject.

[9] See Duffie and Kan (1996) for a rigorous theoretical foundation of affine factor models of the term structure. The origins of the prolific literature on factor models can be traced to the intertemporal capital asset pricing model by Merton (1973) and to the arbitrage pricing theory developed by Ross (1976). Readers will surely wonder how this class of models has gone about selecting the factors that enter into the pricing kernel. Again, there are two approaches. The reverse-engineering approach tries to back out the underlying factors from actually observed prices, while the straight-engineering approach starts out with factors suggested by theory, subsequently checking if the government bond prices thus generated resemble those actually observed. A more detailed discussion of these two approaches can be found in Appendix A.

proliferation of studies that attempt to approximate the observed prices of all sorts of assets and securities, including government bonds, with more sophisticated versions of (17), summarized by the general expression:

$$-\ln m_{t+1} = \Psi + \sum_{h=1}^{q} \gamma_h F_h + \sum_{h=1}^{q} \lambda_h \sqrt{\alpha_h + \sum_{l=1}^{q} \beta_{h,l} F_l}\, \varepsilon_{h,t+1}, \qquad (18)$$

where Ψ is a constant; the coefficients γ_h are the factor loadings corresponding to each of the q factors F_h; the coefficients λ_h will be interpreted later (see page 22) as the price of risk associated with each source of risk, captured by the stochastic unforecastable shocks ε_h, assumed to be identically and independently normally distributed with mean 0 and variance 1; and α_h and $\beta_{h,l}$ are deterministic scalars that are greater than or equal to 0.

Odd as it may seem, the occurrence of a square root in (18) introduces in the pricing kernel the flexibility to generate government prices whose risk premia components are either constant or time varying.[10] For reasons that will be fully understood later, when all of the scalars $\beta_{h,l}$ are switched off—that is, when $\beta_{h,l} = 0$ for all h and l—the pricing kernel will satisfy the expectation hypothesis and induce government bond prices with constant risk premia components. Otherwise, when at least one of those scalars is strictly positive, the model will generate government bond prices with time-varying risk premia components.

In fact, it can be verified that the pricing kernel (17) is a special case of (18) because it is the result of multiplying both sides of the last expression by -1 and setting $\Psi = \frac{1}{2}\sigma_\varepsilon^2$, $q = 1$, $\gamma_1 = \gamma$, $F_1 = \ln Y_t^{(1)}$, $\lambda_1 = \lambda$, $\alpha_1 = 1$, $\beta_{1,1} = 0$. Given this last parameter value, it follows from the preceding discussion that the one-factor model of the pricing kernel (17) will generate bond yields characterized by constant risk premia, a useful benchmark with which to explore in the next section the attractive computational and analytical features of affine models in general.

A One-Factor Model of the Yield Curve

This section illustrates how the one-factor pricing kernel (17) generates prices for government bonds of any maturity—and, therefore, the yield curve—with methods that exploit the recursive structure of the basic zero-coupon bond pricing formula (5).

To that effect, recall that ε_{t+1} is normally distributed, which via (17) implies that m_{t+1} is conditionally lognormally distributed. The price for the one-period government bond can therefore be easily calculated by ex-

[10] The motivation for the square root trick originally introduced by Cox, Ingersoll, and Ross (1985) is that, in continuous time, this specification ensures the nominal pricing kernel doesn't generate *negative* nominal interest rates. The fact that this possibility is not ruled out in discrete time frameworks like the one adopted in this paper doesn't necessarily pose a problem. Recall that $Y_t^{(1)}$ represents a *gross* yield $1 + R_t$, so as long as the negative interest rate is not too large, say -5 percent, the yield will be positive and $\ln Y_{t+i}^{(1)}$ will still be a well-defined operation.

ploiting the lognormal distribution properties mentioned in the discussion following equation (8):

$$P_t^{(1)} = E_t[m_{t+1}] = e^{E_t[\ln m_{t+1}]+\frac{1}{2}Var[\ln m_{t+1}]} = e^{-\frac{1}{2}\sigma_\varepsilon^2 \; -\gamma \, \ln Y_t^{(1)} + \frac{1}{2}\lambda^2\sigma_\varepsilon^2}.$$

Taking logarithms and considering definition (3) results in the one-period government bond price:

$$\begin{aligned} -\ln Y_t^{(1)} &\equiv \ln P_t^{(1)} = E_t\left[\ln m_{t+1}\right] + \frac{1}{2}E_t\left[\ln m_{t+1} - E_t[\ln m_{t+1}]\right]^2 \\ &= -\frac{1}{2}\sigma_\varepsilon^2 - \gamma \ln Y_t^{(1)} + \frac{1}{2}\lambda^2\sigma_\varepsilon^2. \qquad (19) \end{aligned}$$

This equation imposes the restriction $\gamma = 1$ and $\lambda = 1$ because those are the values for the factor-loading parameter and the price-of-risk parameter (with interpretations to be given later) that make both ends of the equation equal.[11]

Consider next the problem of pricing the two-period zero-coupon bond with the version of the pricing formula (5) specific to that bond:

$$P_t^{(2)} = E_t\left[m_{t+1}.P_{t+1}^{(1)}\right].$$

The two-period bond price can be readily computed with this formula because the composite random variable within brackets inherits from its components the conditionally lognormal distribution implied by the identity $\ln P_{t+1}^{(1)} \equiv -\ln Y_{t+1}^{(1)}$ and equations (6) and (17). Exploiting once more the statistical property for lognormally distributed variables,

$$\begin{aligned} P_t^{(2)} &= E_t\left[m_{t+1}.P_{t+1}^{(1)}\right] \\ &= e^{E_t[\ln(m_{t+1}P_{t+1}^{(1)})]+\frac{1}{2}E_t\{\ln(m_{t+1}P_{t+1}^{(1)})-E_t[\ln(m_{t+1}P_{t+1}^{(1)})]\}^2}. \end{aligned}$$

Algebraic manipulations after taking logs on both sides of the equation result in the following relationship:

$$\begin{aligned} \ln P_t^{(2)} &= E_t\left[\ln m_{t+1}\right] + \frac{1}{2}Var_t[\ln m_{t+1}] + E_t\left[\ln P_{t+1}^{(1)}\right] + \\ &\quad \frac{1}{2}Var_t[\ln P_{t+1}^{(1)}] + Cov_t[\ln m_{t+1}, \ln P_{t+1}^{(1)}], \qquad (20) \end{aligned}$$

where Var_t and Cov_t denote, respectively, the conditional variance and covariance of the random variables within the corresponding brackets.

Taking into account that the sum of the first two terms in (20) is equal to $-\ln Y_t^{(1)}$ (see equation 19), the two-period bond price can be written in the more compact notation:

[11] As discussed in more detail in Appendix A, the factor $\ln Y_t$ could have been kept latent; that is, as an unidentified factor $\ln X_t$ that equation (19) would have later revealed to be the one-period bond yield with factor loading and price of risk equal to 1. More complicated versions of this latent-variables procedure are widely employed in the empirical literature.

$$\ln P_t^{(2)} = -\ln Y_t^{(1)} + E_t[\ln P_{t+1}^{(1)}] + \frac{1}{2} Var_t[\ln P_{t+1}^{(1)}] + Cov_t[\ln m_{t+1}, \ln P_{t+1}^{(1)}]. \tag{21}$$

All that is needed to obtain an expression for the two-year bond in terms of fundamental parameters is to carry out the expectations and variances calculations indicated on the right-hand side of the formula. To that effect, notice that definition (3) and equation (6) imply:

$$\ln P_{t+1}^{(1)} = -\ln Y_{t+1}^{(1)} = -(1-\rho)(\delta - \frac{1}{2}\sigma_\varepsilon^2) - \rho \; \ln Y_t^{(1)} - (\theta+\rho) \; \varepsilon_{t+1},$$

which, in turn, implies:

$$\begin{aligned} E_t \ln P_{t+1}^{(1)} &= -(1-\rho)(\delta - \frac{1}{2}\sigma_\varepsilon^2) - \rho \; \ln Y_t^{(1)}; \qquad (22) \\ Var_t(\ln P_{t+1}^{(1)}) &= E_t \left[\ln P_{t+1}^{(1)} - E_t[\ln P_{t+1}^{(1)}] \right]^2 \\ &= Var_t(\ln Y_{t+1}^{(1)}) = (\theta+\rho)^2 \sigma_\varepsilon^2; \qquad (23) \\ Cov_t(\ln m_{t+1}, \ln P_{t+1}^{(1)}) &= E_t \Big[(\ln m_{t+1} - E_t[\ln m_{t+1}]) \\ &\qquad (\ln P_{t+1}^{(1)} - E_t[\ln P_{t+1}^{(1)}]) \Big] \\ &= \lambda(\theta+\rho)\sigma_\varepsilon^2. \qquad (24) \end{aligned}$$

Plugging (22), (23), and (24) into (21), adding and subtracting $\delta - \frac{1}{2}\sigma_\varepsilon^2$ to and from the right-hand side of the resulting expression, and rearranging terms delivers the two-period yield as a function of primitive parameters:

$$\begin{aligned} -\ln P_t^{(2)} &= 2 \; \ln Y_t^{(2)} = 2(\delta - \frac{1}{2}\sigma_\varepsilon^2) + (1+\rho)[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)] - \\ &\quad \lambda(\theta+\rho)\sigma_\varepsilon^2 - \frac{1}{2}(\theta+\rho)^2\sigma_\varepsilon^2. \qquad (25) \end{aligned}$$

The prices for all subsequent maturities could be calculated with the bond pricing formula (5) by repeated application of the same steps followed to calculate the price of a two-period bond. In particular, the price of an n-period bond could be calculated with the formula

$$P_t^{(n)} = E_t \left[m_{t+1}.P_{t+1}^{(n-1)} \right] = e^{E_t[\ln(m_{t+1}.P_{t+1}^{(n-1)})] + \frac{1}{2} Var_t[\ln(m_{t+1}.P_{t+1}^{(n-1)})]},$$

which, after taking logarithms on both sides and applying some algebra, results in the following generic expression for the log of the price of a bond with arbitrary maturity:

$$\begin{aligned}\ln P_t^{(n)} &= -\ln Y_t^{(1)} + E_t[\ln P_{t+1}^{(n-1)}] + \\ &\frac{1}{2}E_t\left[\ln P_{t+1}^{(n-1)} - E_t[\ln P_{t+1}^{(n-1)}]\right]^2 + \\ &E_t\left[(\ln m_{t+1} - E_t[\ln m_{t+1}])(\ln P_{t+1}^{(n-1)} - E_t[\ln P_{t+1}^{(n-1)}])\right].\end{aligned} \tag{26}$$

However, this method for generating the yield curve requires the successive substitution of prices represented in terms of variables dated at time t or earlier on the right-hand side of the formula, a tedious process that becomes unwieldy rather quickly. It turns out that the affine structure of the model makes it possible to accomplish that same task with a shortcut.

A more expedient procedure to compute bond prices for all maturities and to generate the yield curve is to exploit the observation that, after adding and subtracting the term $\frac{1}{2}\lambda^2\sigma_\varepsilon^2$ on the right-hand side of (25), the nominal price of the two-period bond can be rewritten as a linear function of the only factor, the one-period government bond yield, and of its conditional variance:

$$-\ln P_t^{(2)} = 2\ln Y_t^{(2)} = A_2 + B_2\left[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)\right] + C_2 Var_t[\ln Y_{t+1}^{(1)}], \tag{27}$$

where:

$$\begin{aligned}A_2 &= 2(\delta - \frac{1}{2}\sigma_\varepsilon^2) + \frac{1}{2}\lambda^2\sigma_\varepsilon^2; \\ B_2 &= 1 + \rho; \\ C_2 &= -\frac{1}{2}\left(\frac{\lambda}{\theta+\rho} + 1\right)^2.\end{aligned}$$

Given that $Var_t[\ln Y_{t+1}^{(1)}]$ is the constant $(\theta+\rho)^2\sigma_\varepsilon^2$, it follows that factor models of the pricing kernel impart their affine structure as well to the prices of the securities under study—government bonds, in this case. As demonstrated next, this feature makes possible the calculation of government bond prices for all maturities using a recursive formula that is easy to implement in a computer program.

To see the recursive structure of government bond prices generated by such models, make the guess that the formula for the price of any government bond is the following natural extension of the shortcut formula (27):

$$-\ln P_t^{(n)} = n\ln Y_t^{(n)} = A_n + B_n\left[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)\right] + C_n Var_t[\ln Y_{t+1}^{(1)}]. \tag{28}$$

Next, substitute $P_{t+1}^{(n-1)}$ on the right-hand side of formula (26) for the right-hand side of (28), with subscripts and superscripts relabeled in a consistent manner. Take into account (6), (17), and the equality $Var_{t+i}[\ln Y_{t+i+1}^{(1)}] = Var_t[\ln Y_{t+1}^{(1)}]$ for $i \geqq 0$ implied by the former equation, and subsequently match the coefficients in the resulting expression with those for the corresponding terms on the right-hand side of (28), to reveal the following recursive structure for the coefficients A_n, B_n, and C_n:

$$\begin{aligned} A_n &= A_{n-1} + \delta - \frac{1}{2}\sigma_\varepsilon^2 + \frac{1}{2}\lambda^2\sigma_\varepsilon^2; \qquad (29) \\ B_n &= 1 + \rho B_{n-1}; \\ C_n &= C_{n-1} - \frac{1}{2}\Big(\frac{\lambda}{\theta+\rho} + B_{n-1}\Big)^2, \qquad (30) \end{aligned}$$

with initial conditions $A_0 = B_0 = C_0 = 0$ obtained from the requirement that the recursive formula (28) must deliver the price of the one-period bond when $n = 1$ (recall that $\ln Y_t^{(1)} \equiv -\ln P_t^{(1)}$).

Because the recursion for the coefficient B_n implies that it behaves like the truncated geometric progression $\frac{1-\rho^n}{1-\rho}$, for consistency with the other coefficients, it will be more convenient to use the following alternative equivalent representation:

$$B_n = B_{n-1} + \rho^{n-1}. \qquad (31)$$

Two features of the recursive price formula (28) are worth highlighting. First, the coefficients for each maturity are successively built up, as it were, by adding to the coefficients of the immediately preceding maturity the increments represented by the terms following the first coefficient on the right-hand sides of expressions (29), (30), and (31). This feature will become handy at the time of computing forward rates.

Second, as it follows from (23), for the specific one-factor model analyzed in this section, the conditional variance of the factor is the constant $(\theta+\rho)^2\sigma_\varepsilon^2$. As a result of this property, known as homoskedasticity in the statistical jargon, the term $C_n Var_t[\ln Y_{t+1}^{(1)}]$ in (28) is for each maturity the constant $C_n(\theta+\rho)^2\sigma_\varepsilon^2$, which could be consolidated into the constant A_n, a standard practice not adopted here, for expositional reasons.

It is a straightforward (albeit tedious) exercise to show that the recursive application of (28) to bonds of successive maturities generates the following expressions for the price (and associated yields) of two-, three-, and four-period bonds:

$$\begin{aligned} -\ln P_t^{(2)} &= 2\ln Y_t^{(2)} = 2(\delta - \frac{1}{2}\sigma_\varepsilon^2) + \\ &\quad (1+\rho)\Big[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)\Big] - \\ &\quad \frac{1}{2}(\theta+\rho)^2\sigma_\varepsilon^2 - \lambda(\theta+\rho)\sigma_\varepsilon^2; \qquad (32) \end{aligned}$$

$$
\begin{aligned}
-\ln P_t^{(3)} &= 3\ln Y_t^{(3)} = 3(\delta - \frac{1}{2}\sigma_\varepsilon^2) + \\
&(1+\rho+\rho^2)\Big[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)\Big] - \\
&\frac{1}{2}\Big[1+(1+\rho)^2\Big](\theta+\rho)^2\sigma_\varepsilon^2 - \lambda\Big[1+(1+\rho)\Big](\theta+\rho)\sigma_\varepsilon^2;
\end{aligned} \tag{33}
$$

$$
\begin{aligned}
-\ln P_t^{(4)} &= 4\ln Y_t^{(4)} = 4(\delta - \frac{1}{2}\sigma_\varepsilon^2) + \\
&(1+\rho+\rho^2+\rho^3)\Big[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)\Big] - \\
&\frac{1}{2}\Big[1+(1+\rho)^2+(1+\rho+\rho^2)^2\Big](\theta+\rho)^2\sigma_\varepsilon^2 - \\
&\lambda\Big[1+(1+\rho)+(1+\rho+\rho^2)\Big](\theta+\rho)\sigma_\varepsilon^2.
\end{aligned} \tag{34}
$$

Expression (32) reproduces the price for the two-period bond (25) obtained with the brute-force approach pursued earlier to motivate the recursive formula (28).

Notice also that the three formulas are a rigorous restatement of the stylized formula (1) presented in section 1. This can be seen more clearly by taking, for example, the formula for the four-period bond, dividing both sides of it by 4 and considering the relationship $\ln(1+i_t) = E_t[\ln(1+r_{t+1})] + E_t[\ln(1+\Pi_{t+1})]$, the popular logarithmic variation of the Fisher equation proposed by Irving Fisher (1896; 1930) in his attempts to formalize a theory on the determination of interest rates. Under the assumption that the short-term nominal interest rate has a well-defined long-run or unconditional mean (formally, that $E[\ln(1+i_t)] = E[\ln(1+r_{t+1})] + E[\ln(1+\Pi_{t+1})] = E[\ln(1+r_t)] + E[\ln(1+\Pi_t)]$), the logarithmic version of the Fisher equation and (6) imply $E[\ln Y_t^{(1)}] = \delta - \frac{1}{2}\sigma_\varepsilon^2 = E[\ln(1+r_t)] + E[\ln(1+\Pi_t)]$.

The resulting expression for the yield associated with a four-period bond is:

$$
\begin{aligned}
\ln Y_t^{(4)} &= E\ln(1+\Pi_t) + E\ln(1+r_t) + \\
&\frac{1}{4}(1+\rho+\rho^2+\rho^3)\Big[E_t[\ln(1+\Pi_{t+1}) - E\ln(1+\Pi_t)] + \\
&\qquad E_t[\ln(1+r_{t+1}) - E\ln(1+r_t)]\Big] - \\
&\frac{1}{8}\Big[1+(1+\rho)^2+(1+\rho+\rho^2)^2\Big](\theta+\rho)^2\sigma_\varepsilon^2 - \\
&\frac{1}{4}\lambda\Big[1+(1+\rho)+(1+\rho+\rho^2)\Big](\theta+\rho)\sigma_\varepsilon^2.
\end{aligned}
$$

The parallel of this expression with its heuristic version (1) in section 1 is apparent if the last two terms correspond to risk premia components, which is indeed the case, as demonstrated in the next section.

Risk Premia Components and the Expectations Hypothesis. To see why the last two terms in formulas (32) to (34) contain the risk premia components of government bond yields, start with expression (21) for the price of a two-period government bond. Rearranging terms after taking into account identity (3) results in the following expression:

$$2 \ln Y_t^{(2)} - (\ln Y_t^{(1)} + E_t \ln Y_{t+1}^{(1)}) = -\frac{1}{2} Var_t[\ln P_{t+1}^{(1)}] - Cov_t[\ln m_{t+1}, \ln P_{t+1}^{(1)}]. \quad (35)$$

Notice that the left-hand side of this equation is the differential between the return that an investor would expect from an investment strategy of buying a two-period bond at time t and holding it until maturity, represented by the first term, and the alternative strategy of buying in sequence two one-period bonds, represented by the terms between parentheses. If these two portfolios were exactly equivalent in the eyes of investors, their returns should be the same, which by the no-arbitrage condition implies that the right-hand side of this equation should be equal to 0.

By the same token, a nonzero differential on the right-hand side would imply, in the absence of arbitrage, that investors don't regard the two investment options as equivalent. A negative differential would imply that investors see the two-period bond as a safer investment than the sequence of two one-period bonds and are willing to give up returns or, equivalently, pay an insurance premium, to hold the two-period bond in their portfolios. A positive differential would imply that the two-period bond is considered a riskier investment and that investors will therefore demand a higher return for it; that is, a risk premium over the sequence of two one-period bonds.

Whether investors pay an insurance premium or get paid a risk premium for holding a bond of given characteristics in their portfolios is determined by the right-hand side of equation (35). In other words, the two right-hand-side terms identify the risk premia components of the nominal interest rate heuristically introduced in section 1.

An entirely analogous procedure identifies the risk premia components of a government bond of arbitrary maturity n with the right-hand side of the following general version of (35):

$$\text{Risk premia components} = n \ln Y_t^{(n)} - \sum_{i=0}^{n-1} E_t \ln Y_{t+i}^{(1)}$$

$$= -\frac{1}{2} \sum_{i=0}^{n-2} Var_t[\ln P_{t+1+i}^{(1)}] - \sum_{i=0}^{n-2} Cov_t[\ln m_{t+1+i}, \ln P_{t+1+i}^{(1)}]. \quad (36)$$

It can be verified that substituting the recursive price formula (28) in the right-hand side of the expression, taking into account (6), generates risk premia components for the two- to four-period bonds that exactly replicate the third and fourth terms of the price formulas for those maturities given by (32) through (34). The different components of those formulas can, therefore, be given meaningful economic interpretations.

The first term on the right-hand side of formulas (32) to (34), $\delta - \frac{1}{2}\sigma_\varepsilon^2$, captures the value to which the short nominal interest rate—the one-period yield—reverts in the long run according to (6). If it were up to this term, the yield curve would be flat on average, as is apparent from dividing both sides of equations (32) through (34) by the corresponding maturity length.

The second term on the right-hand side of those formulas—the term involving $[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)]$—captures the change in yields induced by deviations of the time t one-period yield from its stationary value or long-run mean. The size of the response of each yield in the term structure to movements in the one-period yield is ultimately controlled by the sequence of coefficients $\{B_n/n\}$ in the appropriately modified version of (28) for yields. In light of the arguments preceding equation (31), those coefficients behave like the sequence $\left\{\frac{\frac{1-\rho^n}{1-\rho}}{n}\right\}$, which decreases monotonically to 0 because, from one member of the sequence to the next, the denominator is always augmented by 1, while the numerator is always augmented by a smaller number that decreases to 0 at the geometric rate ρ. This pattern implies that the impact that deviations of the one-period yield have on subsequent yields dies out with the maturity of the bond. In particular, yields for maturities in the distant future do not respond at all to those deviations (i.e., $\lim_{n\to\infty} \frac{1}{n}\frac{1-\rho^n}{1-\rho} = \frac{1}{1-\rho}\lim_{n\to\infty}\frac{1}{n} = 0$). Regardless of how they propagate throughout the term structure, positive and negative deviations of the one-period yield from its long-run mean cancel each other on average—that is, $E[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)] = 0$. So if it were up to the second term in formulas (32) through (34) just analyzed, the yield curve would still be flat on average instead of upward sloping, as Figure 2 suggests.

Thus, the forces that induce a positive slope in the yield curve must be contained in the third and fourth terms on the right-hand side of formulas (32) to (34) which, for the arguments given when interpreting equation (35), capture the risk premia components built into the bond yield for the associated maturity. By comparison with expression (36), the third term corresponds to the sum of the variances generated by each element of the price sequence $\{\ln P_{t+i}^{(1)}\}$, while the fourth corresponds to the sum of the covariances of each of those elements with the contemporaneous pricing kernel.

The third term is a statistical artifact that emerges in the process of going from $E_t \ln P_{t+i}^{(1)}$ to $E_t P_{t+i}^{(1)}$.[12] The generic representation of this

[12]For that reason, this variance-related term is sometimes referred to as a Jensen's inequality component, in reference to the result established by Jensen that if $f(.)$ is a concave function of the stochastic variable x, then $E_t f(x) < f(E_t x)$. Since $\ln(x)$ is a concave function of x, it follows that $E_t \ln P_{t+1}^{(1)} < \ln(\mathrm{E}_t P_{t+1}^{(1)})$. This condition and the definition of logarithm imply that $\mathrm{E}_t P_{t+1}^{(1)} > e^{E_t \ln P_{t+1}^{(1)}}$, an inequality that is satisfied in this case because the variance is a strictly positive number and lognormally distributed prices imply $E_t P_{t+1}^{(1)} = e^{E_t \ln P_{t+1}^{(1)} + \frac{1}{2}Var(\ln P_{t+1}^{(1)})}$. Readers may wonder what happened to the variance, or Jensen's inequality component, corresponding to the pricing kernel. The answer is that it is hidden under the guise $\frac{1}{2}\sigma_\varepsilon^2$ in the first term of formulas (32) to (34).

variance-related term, $-(\theta+\rho)^2\sigma_\varepsilon^2 \sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})^2$, implies that the absolute value of its counterpart in the yield formula is increasing in the maturity of the bond. As a result of this property, this third term introduces a downward bias in the yield curve.[13]

That leaves the fourth term on the right-hand side of (32) to (34) as the only one capable of making the yield curve slope up on average, as it does in the U.S. data, according to Figure 2. The generic representation of the counterpart of this variance-related term in the yield formula is $-\lambda(\theta+\rho)\sigma_\varepsilon^2 \frac{1}{n} \sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})$. It is easy to verify that the arguments in footnote 13 apply also to this term and that its absolute value, therefore, increases as well with the maturity of the bond. It follows that this fourth term cannot induce an upward-sloping yield curve unless it is made overall positive by a negative covariance between the pricing kernel and the contemporaneous one-period bond price; that is, unless $\lambda(\theta+\rho)\sigma_\varepsilon^2 < 0$. For the particular example under consideration, this condition is equivalent to the requirement that the coefficients on the unforecastable errors of the autoregressive process (6), $(\theta+\rho)$, and of the pricing kernel (17), λ, have the opposite sign.

The negative covariance condition is necessary but not sufficient to generate an upward-sloping yield curve because the size of the covariance matters, too: An overall positive but small covariance-related fourth term might be overwhelmed by the overall negative variance-related third term. In line with the discussion following equation (35), this would imply that the bonds' risk decreases with their maturity and that investors would be willing to accept lower yields on long-term bonds than on short-term bonds. Only a large enough negative covariance will make a bond of any given maturity riskier than the preceding ones in the term structure and, therefore, induce a positive slope in the yield curve.

The role that the sign and size of the covariance between $\ln m_{t+1}$ and $\ln P_{t+1}^{(1)}$ play in shaping the slope of the yield curve makes apparent why the literature refers to the parameter λ as the price of risk: The last term in formulas (32) to (34) can be interpreted as the total cost of risk associated with a bond of a given maturity, obtained from multiplying the unit price of risk λ, which is the same for all maturities, by the quantity of risk contained in each bond, measured by the remaining parameters in that last term, which is typically different for each maturity: $(\theta+\rho)\sigma_\varepsilon^2/2$ for the two-period bond yield, $(\theta+\rho)\sigma_\varepsilon^2\frac{1}{3}[1+(1+\rho)]$ for the three-period bond

[13] To see this, divide both sides of formulas (32) to (34) by the length of the associated maturity. Simple algebra shows that the size of the generic third term on the right-hand side of the resulting expressions increases with the maturity of the bond: $\frac{1}{n+1}\sum_{h=1}^{n+1}(\frac{1-\rho^h}{1-\rho})^2 - \frac{1}{n}\sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})^2 = \frac{1}{(n+1)n}\left[n\sum_{h=1}^{n+1}(\frac{1-\rho^h}{1-\rho})^2 - (n+1)\sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})^2\right] = \frac{1}{(n+1)n}\left[n(\frac{1-\rho^{n+1}}{1-\rho})^2 + n\sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})^2 - n\sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})^2 - \sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})^2\right] = \frac{1}{(n+1)n}\left[n(\frac{1-\rho^{n+1}}{1-\rho})^2 - \sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})^2\right] > 0$, because $0 < \rho < 1$ and, therefore, each of the n elements in the last sum is smaller than $(\frac{1-\rho^{n+1}}{1-\rho})^2$.

yield, $(\theta+\rho)\sigma_\varepsilon^2\frac{1}{4}[1+(1+\rho)+(1+\rho+\rho^2)]$ for the four-period bond yield, and $(\theta+\rho)\sigma_\varepsilon^2\frac{1}{n}\sum_{h=1}^{n}(\frac{1-\rho^h}{1-\rho})$ in general for the n-period bond.[14]

Thus, this analysis has established that the risk premia components are different for each maturity and also, important for the arguments to follow, that they remain the same over time *for any given maturity.* This is the key prediction of the expectations hypothesis. The fact that this hypothesis holds for the particular one-factor model studied in this section can be exploited to check whether long-run inflation expectations are well-anchored with the forward rates method discussed next.[15]

3. FORWARD RATES AND LONG-RUN INFLATION EXPECTATIONS

The Constant Risk Premia Case

Imagine an investor that at time t simultaneously performs the following financial transactions: He buys one n-period zero-coupon bond and sells $(P_t^{(n)}/P_t^{(n+j)})$ units of an $(n+j)$-period zero-coupon bond. This operation guarantees that the investor n periods later will receive \$1, the face value of the n-period bond he purchased at t, for which he will have to pay $n+j$ periods later the amount $[(P_t^{(n)}/P_t^{(n+j)})\,.\,\$1]$; that is, the quantity of $(n+j)$-period bonds he sold at t times their \$1 face value.

Note that this sequence of financial transactions is equivalent to a contract that guarantees to the investor at time t that he will be able to get a loan of \$1 at time $t+n$, provided he commits to pay back $\$(P_t^{(n)}/P_t^{(n+j)})$ at time $t+n+j$. The total return on the loan over its j-period duration is given by the expression $(P_t^{(n)}/P_t^{(n+j)})/\$1$, the annualized version of which is the forward rate, in reference to the underlying forward loan contract. Formally,

$$FN_t^{(n)\to(n+j)} = \left(\frac{P_t^{(n)}}{P_t^{(n+j)}}\right)^{\frac{1}{j}},$$

where FN_t stands for nominal forward rate, given that nominal prices appear in the right-hand side of the expression. In what follows, it will be more convenient to work with the logarithmic version of this expression,

$$\ln FN_t^{(n)\to(n+j)} = \frac{\ln P_t^{(n)} - \ln P_t^{(n+j)}}{j}. \tag{37}$$

Replacing the recursive formula (28) in the right-hand side of (37) for the case in which $j = 1$ results in the following expression for the

[14] Fisher (2001) offers an accessible and more detailed explanation of the different forces that shape the yield curve. See Backus and Zin (1994) for a rigorous statement of the conditions that guarantee an upward-sloping yield curve. Piazzesi and Schneider (2006) show that those conditions can be rationalized with microfoundations if inflation is, on average, bad news for consumption growth.

[15] In fact, the expectations hypothesis is often formally stated using equation (36) but with its right-hand side set equal to some constant. See chapter 10 of Campbell, Lo, and MacKinlay (1997) for the different forms the expectations hypothesis has taken over time.

n-maturity one-period-ahead forward nominal rate:

$$\begin{aligned}\ln FN_t^{(n)\to(n+1)} &= A_{n+1} - A_n + (B_{n+1} - B_n)\left[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)\right] + \\ &\quad (C_{n+1} - C_n)(\theta+\rho)^2\sigma_\varepsilon^2 \\ &= \delta - \frac{1}{2}\sigma_\varepsilon^2 + \rho^n\left[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)\right] - \\ &\quad \frac{1}{2}\left(\frac{1-\rho^n}{1-\rho}\right)^2(\theta+\rho)^2\sigma_\varepsilon^2 - \lambda\frac{1-\rho^n}{1-\rho}(\theta+\rho)\sigma_\varepsilon^2, \qquad (38)\end{aligned}$$

where the second equality has been obtained after taking into account (29), (30), and (31) and developing the square in the right-hand side of (30).

The expression explains the popularity of forward rates: The trick of subtracting successive bond prices retains in the resulting expression only the increments by which the coefficients of the price formula (28) are augmented in the process of going from one maturity to the next. As a result, forward rates tend to approximate the long-run means of the real interest rate and inflation more accurately than the underlying yields themselves.

This assertion can be verified by writing formula (38) for the concrete case $n = 3$:

$$\begin{aligned}\ln FN_t^{(3)\to(4)} &= \delta - \frac{1}{2}\sigma_\varepsilon^2 + \rho^3\left[\ln Y_t^{(1)} - (\delta - \frac{1}{2}\sigma_\varepsilon^2)\right] - \\ &\quad \frac{1}{2}(1+\rho+\rho^2)^2(\theta+\rho)^2\sigma_\varepsilon^2 - \lambda(1+\rho+\rho^2)\Big)(\theta+\rho)\sigma_\varepsilon^2.\end{aligned}$$

Comparison of this expression with (34) confirms that only the tail of the coefficients A_4, B_4, and C_4 enters into the forward rate. This is a direct consequence, of course, of the incremental nature of the recursions (29), (30), and (31).

Furthermore, for forward rates far into the future, the term associated with deviations from long-run means will tend to be negligible. For example, if $\rho = 0.8$, a value not too far from those used in the literature at annual frequencies, the coefficient for the third term in the equation corresponding to the fourteen-fifteen-year forward rate will be $\rho^{14} \approx 0.04$. Such small numbers suggest that, for practical purposes, the third term in the immediately preceding expression can be ignored when looking at forward rates several periods ahead. For long enough maturities, therefore, the general formula for nominal forward rates (38) can be approximated with the simpler expression:

$$\begin{aligned}\ln FN_t^{(n)\to(n+1)} &\approx E[\ln(1+r_t)] + E[\ln(1+\Pi_t)] - \\ &\quad \frac{1}{2}\left(\frac{1-\rho^n}{1-\rho}\right)^2(\theta+\rho)^2\sigma_\varepsilon^2 - \lambda\frac{1-\rho^n}{1-\rho}(\theta+\rho)\sigma_\varepsilon^2,\end{aligned} \qquad (39)$$

where the first two terms pop up as a result of taking into account that (6) and the assumption $E[\ln(1+i_t)] = E[\ln(1+r_t)] + E[\ln(1+\Pi_t)]$, justified on page 19, imply $\delta - \frac{1}{2}\sigma_\varepsilon^2 = E\ln Y_t^{(1)} = E[\ln(1+r_t)] + E[\ln(1+\Pi_t)]$.

A remarkable implication of approximation (39) is that the plot of a time series of forward rates between any two given maturities in the distant future will look like a flat line if long-run inflation expectations are well anchored at the long-run inflation rate $E\ln(1+\Pi_t)$. This is because, by assumption, the real interest rate has a well-defined long-run mean and because the last two terms are constant for each maturity, given that they correspond to the tails of risk premia components generated by a single-factor model for which the expectations hypothesis holds.

The preceding analysis suggests that there is a relatively straightforward test to check whether long-run inflation expectations are well anchored or not when the expectations hypothesis is satisfied. The inspiration for such a test comes from many studies in the literature, such as that by Rudebusch (1998) and the one by Clarida, Galí, and Gertler (2000) mentioned in the introduction, which maintain the hypothesis in their analyses and, as a result, predict that forward rates far into the future should be unresponsive to monetary policy surprises.

Gürkaynak, Sack, and Swanson (2005) tested that implication by running regressions of U.S. forward rates of different maturities on indicators of monetary policy surprises. They found that even forward rates far into the future move quite a bit in response to unexpected changes in monetary policy. If the real interest rate exhibits mean reversion in the long run and the risk premium component of the forward rate is constant, as the reference models assume, the only thing that can move forward rates in the long end of the term structure is nonanchored inflation expectations or, equivalently, an inflation process that doesn't have a well-defined long-run mean. This could be a hint that the monetary policy regime in place is such that it prompts revisions of long-run inflation expectations every time the monetary authority surprises the markets with a decision.

However, as Gürkaynak, Sack, and Swanson were quick to point out, it is also possible that the assumptions of the reference models are not satisfied in reality. It could be, for example, that contrary to the assumption made in deriving (39), the real interest rate does not fluctuate around any well-defined long-run mean but instead drifts over time in a random walk fashion.[16] It is clear from inspection of (39) that permanent movements of the real interest rate could be responsible for movements of the forward rates in the long end of the term structure.

Alternatively, it could be that the expectations hypothesis does not hold and that the last two terms in (39) appear to be constant only as a result of model misspecification. Observers unaware of this fact would tend to attribute fluctuations in forward rates in the far end of the term structure to unanchored long-term inflation expectations, when the real source of such fluctuations is time-varying risk premia. Because that seems to be a more empirically appealing assumption, as documented earlier, it is imperative to examine in the next section its implications for the interpretation of inflation-expectations indicators constructed from forward rates.

---

[16] As documented in Phillips (1998), the stochastic properties of the real interest rate is a subject of debate among experts.

The Time-Varying Risk Premia Case

Disentangling the Nominal and Real Pricing Kernels. The implications of time-varying risk premia for the interpretation of inflation-expectations indicators constructed from forward rates will become more transparent by distinguishing between nominal and real government bond prices. That can be accomplished by decomposing the nominal pricing kernel as shown in this section into a real part, which prices government bonds in real terms, and an inflation part, which transforms those real prices into nominal ones.

The starting point for that decomposition is the observation that up to now, government bonds have been priced in nominal terms and that the underlying pricing kernel is therefore a nominal pricing kernel, a fact that will be convenient to recognize with the notation $m^{\$}_{t+1}$. With this new convention, formula (5) for pricing nominal government bonds should be rewritten as follows:

$$P^{(n)}_t = E_t\left[m^{\$}_{t+1}.P^{(n-1)}_{t+1}\right], \tag{40}$$

which, for the special case of a one-period nominal zero-coupon bond, reduces to:

$$P^{(1)}_t = E_t\left[m^{\$}_{t+1}.\$1\right] = E_t[m^{\$}_{t+1}]. \tag{41}$$

The real pricing kernel should price bonds according to their value in real terms, which can be calculated by dividing each period's nominal price by the value of a representative basket of goods or a price index, $\mathbb{P}_t$, as follows:

$$\frac{P^{(n)}_t}{\mathbb{P}_t} = E_t\left[m_{t+1}.\frac{P^{(n-1)}_{t+1}}{\mathbb{P}_{t+1}}\right], \tag{42}$$

where m_{t+1} denotes the real pricing kernel.

Multiplying both sides of the last equation by $\mathbb{P}_t$ connects the real pricing kernel with the nominal pricing kernel via the definition

$$m^{\$}_{t+1} = m_{t+1}.\frac{\mathbb{P}_t}{\mathbb{P}_{t+1}},$$

whose logarithmic transformation is:

$$\ln m^{\$}_{t+1} = \ln m_{t+1} + \ln\frac{\mathbb{P}_t}{\mathbb{P}_{t+1}}. \tag{43}$$

Note that if the price level is increasing, the fraction $\mathbb{P}_t/\mathbb{P}_{t+1}$ is a number less than 1, which implies that $\ln(\mathbb{P}_t/\mathbb{P}_{t+1})$ is a negative number. Formula (43) therefore has an intuitive interpretation: It says that in the presence of inflation, nominal bonds should be more heavily discounted than real bonds, with the additional discount approximately equal to the inflation rate. This representation of the pricing kernel has the flavor of the Fisher relationship briefly introduced on page 19.

For subsequent algebraic manipulations, it is more convenient to represent equation (43) in the equivalent form:

$$\ln m^{\$}_{t+1} = \ln m_{t+1} - \ln(1+\Pi_{t+1}), \tag{44}$$

in which, for consistency with previous notation, $1 + \Pi_{t+1} = \frac{\mathbb{P}_{t+1}}{\mathbb{P}_t}$.

The expression suggests that if the logarithms of the stochastic real pricing kernel m_{t+1} and the gross inflation component $(1 + \Pi_{t+1})$ can be represented as linear functions of state variables or factors, the nominal pricing kernel will inherit that property as well. This additivity is one feature of affine factor models that accounts for their tractability because it makes it possible to infer the inflation-expectations component imbedded in bond prices with the simple algebraic operation of subtracting real forward rates from nominal ones.

The Real Pricing Kernel. A first step in disentangling nominal from real variables is to start out with a pricing kernel in the class of affine factor models that can price government bonds in real terms with the asset-pricing formula (42). One such specification of the stochastic real pricing kernel proposed in the literature that will be particularly useful for the purpose of this article is:

$$-\ln m_{t+1} = \delta + \gamma_c \, \ln \frac{c_t}{c_{t-1}} + \gamma_v \, v_t + \lambda_c \sigma_c \sqrt{v_t}\, \varepsilon_{c,\,t+1} + \lambda_v \sigma_v \varepsilon_{v,\,t+1}, \quad (45)$$

where c_t stands for aggregate real consumption and c_t/c_{t-1} for aggregate real consumption growth, v_t is a stochastic variable to be defined later, σ_c and σ_v are strictly positive scalars, and $\varepsilon_{c,\,t+1}$ and $\varepsilon_{v,\,t+1}$ are identically and independently distributed normal random variables with mean 0 and variance equal to 1.[17]

In light of the interpretation of expression (18) offered earlier, it follows that this particular two-factor pricing kernel generates time-varying risk premia because it sets one of the scalars $\beta_{h,l}$ in the radicand of that expression, $\beta_{1,2}$, to a strictly positive value. That implication is given away by the fact that, in contrast with the one-factor model studied in the previous section, one of the unforecastable errors, $\varepsilon_{c,t+1}$, appears multiplied by the square root of the level of one of the factors in the model, v_t.

Although this two-factor real pricing kernel may appear arbitrary, it can be derived, as demonstrated by Gallmeyer et al. (2007), from consumers' choice theory in the spirit of the straight-engineering approach discussed in Appendix A. Moreover, those authors show that the pricing kernel inherits its square-root element from the following stochastic process with which they propose to capture the behavior of real consumption growth in the data:

$$\ln \frac{c_{t+1}}{c_t} = (1 - \phi_c)\theta_c + \phi_c \ln \frac{c_t}{c_{t-1}} + \sigma_c \sqrt{v_t}\, \varepsilon_{c,t+1}, \quad (46)$$

where $0 < \phi_c < 1$ is a coefficient that captures the persistence of real consumption growth, θ_c is the long-run average of that variable (i.e.,

[17] Notice that this pricing kernel can be obtained from the general affine expression (18) after multiplying both sides of that expression by -1 and setting $\Psi = \delta$, $q = 2$, $\gamma_1 = \gamma_c$, $F_1 = \ln \frac{c_t}{c_{t-1}}$, $\gamma_2 = \gamma_v$, $F_2 = v_t$, $\lambda_1 = \lambda_c$, $\alpha_1 = 0$, $\beta_{1,1} = 0$, $\beta_{1,2} = \sigma_c^2$, $\varepsilon_{1,t+1} = \varepsilon_{c,t+1}$, $\lambda_2 = \lambda_v$, $\alpha_2 = \sigma_v^2$, $\beta_{2,1} = 0$, $\beta_{2,2} = 0$, and $\varepsilon_{2,t+1} = \varepsilon_{v,t+1}$. Thus, the stochastic pricing kernel (45) is a two-factor affine model with factors $\ln \frac{c_t}{c_{t-1}}$ and v_t, factor loadings γ_c and γ_v, and prices of risk λ_c and λ_v.

$E[\ln \frac{c_{t+1}}{c_t}] = \theta_c)$, the innovation $\varepsilon_{c,t+1}$ is an identically and independently distributed normal random variable with mean 0 and variance equal to 1, and v_t is a stochastic variable whose dynamics is governed by the following univariate autoregressive process:

$$v_{t+1} = (1 - \phi_v)\theta_v + \phi_v v_t + \sigma_v \varepsilon_{v,t+1}, \tag{47}$$

where $0 < \phi_v < 1$, $\theta_v = E[v_t]$, and $\varepsilon_{v,t+1}$ is, as usual, an identically and independently distributed normal random variable with mean 0 and variance equal to 1.

Notice that the specification for real consumption growth implies that its conditional variance, $Var_t[\ln \frac{c_{t+1}}{c_t}] = E_t[\ln \frac{c_{t+1}}{c_t} - E_t \ln \frac{c_{t+1}}{c_t}]^2 = \sigma_c^2 v_t$, is not constant (homoskedastic) but varies over time with the level of another variable (heteroskedastic). It is this feature of the pricing kernel that introduces time-varying risk premia into government bond prices.[18]

But that is not the only reason for appealing to a square-root specification for the unforecastable error in the stochastic process governing real consumption growth. A more important consideration is empirical in nature: The evidence extensively documented by Bansal and Yaron (2004) strongly suggests that the conditional variance of real consumption growth is time varying.

The Dynamics of the Inflation Rate. As hinted at by the discussion following expression (44), all it takes to go from the real pricing kernel to the nominal one is to specify a stochastic process for the inflation rate that is also a linear function of state variables or factors. Following Gallmeyer et al. (2007), the dynamics of the inflation rate is assumed to be given by the stochastic process:

$$\ln(1 + \Pi_{t+1}) = \bar{\pi} + \pi_c \ln \frac{c_{t+1}}{c_t} + \pi_v v_{t+1} + \pi_s s_{t+1}, \tag{48}$$

where $\bar{\pi}$, π_c, π_v, and π_s are constant coefficients and s_t is a stochastic nominal variable with autoregressive representation:

$$s_{t+1} = \phi_s s_t + \sigma_s \varepsilon_{s,t+1}, \tag{49}$$

where $0 < \phi_s < 1$ and $\varepsilon_{s,t}$ is, as usual, an identically and independently distributed normal random variable with mean 0 and variance equal to 1. Conceptually, s_t captures the nonsystematic component of monetary policy; that is, unanticipated policy moves.

Notice that, according to this specification, positive monetary policy surprises cancel out negative ones on average; that is, $E[s_t] = 0$. As a

[18] In fact, if the conditional variance of real consumption growth were constant over time, the real pricing kernel would collapse to a one-factor model, with real consumption growth the only factor. With some violation of rigor, this assertion can be intuitively verified by setting $v_t = 1$ for all t in (45). The proper procedure, however, is to follow the same steps as in Gallmeyer et al. (2007) and show that the pricing kernel for the homoskedastic case is a version of (45) with $\gamma_v = \lambda_v = 0$, $\nu_t = 1$, and $\delta = \overset{\star}{\delta}$, a different constant.

result, such surprises have no impact on the long-run inflation rate given by:

$$E[\ln(1+\Pi_t)] = \bar{\pi} + \pi_c\theta_c + \pi_v\theta_v. \tag{50}$$

Another feature of the stochastic process characterizing monetary policy is that its conditional variance, $Var_t[s_{t+1}] = E_t[s_{t+1} - E_t(s_t)]^2 = \sigma_s^2$, is homoskedastic.[19]

Taking into account (46), (47), and (49), expression (48) can be rewritten:

$$\begin{aligned} \ln(1+\Pi_{t+1}) &= \bar{\pi} + \pi_c(1-\phi_c)\theta_c + \pi_v(1-\phi_v)\theta_v + \\ &\quad \pi_c\phi_c \ln\frac{c_t}{c_{t-1}} + \pi_v\phi_v v_t + \pi_s\phi_s s_t + \\ &\quad \pi_c\sigma_c\sqrt{v_t}\,\varepsilon_{c,t+1} + \pi_v\sigma_v\varepsilon_{v,t+1} + \pi_s\sigma_s\varepsilon_{s,t+1}. \end{aligned} \tag{51}$$

The Nominal Pricing Kernel. For the reasons given in the subsection "Disentangling the Nominal and Real Pricing Kernels" on page 26, it is legitimate to add the real pricing kernel (45) to the inflation process (51) to obtain a three-factor-model representation of the nominal pricing kernel:

$$\begin{aligned} -\ln m^{\$}_{t+1} &= -\ln m_{t+1} + \ln(1+\Pi_{t+1}) \\ &= \delta + \bar{\pi} + \pi_c(1-\phi_c)\theta_c + \pi_v(1-\phi_v)\theta_v + \\ &\quad (\gamma_c + \pi_c\phi_c)\ln\frac{c_t}{c_{t-1}} + (\gamma_v + \pi_v\phi_v)v_t + \pi_s\phi_s s_t + \\ &\quad (\lambda_c + \pi_c)\sigma_c\sqrt{v_t}\,\varepsilon_{c,t+1} + (\lambda_v + \pi_v)\sigma_v\varepsilon_{v,t+1} + \pi_s\sigma_s\varepsilon_{s,t+1} \\ &= E[\ln(1+\Pi_t)] + \delta - \phi_c\pi_c\theta_c - \phi_v\pi_v\theta_v + \\ &\quad (\gamma_c + \pi_c\phi_c)\ln\frac{c_t}{c_{t-1}} + (\gamma_v + \pi_v\phi_v)v_t + \pi_s\phi_s s_t + \\ &\quad (\lambda_c + \pi_c)\sigma_c\sqrt{v_t}\,\varepsilon_{c,t+1} + (\lambda_v + \pi_v)\sigma_v\varepsilon_{v,t+1} + \pi_s\sigma_s\varepsilon_{s,t+1}. \end{aligned} \tag{52}$$

Several features of this nominal pricing kernel are worth noticing. First, the dependence of that kernel on long-run inflation expectations, represented by the term $E[\ln(1+\Pi_t)]$, has been made explicit by plugging (50) into the second equality in (52). Second, the source of time-varying risk premia in this three-factor model, as shown later, is the term containing $\sqrt{v_t}$. Third, the nominal variable s_t, ultimately responsible for capturing the effects of monetary policy on the short-term fluctuations of inflation around its long-run mean, turns the two-factor model of the real pricing kernel into a three-factor model of the nominal pricing kernel.

[19] Although readers may feel uncomfortable with the seemingly arbitrary assumption that the inflation rate follows the stochastic process (48), again Gallmeyer et al. (2007) show that this dynamics emerges endogenously in an endowment economy characterized by the real pricing kernel (45) and where the monetary policy instrument—the short-term interest rate—is set to the level determined by a Taylor-like rule. The systematic component of that rule is a linear function of the logarithm of real consumption growth, $\ln\frac{c_t}{c_{t-1}}$, and of the logarithm of the inflation rate, $\ln\frac{\mathbb{P}_{t+1}}{\mathbb{P}_t}$.

Nominal Government Bond Prices. Equipped with the nominal and real pricing kernels, it is possible to compute the nominal and real bond prices for all maturities with the pricing formulas (40) and (42), respectively, through the series of tedious steps outlined at the beginning of the section "A One-Factor Model of the Yield Curve" on page 14. Fortunately, given the affine structure of both kernels and the lognormal distribution of the relevant variables, the more handy recursive formula presented at the end of that section is valid as well. Thus, nominal government bond prices can be readily computed with a more general version of formula (28):

$$\begin{aligned} -\ln P_t^{(n)} &= A_n + B_{c,n}\left(\ln \frac{c_t}{c_{t-1}} - \theta_c\right) + B_{v,n}(v_t - \theta_v) + B_{s,n}s_t + \\ & D_{c,n}Var_t[\ln \frac{c_{t+1}}{c_t}] + D_{v,n}Var_t[v_{t+1}] + D_{s,n}Var_t[s_{t+1}], \end{aligned} \quad (53)$$

whose coefficients, following the steps outlined in the section just mentioned, can be calculated sequentially from the recursions:

$$\begin{aligned} A_n &= A_{n-1} + \bar{\pi} + \pi_c\theta_c + \pi_v\theta_v + \delta + \gamma_c\theta_c + \\ & \left[\gamma_v + (1-\phi_v)\sigma_c^2 C_{c,n-1}\right]\theta_v; \\ B_{c,n} &= B_{c,n-1} + \phi_c^{n-1}(\gamma_c + \phi_c\pi_c); \\ B_{v,n} &= B_{v,n-1} + \phi_v^{n-1}(\gamma_v + \phi_v\pi_v); \\ B_{s,n} &= B_{s,n-1} + \phi_s^{n-1}\phi_s\pi_s; \\ D_{c,n} &= \phi_v D_{c,n-1} - \frac{1}{2}(\lambda_c + \pi_c + B_{c,n-1})^2; \\ D_{v,n} &= D_{v,n-1} - \frac{1}{2}(\lambda_v + \pi_v + B_{v,n-1} + \sigma_c^2 D_{c,n-1})^2; \\ D_{s,n} &= D_{s,n-1} - \frac{1}{2}(\pi_s + B_{s,n-1})^2, \end{aligned} \quad (54)$$

with initial conditions $A_0 = B_{c,0} = B_{v,0} = B_{s,0} = D_{c,0} = D_{v,0} = D_{s,0} = 0$ implied by the restriction that the price of the one-period bond generated by formula (53) is equal to that generated by the basic formula (41).[20]

[20] Readers with training in finance might be tempted to compute the nominal price of a bond of a given maturity n with the theoretically equivalent multi-period version of formula (40) that can be obtained after repeatedly substituting the prices that pop up sequentially in the right-hand side of that equation for replicas of that same right-hand side appropriately shifted in time; that is, with the formula $P_t^{(n)} = E_t[m_{t+1}^{\$} m_{t+2}^{\$} m_{t+3}^{\$} ... m_{t+n}^{\$}]$. However, the process of computing prices with this alternative formula is not just tedious but analytically intractable because under the square-root specification for the standard deviation of real consumption growth assumed in (46), the factor $\ln \frac{c_{t+1}}{c_t}$ and the log of the nominal pricing kernel, $\ln \mathrm{m}_{t+1}^{\$}$, are normally distributed *conditional* on $\ln \frac{c_t}{c_{t-1}}$, but that ceases to be the case for $\ln \frac{c_{t+i}}{c_{t+i-1}}$ and $\ln m_{t+i}^{\$}$ when i is greater than 1.

Notice that the formal structure of this price formula—like that for the one-factor model studied earlier—groups terms in three broad categories: a constant, those associated with the short-term dynamics induced by deviations of each factor from its corresponding long-run mean (recall that $E[\ln \frac{c_t}{c_{t-1}}] = \theta_c, E[v_t] = \theta_v$, and $E[s_t] = 0$), and those associated with the conditional variances of the factors.

Inspection of the recursions (54) reveals that in the latter category, only the coefficient corresponding to the variance of the real consumption growth factor, $D_{c,n}$, doesn't follow the pattern of the comparable coefficient C_n in the one-factor model with constant risk premia studied from page 14 onward (see expression 30). The anomaly is that the incremental risk premium between successive maturities attributable to the consumption factor, captured by the difference $D_{c,n} - D_{c,n-1}$, is not independent from the coefficient for the shorter maturity in that differential, $D_{c,n-1}$. That is, each element $D_{c,n}$ in the sequence of coefficients $\{D_{c,n}\}$ is obtained by adding to the immediately preceding element the amount $(\phi_v - 1)D_{c,n-1} - \frac{1}{2}(\lambda_c + \pi_c + B_{c,n-1})^2$, an increment not independent from the previous element in the sequence $D_{c,n-1}$. By contrast, the coefficients in the sequence $\{B_{c,n}\}$ are augmented from one maturity to the next by $\phi_c^{n-1}(\gamma_c + \phi_c\pi_c)$, an increment independent from the coefficient corresponding to the immediately preceding maturity $B_{c,n-1}$.

The out-of-pattern behavior in the coefficients associated with the conditional variance of consumption, $D_{c,n}$, is not by chance. It has to do with the fact that real consumption growth is the only one of the three factors in the specific model studied in this section whose conditional variance changes over time (recall that $Var_t[\ln \frac{c_{t+1}}{c_t}] = \sigma_c^2 v_t$). It is through this real consumption growth factor that time-varying risk premia enter into forward rates, potentially creating havoc with the readings of inflation-expectations indicators constructed using those rates.

Real Government Bond Prices. The preceding section suggests that the same series of steps invoked there could be applied to produce a recursive formula that prices government bonds in real terms using the factors that enter into the pricing kernel (45). Indeed, that is the case, with the resulting formula given by:

$$\begin{aligned} -\ln p_t^{(n)} &= a_n + b_{c,n}(\ln \frac{c_t}{c_{t-1}} - \theta_c) + b_{v,n}(v_t - \theta_v) + \\ & \quad d_{c,n} Var_t[\ln \frac{c_{t+1}}{c_t}] + d_{v,n} Var_t[v_{t+1}], \end{aligned} \tag{55}$$

where $p_t^{(n)}$ stands for the price of a government bond that delivers one unit of the consumption good (rather than \$1) at its expiration at time $t+n$ and where the coefficients can be determined recursively from the following expressions:

$$
\begin{aligned}
a_n &= a_{n-1} + \delta + \gamma_c \theta_c + \left[\gamma_v + (1-\phi_v)\sigma_c^2 d_{c,n-1} \right] \theta_v; \\
b_{c,n} &= b_{c,n-1} + \phi_c^{n-1} \gamma_c; \\
b_{v,n} &= b_{v,n-1} + \phi_v^{n-1} \gamma_v; \\
d_{c,n} &= \phi_v d_{c,n-1} - \frac{1}{2}(\lambda_c + b_{c,n-1})^2; \\
d_{v,n} &= d_{v,n-1} - \frac{1}{2}(\lambda_v + b_{v,n-1} + \sigma_c^2 d_{c,n-1})^2, \qquad (56)
\end{aligned}
$$

with initial conditions $a_0 = b_{c,0} = b_{v,0} = d_{c,0} = d_{v,0} = 0$ implied by the restriction that the real price of the one-period bond generated by the formula must be the same as that generated with the basic real bond pricing formula $p_t^{(1)} = E_t[m_{t+1}]$. Recall that by the no-arbitrage condition, the real price of a bond issued at $t+1$ that pays one unit of the consumption good in that same period, $p_{t+1}^{(0)}$, must be 1.[21]

Because no nominal variables enter into the real pricing kernel, real bond prices depend on just the two real factors present in the two-factor pricing kernel (45): real consumption growth and its volatility.[22] Given the dependence of this pricing formula on the factor captured by real consumption growth, it is not surprising that the corresponding coefficient, $d_{c,n}$, displays the same anomalous pattern detected in its nominal counterpart, $D_{c,n}$, and from the same source: a variance of that factor that changes over time.

The formula for the real price of government bonds completes the list of elements needed to present the theory behind the construction of inflation-expectations indicators like the one in Figure 1.

The Nominal–Real Forward Rate Spread and Its Components. Inspection of the recursions (54) for the nominal price formula reveals that they contain parameters associated with nominal as well as real factors, while the analogous recursions (56) for the real price formula contain parameters related just to real factors. Thus, it seems legitimate to conjecture that subtracting the coefficients pertaining to real prices from those corresponding to nominal prices will produce the desired result of isolating the nominal factors—long-run inflation expectations among them.

In fact, that is the attempt behind popular inflation-expectations indicators, typically calculated as the following nominal–real forward rate spread:

---

[21] The same warning as in footnote 20 applies to the attempt of computing government bond prices in real terms with the iterative pricing formula $p_t^{(n)} = E_t[m_{t+1}m_{t+2}m_{t+3}...m_{t+n}]$.

[22] This is, of course, a consequence of the independence of the real stochastic pricing kernel from inflation postulated in "Disentangling the Nominal and Real Pricing Kernels" on page 26.

$$\ln F\bar{N}_t^{(n)\to(n+j)} - \ln F\bar{R}_t^{(n)\to(n+j)} = \frac{\ln \bar{P}_t^{(n)} - \ln \bar{P}_t^{(n+j)} - [\ln \bar{p}_t^{(n)} - \ln \bar{p}_t^{(n+j)}]}{j}, \quad (57)$$

where $\ln F\bar{N}_t^{(n)\to(n+j)}$ is the forward nominal rate defined in (37) calculated from actually observed nominal prices $\bar{P}_t^{(n)}$, identified by a bar over the corresponding symbol; $\ln F\bar{R}_t^{(n)\to(n+j)}$ is the analogous concept calculated from actually observed real (inflation–indexed) government bond prices $\bar{p}_t^{(n)}$; and n denotes a long maturity, typically five years or more.

Whether the calculations above deliver as intended an indicator that effectively and reliably traces the evolution of long-run inflation expectations over time will have to be established by reading that indicator under the light of its theoretical counterpart, which for $j = 1$ in the case of the three-factor model with time-varying volatility is given by the following expression:

$$\begin{aligned}
\ln F\hat{N}_t^{(n)\to(n+1)} - \ln F\hat{R}_t^{(n)\to(n+1)} &= \ln \hat{P}_t^{(n)} - \ln \hat{P}_t^{(n+1)} - \\
&\quad (\ln \hat{p}_t^{(n)} - \ln \hat{p}_t^{(n+1)}) \\
&= A_{n+1} - A_n - (a_{n+1} - a_n) + \\
&\quad \Big[B_{c,n+1} - B_{c,n} - (b_{c,n+1} - b_{c,n})\Big]\Big(\ln \frac{c_t}{c_{t-1}} - \theta_c\Big) + \\
&\quad \Big[B_{v,n+1} - B_{v,n} - (b_{v,n+1} - b_{v,n})\Big](v_t - \theta_v) + \\
&\quad (B_{s,n+1} - B_{s,n}) s_t + \\
&\quad \Big[D_{v,n+1} - D_{v,n} - (d_{v,n+1} - d_{v,n})\Big] Var_t[v_{t+1}] + \\
&\quad (D_{s,n+1} - D_{s,n}) Var_t[s_{t+1}] + \\
&\quad \Big[D_{c,n+1} - D_{c,n} - (d_{c,n+1} - d_{c,n})\Big] Var_t[\ln \frac{c_{t+1}}{c_t}],
\end{aligned}$$

where the hats over the symbols identifying forward rates and prices indicate variables implied by the model rather than obtained from the data, and the second equality follows from straight application of the theoretical price formulas (53) and (55).

Substituting the recursions (54) and (56) in the expression and working through considerable algebra results in the following alternative representation of the nominal–real forward rate spread:

$$\begin{aligned}
&\ln F\hat{N}_t^{(n)\to(n+1)} - \ln F\hat{R}_t^{(n)\to(n+1)} = E[\ln(1 + \Pi_t)] + \\
&\phi_c^{n+1} \pi_c (\ln \frac{c_{t+1}}{c_t} - \theta_c) + \phi_v^{n+1} \pi_v (v_t - \theta_v) + \phi_s^{n+1} \pi_s \, s_t - \\
&\Omega_{v,n} - \Omega_{s,n} - \Omega_{c,n} \, v_t, \quad (58)
\end{aligned}$$

where the coefficients $\Omega_{i,n}$, $i = v, s, c$, are complicated functions of the fundamental parameters that appear in the recursions (54) and (56). The explicit derivation of those functions can be found in Appendix B.

Given that the autoregressive coefficients ϕ_c, ϕ_v, and ϕ_s are assumed to be less than 1 in absolute value, the second line in this expression can be ignored for long maturities, based on the same arguments made in the discussion preceding the approximate formula (39). The nominal–real forward rate spread can be approximated with the expression:

$$\ln F\hat{N}_t^{(n)\to(n+1)} - \ln F\hat{R}_t^{(n)\to(n+1)} \approx E[\ln(1+\Pi_t)] - \Omega_{v,n} - \Omega_{s,n} - \Omega_{c,n}\, v_t. \tag{59}$$

As already mentioned, two differences stand out between the theoretical nominal forward rate (39) derived for the one-factor model studied on page 14 and the theoretical nominal–real forward rate spread above: The latter is not contaminated by a long-run mean for the real interest rate, and—more critically for the message of this article—its risk premium component, captured by the last three terms in (59), is not constant but moves over time with the level of the factor v_t.

To be precise, movements in the factor v_t will induce in the nominal–real forward rate spread (58) fluctuations on the order of magnitude of $\Omega_{c,n}\, v_t$, the size of the term capturing time-varying risk premia on the right-hand side of that expression. Importantly, that term will be the *only* quantitatively significant source of such fluctuations for long maturities, as is apparent in expression (59), the approximate version of the exact nominal–real forward rate spread. Thus, time-varying risk premia could be largely responsible for the fluctuations in popular inflation-expectations indicators, which are nothing but the empirical counterparts of the nominal–real forward rate spreads derived here. This is precisely the conjecture rigorously examined in the next section.

4. THE NOMINAL–REAL FORWARD RATE SPREAD AND INFLATION-EXPECTATIONS INDICATORS

The popular long-run inflation-expectations indicator typically included in financial and economic reports is obtained by calculating a time series for expression (57) with actually observed zero-coupon prices for five- and ten-year nominal and inflation-adjusted government bonds.[23] The question

[23] The observations in Figure 1 correspond to quarterly averages of daily nominal–real forward rate spreads calculated from zero-coupon prices between March 1999 and March 2009. For the U.S., data for nominal bond yields correspond to series SVENY10 (ten-year nominal bond) and SVENY05 (five-year nominal bond) in Gürkaynak, Sack, and Wright (2006), and data for real bond yields correspond to series TIPSY10 (ten-year real bond) and TIPSY05 (five-year real bond) in Gürkaynak, Sack, and Wright (2008). For the U.K., the data were extracted from the Bank of England Interactive Database in www.bankofengland.co.uk/mfsd/iadb/NewIntermed.asp and correspond to series reported under Interest and Exchange Rates/Yields/Zero Coupon Yields, specifically: IUDMNZC (ten-year nominal yield), IUDSNZC (five-year nominal yield), IUDMRZC (ten-year real yield), and IUDSRZC (five-year real yield). The zero-coupon qualification is important because the forward-rate-based inflation-expectations indicators reported in the press are typically calculated instead from constant maturity bonds which, unlike zero-coupon bonds, are contaminated by differences in interest payments.

raised at the beginning of this article was whether the fluctuations typically exhibited by this indicator reflect movements in long-run inflation expectations or in something else. The purpose of this section is to establish the extent to which time-varying risk premia are quantitatively large enough to account for the movements that inflation-expectations indicators show in the data. That is accomplished with numerical simulations of the three-factor model laid out in the previous sections. The simulations involve the following steps:

1. Choose values for the relevant primitive parameters in the pricing kernel (45) and stochastic processes (46), (47), (48), and (49). To add realism to the exercise, the parameter values are set equal to those selected or implied by the study in Gallmeyer et al. (2007), whose choices were based on other empirical studies and the time-series behavior of the appropriate macroeconomic aggregates, such as real consumption growth (the parameter values are reported in Appendix C).[24]

2. Create an artificial time series for the three unforecastable errors in the model, $\{\varepsilon_{c,t}\}$, $\{\varepsilon_{v,t}\}$, and $\{\varepsilon_{s,t}\}$.

3. Do the same for the three factors in the model using expressions (46), (47), and (49), the time series for the unforecastable errors from the previous step, and initial (time 0) values for the factors equal to their long-run means (with the implication that at time 1, the factors take on values given by the expressions $\ln\frac{c_1}{c_0} = \theta_c + \sigma_c\sqrt{\theta_v}\varepsilon_{c,1}$, $v_1 = \theta_v + \sigma_v\varepsilon_{v,1}$, and $s_1 = \sigma_s\varepsilon_{s,1}$).

4. Use price formulas (53) and (55), recursions (54) and (56), and the artificial time series for the factors obtained in the previous step to calculate artificial time series for the logarithm of nominal and real prices corresponding to zero-coupon government bonds with five and ten years to maturity.

[24]Two parameters of quantitatively significant importance for the simulations reported later, ϕ_s and σ_c^2, were not set to the values chosen or implied by Gallmeyer et al. (2007). The parameter ϕ_s was set equal to 0.85 instead of the 0.922 those authors used, in order to minimize the relatively large fluctuations in the simulated inflation-expectations indicator induced by the term in (58) associated with deviations of the factor s_t from its mean, and to focus the attention on the fluctuations induced by the time-varying risk premium component of that expression. The value for the parameter σ_c^2 was picked with a different procedure from that used by Gallmeyer et al. Those authors derived their theoretical pricing kernel under the assumption $\sigma_c^2 = 1$ ($\sigma_x^2 = 1$ in the authors' notation). In the quantitative exercise section, however, they report to have used a different value given by the formula $Var[\ln\frac{c_{t+1}}{c_t}](1-\phi_c^2)$, with $Var[\ln\frac{c_{t+1}}{c_t}]$ set equal to 0.000023, the sample variance of detrended quarterly consumption growth over the period 1952:2–2005:4 (which implies, of course, $\left[Var[\ln\frac{c_{t+1}}{c_t}]\right]^{\frac{1}{2}} = 0.0048$). That expression is consistent with the formula for the unconditional variance of consumption growth that would be obtained in the standard homoskedastic case, $Var[\ln\frac{c_{t+1}}{c_t}] = \frac{\sigma_c^2}{1-\phi_c^2}$. Given that the time-varying risk premia case is the one of interest for the purpose of this article, the value of the parameter σ_c^2 was set instead to that implied by the formula for the unconditional variance of real consumption growth that would be obtained under heteroskedasticity, $Var[\ln\frac{c_{t+1}}{c_t}] = E[\ln\frac{c_{t+1}}{c_t} - E\ln\frac{c_{t+1}}{c_t}]^2 = \frac{\sigma_c^2}{1-\phi_c^2}\theta_v$.

5. Calculate the theoretical nominal–real forward rate spread with expression (57) for $n = 5$ and $j = 5$, using the artificial time series for prices generated in the previous step rather than actually observed prices.

The outcome of the last step, a time series for the five-year nominal–real forward rate spread, is reported in Figure 4.

Figure 4: Simulated Nominal–Real Forward Rate Spread (Time-Varying Risk Premia Case)

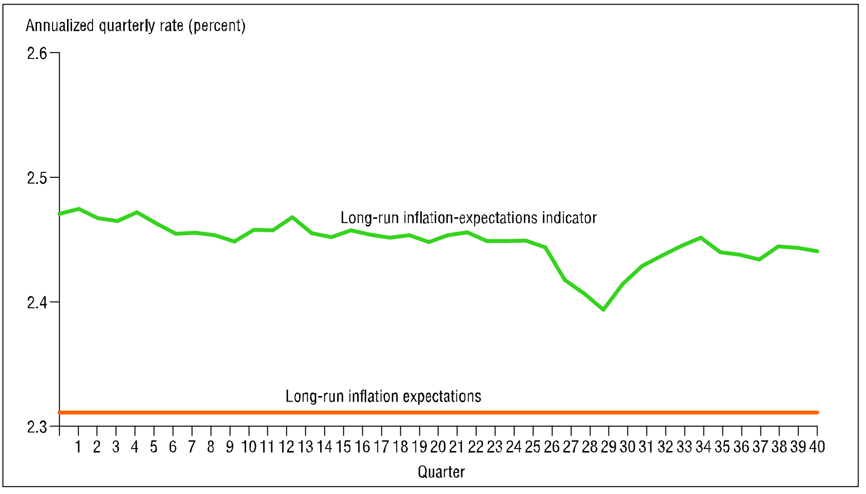

The plot is nothing but the theoretical counterpart of the empirical inflation-expectations indicator calculated from actually observed prices. The important insight from this plot is that the synthetic long-run inflation-expectations indicator fluctuates even if long-run inflation expectations do not because, by construction, long-run inflation expectations stay firmly anchored at the 2.31 percent annual inflation target implied by the parameter values adopted for the inflation process (48).[25] This implies that the larger-than-normal drop in the synthetic indicator toward the end of the artificial sample is the result of time-varying risk premia induced by an abnormally low realization of the factor v_t, which can be seen from looking at the approximate expression for the nominal–real forward rate spread (59). Thus, the simulation of the three-factor model serves as a warning against interpreting fluctuations in empirical nominal–real forward rates as evidence of a change of equal size in long-run inflation expectations.

But that warning must be balanced by some obvious limitations of the quantitative exercise summarized by Figure 4. First, as already noted, the artificially generated time series of the factors includes realizations of v_t that exceed four standard errors and are, therefore, extremely rare. Were it not for those infrequent values, the theoretical nominal–forward rate spread would look most of the time like the smooth line in the first two-thirds of

[25] The sample mean of the artificial inflation-expectations indicator plotted in Figure 4 (top line) is slightly above the long-run inflation expectations underlying its calculation (bottom line) because the risk premia components left over in the nominal–real forward rate spread introduce a small upward bias in that indicator.

Figure 4. A more benign assessment might not find enough reasons to completely invalidate the empirical relevance of the exercise because, after all, the decline that the indicator in Figure 1 exhibits toward the end of 2008 and beginning of 2009 indeed took place under a rare circumstance—a financial turmoil of a geographical extension, depth, and intensity not seen since the Great Depression, by many accounts.

The issue of the frequency of unusual events does seem minor relative to that of orders of magnitude: the trough of the theoretical inflation-expectations indicator displayed in Figure 4 is just 0.08 percentage point below the peak of that indicator at the beginning of the artificial sample, whereas the comparable difference in the empirical counterpart of that spread in Figure 1 is about six times as large for the U.S. This inability to produce fluctuations in the nominal–real forward rate spread of a magnitude similar to those observed in the data is not limited to the three-factor model just simulated. It is a problem common to virtually all models of the term structure, as documented by Rudebusch and Swanson (2009): "The term premium on nominal long-term bonds in the standard dynamic stochastic general equilibrium (DSGE) model used in macroeconomics is far too small and stable relative to empirical measures obtained from the data—an example of the 'bond premium puzzle.' "

Judging from the numerical simulation of the particular three-factor model studied in this article, the puzzle applies as well to inflation-expectations indicators calculated with the forward-rates technique: The time-varying risk premia implied by observed macroeconomic aggregates and a monetary policy targeting a time-invariant long-run inflation rate are too small and stable to account for the volatility of those indicators. This quantitative finding seems to favor the alternative hypothesis that the large swings typically observed in inflation-expectations indicators are induced mostly by changes in long-run inflation expectations.

However, many policymakers and researchers are likely to dispute that conclusion on the grounds that it is hard to believe that long-run inflation expectations can go up or down by as much as 0.5 percentage point or more from one quarter to the next in countries with a track record of low and stable inflation, such as the U.S. and the U.K. A change in long-run inflation expectations of that magnitude would typically be associated with an actual or prospective change of monetary policy regime, as signaled by a modification to the legal status of the monetary authority, the replacement of the individuals ultimately in charge of monetary policy, or political upheaval.

In the case of the U.S., it is far from obvious that any of those contingencies were behind the 0.6 percentage-point jump that the indicator registered, according to Figure 1, in the brief period spanning the second and fourth quarters of 2003. Chairman Greenspan was at the helm of the Fed all that time, and nothing suggests that his preferred long-run inflation rate changed over those six months. Moreover, nothing indicated that he would be removed before the expiration of his mandate about two years later. Consistent with this view, inflation-expectations indicators based on surveys rather than forward rates didn't move between those two quarters. For example, according to the Federal Reserve Bank of Philadelphia's

Survey of Professional Forecasters, the long-term (ten-year) forecasts for the inflation rate remained firmly anchored at a 2.5 percent annual rate throughout 2003.

These considerations and similar ones in the literature suggest that the current understanding of the determinants of government bond prices is too limited to establish with any confidence which fraction of the relatively large variations in inflation-expectations indicators based on forward rates can be attributed to actual changes in long-run inflation expectations and which to time-varying risk premia. Unfortunately, chances are slim that the existing difficulties in interpreting inflation-expectations indicators constructed with the forward-rates technique will be resolved soon. The assessment by Campbell, Lo, and MacKinlay (1997, pg. 455) that "no one model has yet emerged as a consensus choice for modeling the term structure" continues to reflect the situation as accurately today as it did more than a decade ago.

5. CONCLUSION

Protecting the value of currency is a mission almost all societies entrust to their monetary authorities. For that reason, policymakers are always on the lookout for long-run inflation-expectations indicators that can provide early warnings when those expectations are about to become unanchored. The materialization of that prospect would be a nightmare for any central banker: Once the inflation-expectations genie is let out of the bottle, it will be hard to put it back in again without risking a recession.

This article discusses the challenges of constructing such long-term inflation-expectations indicators from available data on nominal and real government bond yields. A necessary step in the process is understanding exactly how long-term inflation expectations are priced into such bonds. To that end, the article adopts an asset-pricing perspective revealing that, under the lens of affine factor models, government yields can be decomposed into long-term inflation-expectations and risk premia components. The subsequent analysis focuses on the difficult task of identifying the former under different assumptions about the latter.

It turns out that producing indicators that can reliably track the evolution of long-run inflation expectations from one period to the next is fraught with significant theoretical and empirical challenges in the presence of time-varying risk premia, which seem to be an undisputed feature of the data. The disturbing property of risk premia that move around over time is that they can severely distort popular inflation-expectations indicators calculated from nominal–real forward rate spreads. As a result, such indicators could give the wrong impression that long-run inflation expectations have switched dangerously to a deflationary mood when, in reality, that is a mirage produced by declining risk premia. Yet in different economic circumstances, those same falling risk premia might mask a rise in long-term inflation expectations.

Gauging long-term inflation expectations correctly requires a rather precise quantification of the risk premia components of nominal and real interest rates. Unfortunately, the behavior of those components is still poorly understood. The state of the art in this subject has been aptly summarized by Federal Reserve Vice Chairman Donald Kohn (2005): "...the separation of market prices into distinct pieces reflecting expected values and risk is a difficult task that relies heavily on modeling assumptions about underlying processes and investor behavior."

The recognition that the art of eliciting inflation expectations from the yield curve is difficult suggests that some time will pass before many of the remaining theoretical and empirical issues relevant to the construction of reliable long-run inflation-expectations indicators are sorted out. In the meantime, policymakers are well advised not to attribute the relatively ample fluctuations observed in popular long-run inflation-expectations indicators to actual changes in those expectations. That interpretation would have some merit if risk premia were constant, but that view of the world seems no longer tenable given the overwhelming evidence accumulated by now against the empirical relevance of the expectations hypothesis. At the same time, policymakers should refrain from attributing all fluctuations in inflation-expectations indicators to time-varying risk premia. Although time-varying risk premia are a natural suspect, the quantitative evidence accumulated so far is simply too weak and circumstantial to lead to a verdict on that count beyond a reasonable doubt.

APPENDIXES

A. Selection of the Factors in Affine Factor Models

A Reverse-Engineering Approach. As readers will suspect, the reverse-engineering procedure for selecting the factors that characterize the stochastic properties and dynamics of the pricing kernel (18) in the paper is basically the same as the one described on page 10, but with a twist.

In the case of the one-factor model, the twist consists of replacing $\ln Y_t^{(1)}$ in expression (17) with an unspecified, or latent, factor x_t to obtain:

$$\ln m_{t+1} = -\frac{1}{2}\sigma_\varepsilon^2 - \gamma x_t - \lambda\varepsilon_{t+1}. \tag{60}$$

The basic idea is that the identity of the latent factor x_t can be subsequently inferred with econometric techniques that exploit the implications of (60) for government bond prices.

To see how this is accomplished, recall that ε_{t+1} is normally distributed, which implies that m_{t+1} is conditionally lognormally distributed. Therefore, the price for the one-period government bond is given by the expression:

$$P_t^{(1)} = E_t[m_{t+1}] = e^{E_t[\ln m_{t+1}]+\frac{1}{2}Var[\ln m_{t+1}]} = e^{-\frac{1}{2}\sigma_\varepsilon^2\ -\ \gamma x_t\ +\frac{1}{2}\lambda^2\sigma_\varepsilon^2}.$$

Taking logarithms on both sides and considering definition (3) results in the one-period government bond price (or yield):

$$\begin{aligned} -\ln Y_t^{(1)} &= \ln P_t^{(1)} = E_t\left[\ln m_{t+1}\right] + \frac{1}{2}E_t\Big[m_{t+1} - E_t[m_{t+1}]\Big]^2 \\ &= -\frac{1}{2}\sigma_\varepsilon^2 - \gamma x_t + \frac{1}{2}\lambda^2\sigma_\varepsilon^2. \end{aligned} \tag{61}$$

This theoretical relationship can be given empirical meaning by requiring that the unknown parameters and time series for the latent factor x_t on the right-hand side of the theoretical bond-pricing formula acquire values such that they are close, according to some criteria, to the time series of actually observed one-period government bond prices over the period under analysis. Many auxiliary assumptions about the stochastic properties of the factor x_t , as well as computational and econometric techniques such as the Kalman filter, are involved in the process. However, a detailed discussion of them is beyond the scope of this paper.

More important for readers is to gain intuition of the logic behind these techniques by considering the hypothetical situation in which application of these procedures results in point estimates of 1 for the unknown parameters γ and λ. Suppose that $\hat{\gamma} = 1$ and $\hat{\lambda} = 1$. After substitution of these point estimates in (61), it is revealed that the time series for the latent factor that would generate one-period bond prices (in logs) identical to the actual ones is the observed time series of one-period government bond prices itself. That is, the latent factor x_t underlying the pricing kernel (60) is the log of the one-period government bond price $\ln Y_t^{(1)}$!

Of course, this naive example will never be encountered in practice, and reverse engineering the latent factors that enter into the pricing kernel from observed variables will require the application of fairly advanced quantitative and econometric techniques. In fact, a large body of the empirical financial literature revolves around the issue of reverse engineering pricing kernels that match the observed time-series and cross-section features of government bond prices better than existing pricing kernels do. By contrast, the straight-engineering approach relies mostly on axioms and results from choice theory to decide which factors should enter into the pricing kernel, as discussed next.

A Straight-Engineering Approach. The straight-engineering approach starts from the theoretical premise that households don't value assets for their own sake but for the consumption goods that the payoffs from those assets allow them to acquire. Therefore, households must discount asset payoffs according to their implicit consumption value. This approach requires being specific about how households value consumption streams. In particular, it requires the specification of a function that measures the satisfaction, or utility, that the representative household gets out of a particular consumption bundle.

For example, it is frequently assumed that a household derives utility $U(c_t)$ from consuming c_t units of a composite consumption good at time t and $\beta U(c_{t+1})$ from consuming c_{t+1} units of that same good, but at time $t+1$, where $0 < \beta < 1$. At time t, such a household will typically face the choice of allocating one extra dollar of its budget to time t consumption, or of buying, at the certain price P_t, additional units of an asset that will pay stochastic dividends of x_{t+1} and be worth the uncertain price P_{t+1} at time $t+1$.

The utility value that the household gets out of these two choices must be the same; otherwise, the household could improve its utility by consuming more or saving more, as the case may be, until indifference between the options is restored. Formally, the solution to this decision problem is represented by the following equilibrium condition:

$$P_t U'(c_t) = E_t\left[\beta U'(c_{t+1})(x_{t+1} + P_{t+1})\right],$$

where $U'(c_t)$ measures the additional utility that a household would get from an additional time t unit of consumption. Thus, the left-hand side in the equation is a measure of the utility that the representative household would give up if it purchased an asset at a price P_t instead of consumption goods for the same amount. This utility loss must be exactly offset by the utility gains that the household expects to enjoy next period by selling the asset at the price P_{t+1} and using the proceeds, along with the dividends x_{t+1} paid by the asset, to purchase consumption goods then.

Dividing both sides of the equation by the marginal utility of consumption at time t, $U'(c_t)$, results in the following expression:

$$P_t = E_t\left[\beta \frac{U'(c_{t+1})}{U'(c_t)}(x_{t+1} + P_{t+1})\right]. \tag{62}$$

Comparison of this equation with (4) in the paper reveals the identity of the pricing kernel as:

$$m_{t+1} = \beta \frac{U'(c_{t+1})}{U'(c_t)}.$$

Notice that under the common and sensible assumption that marginal utility is decreasing in consumption—that is, that households don't derive as much utility from the second unit of a consumption item as from the first—the pricing kernel will be particularly high at times when c_{t+1} is particularly low relative to c_t (when consumption growth, $\frac{c_{t+1}}{c_t}$, is particularly low and vice-versa). This implies that households are willing to pay an insurance premium for assets that offer high payoffs in periods in which consumption growth is expected to be relatively low and the utility from buying additional goods is, therefore, particularly high. Conversely, they'll demand a risk premium on the return of assets that offer low payoffs in low consumption-growth states, precisely when households need high payoffs to boost consumption.

This can be seen more clearly by adopting the common assumption in the literature that the level of utility the representative household derives from consumption can be captured by the following power function:

$$U(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma},$$

for $\gamma > 0$, in which case the pricing kernel formula takes the form:

$$m_{t+1} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma}.$$

Taking logarithms of both sides of the expression produces the following one-factor representation of the pricing kernel:

$$\ln m_{t+1} = \ln \beta - \gamma \ln x_t + \xi_{t+1}, \tag{63}$$

where the factor x_t is consumption growth and where ξ_{t+1}, an identically and independently normally distributed unforecastable error orthogonal to $\ln c_{t+1}/c_t$, has been introduced to capture the stochastic nature of the pricing kernel.

Comparison of (60) and (63) reveals that the choice-theory-based straight-engineering approach rationalizes the reverse-engineering approach assumption that the stochastic pricing kernel can be represented by a function linear in state variables plus an innovation orthogonal to them. Each approach has advantages and disadvantages extensively discussed in the debate motivated by the so-called equity premium puzzle that Mehra and Prescott (1985) brought to the attention of the profession. (See Cochrane 2005 for a good introductory overview of the debate.)

B. Derivation of the Nominal–Real Forward Rate Spread

This appendix presents the sequence of algebraic steps that led to the compact formula (58) for the nominal–real forward rate spread as well as the explicit formulas representing the coefficients in that expression as functions of the fundamental parameters of the model.

The first step in the process is to represent the nominal and real forward rates in terms of fundamental parameters. For the one-period-ahead nominal forward rate, this entails substituting the recursions (54) in the appropriate places in definition (37) with $j = 1$:

$$
\begin{aligned}
\ln F\hat{N}_t^{(n)\to(n+1)} &= \ln \hat{P}_t^{(n)} - \ln \hat{P}_t^{(n+1)} = A_{n+1} - A_n + \\
&(B_{c,n+1} - B_{c,n})(\ln \frac{c_t}{c_{t-1}} - \theta_c) + (B_{v,n+1} - B_{v,n})(v_t - \theta_v) + \\
&(B_{s,n+1} - B_{s,n})s_t + (D_{v,n+1} - D_{v,n})Var_t[v_{t+1}] + \\
&(D_{s,n+1} - D_{s,n})Var_t[s_{t+1}] + (D_{c,n+1} - D_{c,n})Var_t[\ln \frac{c_{t+1}}{c_t}] \\
&= \delta + \bar{\pi} + (\pi_c + \gamma_c)\theta_c + (\pi_v + \gamma_v)\theta_v + (1 - \phi_v)\sigma_c^2 D_{c,n}\theta_v + \\
&\phi_c^n(\gamma_c + \phi_c\pi_c)(\ln \frac{c_t}{c_{t-1}} - \theta_c) + \\
&\phi_v^n(\gamma_v + \phi_v\pi_v)(v_t - \theta_v) + \phi_s^n\phi_s\pi_s(s_t - 0) - \\
&\left[(1 - \phi_v)D_{c,n} + \frac{1}{2}(\lambda_c + \pi_c + B_{c,n})^2\right]\sigma_c^2 v_t - \\
&\frac{1}{2}(\lambda_v + \pi_v + B_{v,n} + \sigma_c^2 D_{c,n})^2\sigma_v^2 - \frac{1}{2}(\pi_s + B_{s,n})^2\sigma_s^2 \\
&= \delta + \bar{\pi} + (\pi_c + \gamma_c)\theta_c + (\pi_v + \gamma_v)\theta_v - \\
&(1 - \phi_v)\sigma_c^2 D_{c,n}(v_t - \theta_v) + \phi_c^n(\gamma_c + \phi_c\pi_c)(\ln \frac{c_t}{c_{t-1}} - \theta_c) + \\
&\phi_v^n(\gamma_v + \phi_v\pi_v)(v_t - \theta_v) + \phi_s^n\phi_s\pi_s s_t - \\
&-\frac{1}{2}\left[(\lambda_c + \pi_c)^2 + 2B_{c,n}(\lambda_c + \pi_c) + B_{c,n}^2\right]\sigma_c^2 v_t - \\
&\frac{1}{2}\left[(\lambda_v + \pi_v)^2 + 2(B_{v,n} + \sigma_c^2 D_{c,n})(\lambda_v + \pi_v) + (B_{v,n} + \sigma_c^2 D_{c,n})^2\right]\sigma_v^2 - \\
&\frac{1}{2}(\pi_s + B_{s,n})^2\sigma_s^2 =
\end{aligned}
$$

$$
\begin{aligned}
= \;& \delta + \bar{\pi} + \pi_c\theta_c + \gamma_c\theta_c + \pi_v\theta_v + \gamma_v\theta_v - \frac{1}{2}(\lambda_c + \pi_c)^2\sigma_c^2 v_t - \\
& \frac{1}{2}(\lambda_v + \pi_v)^2\sigma_v^2 + \phi_c^n(\gamma_c + \phi_c\pi_c)(\ln\frac{c_t}{c_{t-1}} - \theta_c) + \\
& \phi_v^n(\gamma_v + \phi_v\pi_v)(v_t - \theta_v) + \phi_s^n\phi_s\pi_s s_t - (1-\phi_v)\sigma_c^2 D_{c,n}(v_t - \theta_v) - \\
& \left[(B_{v,n} + \sigma_c^2 D_{c,n})(\lambda_v + \pi_v) + \frac{1}{2}(B_{v,n} + \sigma_c^2 D_{c,n})^2\right]\sigma_v^2 - \\
& \frac{1}{2}(\pi_s + B_{s,n})^2\sigma_s^2 - \left[B_{c,n}(\pi_c + \lambda_c) + \frac{1}{2}B_{c,n}^2\right]\sigma_c^2 v_t\,. \qquad (64)
\end{aligned}
$$

The corresponding expression for the real forward rate can be derived following the same steps, starting from the definition for that rate given in the paper. However, a more expedient way to obtain the same result is to set all the parameters related to the inflation rate stochastic process equal to 0 in the above expression—that is, $\bar{\pi} = \pi_c = \pi_v = \pi_s = 0$— and replace the coefficients associated with the nominal price with the analogous ones associated with the real price. That shortcut results in the following expression for the real forward rate:

$$
\begin{aligned}
\ln F\hat{R}_t^{(n)\to(n+1)} &= \ln\hat{p}_t^{(n)} - \ln\hat{p}_t^{(n+1)} \\
&= \delta + \gamma_c\theta_c + \gamma_v\theta_v - \frac{1}{2}\lambda_c^2\sigma_c^2 v_t - \frac{1}{2}\lambda_v^2\sigma_v^2 + \\
& \quad \phi_c^n\gamma_c(\ln\frac{c_t}{c_{t-1}} - \theta_c) + \phi_v^n\gamma_v(v_t - \theta_v) - \\
& \quad (1-\phi_v)\sigma_c^2 d_{c,n}(v_t - \theta_v) - \left(b_{c,n}\lambda_c + \frac{1}{2}b_{c,n}^2\right)\sigma_c^2 v_t - \\
& \quad \left[(b_{v,n} + \sigma_c^2 d_{c,n})\lambda_v + \frac{1}{2}(b_{v,n} + \sigma_c^2 d_{c,n})^2\right]\sigma_v^2. \qquad (65)
\end{aligned}
$$

Subtracting (65) from (64) results in the following one-period-ahead nominal–real forward rate spread:

$$
\begin{aligned}
& \ln F\hat{N}_t^{(n)\to(n+1)} - \ln F\hat{R}_t^{(n)\to(n+1)} = \bar{\pi} + \pi_c\theta_c + \pi_v\theta_v - \lambda_c\pi_c\sigma_c^2 v_t - \\
& \frac{1}{2}\pi_c^2\sigma_c^2 v_t - \lambda_v\pi_v\sigma_v^2 - \frac{1}{2}\pi_v^2\sigma_v^2 + \phi_c^{n+1}\pi_c(\ln\frac{c_t}{c_{t-1}} - \theta_c) + \\
& \phi_v^{n+1}\pi_v(v_t - \theta_v) + \phi_s^{n+1}\pi_s s_t - (1-\phi_v)\sigma_c^2(D_{c,n} - d_{c,n})(v_t - \theta_v) - \\
& \left[B_{c,n}(\pi_c + \lambda_c) + \frac{1}{2}B_{c,n}^2 - (b_{c,n}\lambda_c + \frac{1}{2}b_{c,n}^2)\right]\sigma_c^2 v_t - \\
& \Bigg[(B_{v,n} + \sigma_c^2 D_{c,n})(\lambda_v + \pi_v) + \frac{1}{2}(B_{v,n} + \sigma_c^2 D_{c,n})^2 - \\
& \left[(b_{v,n} + \sigma_c^2 d_{c,n})\lambda_v + \frac{1}{2}(b_{v,n} + \sigma_c^2 d_{c,n})^2\right]\Bigg]\sigma_v^2 - \frac{1}{2}(\pi_s + B_{s,n})^2\sigma_s^2 =
\end{aligned}
$$

$$
\begin{aligned}
&= \bar{\pi} + \pi_c\theta_c + \pi_v\theta_v + \\
&\phi_c^{n+1}\pi_c(\ln \frac{c_t}{c_{t-1}} - \theta_c) + \phi_v^{n+1}\pi_v(v_t - \theta_v) + \phi_s^{n+1}\pi_s s_t - \\
&(\lambda_c + \frac{1}{2}\pi_c)\pi_c\sigma_c^2 v_t - \left[\pi_c B_{c,n} + (B_{c,n} - b_{c,n})\lambda_c + \frac{1}{2}(B_{c,n}^2 - b_{c,n}^2)\right]\sigma_c^2 v_t - \\
&(1-\phi_v)\sigma_c^2(D_{c,n} - d_{c,n})(v_t - \theta_v) - (\lambda_v + \frac{1}{2}\pi_v)\pi_v\sigma_v^2 - \frac{1}{2}(\pi_s + B_{s,n})^2\sigma_s^2 - \\
&\Big[\pi_v B_{v,n} + (B_{v,n} - b_{v,n})\lambda_v + \frac{1}{2}(B_{v,n}^2 - b_{v,n}^2) + \pi_v\sigma_c^2 D_{c,n} + \\
&(D_{c,n} - d_{c,n})\sigma_c^2\lambda_v + \frac{1}{2}(D_{c,n}^2 - d_{c,n}^2)\sigma_c^4 + (B_{v,n}D_{c,n} - b_{v,n}d_{c,n})\sigma_c^2\Big]\sigma_v^2. \qquad (66)
\end{aligned}
$$

As mentioned in the paper, the coefficients $B_{i,n}$, for $i = c$, v, s, evolve like truncated geometric sums because each new coefficient in the sequence $\{B_{i,n}\}$ is formed by adding the geometric increment $\phi_i^n(\gamma_i + \phi_i\pi_i)$ to the previous element in the sequence. That is,

$$B_{i,n} = (\gamma_i + \phi_i\pi_i)\frac{1-\phi_i^n}{1-\phi_i}. \qquad (67)$$

The same logic results in the following analogous expression for the corresponding coefficients in the formula for the real bond prices:

$$b_{i,n} = \gamma_i\frac{1-\phi_i^n}{1-\phi_i}. \qquad (68)$$

The last two equations imply the following useful relationships:

$$B_{i,n} - b_{i,n} = \phi_i\frac{1-\phi_i^n}{1-\phi_i}\pi_i; \qquad (69)$$

$$
\begin{aligned}
\frac{1}{2}(B_{i,n}^2 - b_{i,n}^2) &= \frac{1}{2}\left[\left(b_{i,n} + \phi_i\frac{1-\phi_i^n}{1-\phi_i}\pi_i\right)^2 - b_{i,n}^2\right] \\
&= b_{i,n}\phi_i\frac{1-\phi_i^n}{1-\phi_i}\pi_i + \frac{1}{2}(\phi_i\pi_i)^2\left(\frac{1-\phi_i^n}{1-\phi_i}\right)^2 \\
&= \gamma_i\phi_i\left(\frac{1-\phi_i^n}{1-\phi_i}\right)^2\pi_i + \frac{1}{2}(\phi_i\pi_i)^2\left(\frac{1-\phi_i^n}{1-\phi_i}\right)^2 \\
&= \phi_i\left(\frac{1-\phi_i^n}{1-\phi_i}\right)^2\pi_i(\gamma_i + \frac{1}{2}\phi_i\pi_i); \qquad (70)
\end{aligned}
$$

$$\begin{aligned} B_{v,n}D_{c,n} - b_{v,n}d_{c,n} &= \left(b_{v,n} + \phi_v \frac{1-\phi_v^n}{1-\phi_v}\pi_v\right) D_{c,n} - b_{v,n}d_{c,n} \\ &= b_{v,n}(D_{c,n} - d_{c,n}) + \phi_v \frac{1-\phi_v^n}{1-\phi_v}\pi_v D_{c,n} \\ &= \gamma_v \frac{1-\phi_v^n}{1-\phi_v}(D_{c,n} - d_{c,n}) + \phi_v \frac{1-\phi_v^n}{1-\phi_v}\pi_v D_{c,n} \\ &= \frac{1-\phi_v^n}{1-\phi_v}\Big[\gamma_v(D_{c,n} - d_{c,n}) + \phi_v \pi_v D_{c,n}\Big]. \quad (71) \end{aligned}$$

Substituting (67) to (71) into (66) results in the expression:

$$\begin{aligned} &\ln F\hat{N}_t^{(n)\to(n+1)} - \ln F\hat{R}_t^{(n)\to(n+1)} = \bar{\pi} + \pi_c\theta_c + \pi_v\theta_v + \\ &\quad \phi_c^{n+1}\pi_c(\ln \frac{c_t}{c_{t-1}} - \theta_c) + \phi_v^{n+1}\pi_v(v_t - \theta_v) + \phi_s^{n+1}\pi_s s_t + \\ &\quad (1-\phi_v)(D_{c,n} - d_{c,n})\ \theta_v\sigma_c^2 - (1-\phi_v)(D_{c,n} - d_{c,n})\ \sigma_c^2 v_t - \\ &\quad (\lambda_c + \frac{1}{2}\pi_c)\pi_c\sigma_c^2 v_t - \\ &\quad \Big[\gamma_c + \phi_c(\lambda_c + \pi_c) + \phi_c\frac{1-\phi_c^n}{1-\phi_c}(\gamma_c + \frac{1}{2}\phi_c\pi_c)\Big] \frac{1-\phi_c^n}{1-\phi_c}\pi_c\sigma_c^2 v_t - \\ &\quad (\lambda_v + \frac{1}{2}\pi_v)\pi_v\sigma_v^2 - \\ &\quad \Bigg\{ \Big[\gamma_v + \phi_v(\lambda_v + \pi_v) + \phi_v\frac{1-\phi_v^n}{1-\phi_v}(\gamma_v + \frac{1}{2}\phi_v\pi_v)\Big] \frac{1-\phi_v^n}{1-\phi_v}\pi_v + \\ &\quad D_{c,n}\pi_v\sigma_c^2 + (D_{c,n} - d_{c,n})\sigma_c^2\lambda_v + \frac{1}{2}(D_{c,n}^2 - d_{c,n}^2)\sigma_c^4 + \\ &\quad \frac{1-\phi_v^n}{1-\phi_v}\Big[\gamma_v(D_{c,n} - d_{c,n}) + \phi_v\pi_v D_{c,n}\Big]\sigma_c^2 \Bigg\}\sigma_v^2 - \\ &\quad \frac{1}{2}(\pi_s + B_{s,n})^2\sigma_s^2 = \end{aligned}$$

$$
\begin{aligned}
&= \bar{\pi} + \pi_c\theta_c + \pi_v\theta_v + \\
&\phi_c^{n+1}\pi_c(\ln\frac{c_t}{c_{t-1}} - \theta_c) + \phi_v^{n+1}\pi_v(v_t - \theta_v) + \phi_s^{n+1}\pi_s s_t + \\
&(1-\phi_v)(D_{c,n} - d_{c,n})\ \theta_v\sigma_c^2 - \Bigg\{(1-\phi_v)(D_{c,n} - d_{c,n}) + \Bigg[\lambda_c + \frac{1}{2}\pi_c + \\
&\Bigg[\gamma_c + \phi_c(\lambda_c + \pi_c) + \phi_c\frac{1-\phi_c^n}{1-\phi_c}(\gamma_c + \frac{1}{2}\phi_c\pi_c)\Bigg]\frac{1-\phi_c^n}{1-\phi_c}\Bigg]\pi_c\Bigg\}\sigma_c^2 v_t - \\
&\Bigg\{\Bigg[\lambda_v + \frac{1}{2}\pi_v + D_{c,n}\sigma_c^2 + \\
&\Bigg[\gamma_v + \phi_v(\lambda_v + \pi_v) + \phi_v\frac{1-\phi_v^n}{1-\phi_v}(\gamma_v + \frac{1}{2}\phi_v\pi_v)\Bigg]\frac{1-\phi_v^n}{1-\phi_v}\Bigg]\pi_v + \\
&\Bigg[\lambda_v(D_{c,n} - d_{c,n}) + \frac{1}{2}(D_{c,n}^2 - d_{c,n}^2)\sigma_c^2 + \\
&\frac{1-\phi_v^n}{1-\phi_v}\Bigg[\gamma_v(D_{c,n} - d_{c,n}) + \phi_v\pi_v D_{c,n}\Bigg]\Bigg]\sigma_c^2\Bigg\}\sigma_v^2 - \frac{1}{2}\Bigg[1 + \phi_s\frac{1-\phi_s^n}{1-\phi_s}\Bigg]^2\pi_s^2\sigma_s^2.
\end{aligned}
$$

Taking into account that $\bar{\pi} + \pi_c\theta_c + \pi_v\theta_v = E[\ln(1+\Pi_t)]$, this expression is equivalent to the compact representation of the nominal–real forward rate spread (58), with the coefficients $\Omega_{i,n}$ defined as follows:

$$
\begin{aligned}
\Omega_{v,n} \quad = \quad & -(1-\phi_v)(D_{c,n} - d_{c,n})\ \theta_v\sigma_c^2 + \Bigg\{\Bigg[\lambda_v + \frac{1}{2}\pi_v + D_{c,n}\sigma_c^2 + \\
&\Bigg[\gamma_v + \phi_v(\lambda_v + \pi_v) + \phi_v\frac{1-\phi_v^n}{1-\phi_v}(\gamma_v + \frac{1}{2}\phi_v\pi_v)\Bigg]\frac{1-\phi_v^n}{1-\phi_v}\Bigg]\pi_v + \\
&\Bigg[\lambda_v(D_{c,n} - d_{c,n}) + \frac{1}{2}(D_{c,n}^2 - d_{c,n}^2)\sigma_c^2 + \\
&\frac{1-\phi_v^n}{1-\phi_v}\Bigg[\gamma_v(D_{c,n} - d_{c,n}) + \phi_v\pi_v D_{c,n}\Bigg]\Bigg]\sigma_c^2\Bigg\}\sigma_v^2;
\end{aligned}
$$

$$
\Omega_{s,n} \quad = \quad \frac{1}{2}\left(1 + \phi_s\frac{1-\phi_s^n}{1-\phi_s}\right)^2\pi_s^2\sigma_s^2;
$$

$$
\begin{aligned}
\Omega_{c,n} \quad = \quad & \Bigg\{(1-\phi_v)(D_{c,n} - d_{c,n}) + \Bigg[\lambda_c + \frac{1}{2}\pi_c + \\
&\Bigg[\gamma_c + \phi_c(\lambda_c + \pi_c) + \phi_c\frac{1-\phi_c^n}{1-\phi_c}(\gamma_c + \frac{1}{2}\phi_c\pi_c)\Bigg]\frac{1-\phi_c^n}{1-\phi_c}\Bigg]\pi_c\Bigg\}\sigma_c^2.
\end{aligned}
$$

C. Parameter Values for the Simulations of the Three-Factor Model

This appendix presents and discusses the parameter values of the three-factor model underlying the simulation presented in Figure 4 of the paper.

As mentioned in the paper, most of the parameters have been set equal to those corresponding to experiment C in Tables 1 and 2 of Gallmeyer et al. (2007). The exceptions, justified later, are denoted with an inequality symbol ($\neq$) within parentheses next to the value in the list below. When the parameter values in that list are implied by theoretical relationships between other parameters, the relevant formulas are explicitly displayed. The parameters entering into those formulas that were left implicit in the paper are identified with an asterisk within parentheses, (*), and their correspondence with those in Gallmeyer et al. made apparent by keeping the same notation used by those authors.

Parameter values:

$$
\begin{aligned}
\theta_c &= 0.006; \\
\theta_v &= 0.0001825; \\
\phi_c &= 0.36; \\
\phi_v &= 0.973; \\
\phi_s &= 0.85(\neq); \\
\sigma_c &= \left[Var[\ln \frac{c_{t+1}}{c_t}](1-\phi_c^2)\frac{1}{\theta_v}\right]^{\frac{1}{2}} = 0.0048\left[(1-\phi_c^2)\frac{1}{\theta_v}\right]^{\frac{1}{2}} \\
&= 0.331489\,(\neq); \\
\sigma_v &= 0.98841\,.10^{-5}; \\
\sigma_s &= (0.023\,.10^{-4})^{\frac{1}{2}}; \\
\rho &= 0.5\,(*); \\
\alpha &= -4.911\,(*); \\
\beta &= 0.994\,(*); \\
\kappa &= 0.994\,(*); \\
\bar{\tau} &= -0.015\,(*); \\
\tau_x &= 3.064\,(*); \\
\tau_p &= 2.006\,(*); \\
\gamma_c &= (1-\rho)\phi_c = 0.18; \\
\gamma_v &= \frac{\alpha}{2}(\alpha-\rho)\left(\frac{1}{1-\kappa\phi_c}\right)^2\sigma_c^2 = 3.540538\ (\neq); \\
\lambda_c &= (1-\alpha)-(\alpha-\rho)\frac{\kappa\phi_c}{1-\kappa\phi_c} = 8.9262489; \\
\lambda_v &= -\frac{\kappa}{1-\kappa\phi_v}\gamma_v = -107.171410\ (\neq);
\end{aligned}
$$

$$
\begin{aligned}
\pi_c &= \frac{\gamma_c - \tau_x}{\tau_p - \phi_c} = -1.752126; \\
\pi_v &= \frac{\gamma_v - \frac{1}{2}(\lambda_c + \pi_c)^2 \sigma_c^2}{\tau_p - \phi_v} = 0.689987\ (\neq); \\
\pi_s &= -\frac{1}{\tau_p - \phi_s} = -0.865052\ (\neq); \\
\bar{\pi} &= \frac{\delta - \bar{\tau} + \pi_c(1-\phi_c)\theta_c + \pi_v(1-\phi_v)\theta_v - \frac{1}{2}(\lambda_v + \pi_v)^2\sigma_v^2 - \frac{1}{2}\pi_s^2\sigma_s^2}{\tau_p - 1} \\
&= 0.0161157\ (\neq); \\
\delta &= -\ln(\beta) + (1-\rho)(1-\phi_c)\theta_c + \\
&\quad \frac{\alpha}{2}(\alpha - \rho)\left[\frac{\kappa}{1-\kappa\phi_v}\frac{\alpha}{2}(\frac{1}{1-\kappa\phi_c})^2\sigma_c^2\right]^2 \sigma_v^2 = 0.007939\ (\neq).
\end{aligned}
$$

This list makes obvious that the parameters ϕ_s and σ_c^2, as well as those that depend on them, are set to different values than those selected by Gallmeyer et al.

The purpose of the different choice of value for ϕ_s (0.85 instead of 0.922) is to isolate the fluctuations in the inflation-expectations indicator induced by time-varying risk premia from those induced by the short-term dynamics of the monetary policy factor, s_t, captured by the last term in the second line of (58). Numerical experimentations with the model reveal that deviations of that factor from its long-run mean (0 by assumption) are responsible for a sizable fraction of the fluctuations in the simulated inflation-expectations indicator when the parameter ϕ_s takes on values of 0.90 or higher. Notice that changing the values of this parameter has no impact on the fluctuations of the inflation-expectations indicator that can be traced to time-varying risk premia in the model. That is because, as demonstrated in Appendix B, the autoregressive parameter ϕ_s doesn't enter into the coefficient $\Omega_{c,n}$ in (58), the one that ultimately controls the fluctuations of the forward nominal–real rate spread originated in risk premia that change over time.

As noted in the text, the variance of the innovation associated with real consumption growth, σ_c^2, was set to the value implied by the *unconditional* variance of that variable under time-varying risk premia. It can be shown that this unconditional variance is given by the expression:

$$
Var[\ln\frac{c_{t+1}}{c_t}] = E\left[\ln\frac{c_{t+1}}{c_t} - E[\ln\frac{c_{t+1}}{c_t}]\right]^2 = \frac{\sigma_c^2}{1-\phi_c^2}\theta_v.
$$

This method seems to be more in line with the empirical evidence than the seemingly arbitrary assumption that $\sigma_c^2 = 1$, under which Gallmeyer et al. derived their theoretical formulas. As mentioned in footnote 24 of the paper, however, those authors make reference later in their study (bottom of pg. 313) to a formula for this parameter that seems to be related to the unconditional variance of real consumption growth for the constant risk premia case.

As should be obvious from the list of parameter values, the choice of a different value for σ_c^2 from the one selected by Gallmeyer et al. in their theoretical section has a cascade effect because that parameter enters into the formulas that determine the value of many other parameters of the model. Consequently, the quantitative differences with the results obtained by those authors may not be trivial. For example, if $\sigma_c^2 = 1$, as in Gallmeyer et al., the value of the parameter γ_v—the factor loading on the real pricing kernel corresponding to the factor v_t —would be 32.22, the same as those authors report for experiment C in their Table 1. This value is about ten times larger than the one implied by the underlying relationships between the relevant parameters in this paper.

Incidentally, the difference in the value of the parameter π_s in the paper, -0.86, from that in Gallmeyer et al. is not entirely due to the choice of a different value for the parameter ϕ_s because, even for $\phi_s = 0.922$ (the value chosen by those authors), application of the relevant formula in the list of parameter values would imply $\pi_s = -0.92251$ instead of the value -0.61 that Gallmeyer et al. report in experiment C in their Table 2.

REFERENCES

Backus, David K., and Jonathan H. Wright (2007), "Cracking the Conundrum," *Brookings Papers on Economic Activity* 38 (1): 293–329.

Backus, David K., and Stanley E. Zin (1994), "Reverse Engineering the Yield Curve," NBER Working Paper Series, no. 4676 (Cambridge, Mass., National Bureau of Economic Research, March).

Bansal, Ravi, and Amir Yaron (2004), "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance* 59 (4): 1481–509.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay (1997), *The Econometrics of Financial Markets* (Princeton, N.J.: Princeton University Press).

Campbell, John Y., and Robert J. Shiller (1991), "Yield Spreads and Interest Rate Movements: A Bird's Eye View," NBER Working Paper Series, no. 3153 (Cambridge, Mass., National Bureau of Economic Research, September).

Clarida, Richard, Jordi Galí, and Mark Gertler (2000), "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* 115 (1): 147–80.

Cochrane, John H. (2005), *Asset Pricing, Revised Edition* (Princeton, N.J.: Princeton University Press).

——(2007), "Commentary on 'Macroeconomic Implications of Changes in the Term Premium,'" Federal Reserve Bank of St. Louis *Review* 89, July/August, 271–82.

Cochrane, John H., and Monika Piazzesi (2005), "Bond Risk Premia," *The American Economic Review* 95 (1): 138–60.

Cook, Timothy, and Thomas Hahn (1990), "Interest Rate Expectations and the Slope of the Money Market Yield Curve," Federal Reserve Bank of Richmond *Economic Review*, September/October, 3–26.

Cox, John C., Jonathan E. Ingersoll Jr., and Stephen A. Ross (1985), "A Theory of the Term Structure of Interest Rates," *Econometrica* 53 (2): 385–407.

Duffie, Darrell, and Rui Kan (1996), "A Yield-Factor Model of Interest Rates," *Mathematical Finance* 6 (4): 379–406.

Fama, Eugene F. (1990), "Term-Structure Forecasts of Interest Rates, Inflation, and Real Returns," *Journal of Monetary Economics* 25 (1): 59–76.

Fisher, Irving (1896), "Appreciation and Interest," *AEA Publications* 3 (11): 331–442. (Reprinted, New York: Augustus M. Kelley, 1961.)

——(1930), *The Theory of Interest* (New York: Macmillan).

Fisher, Mark (2001), "Forces That Shape the Yield Curve," Federal Reserve Bank of Atlanta *Economic Review* 86, First Quarter, 1–15.

Gallmeyer, Michael F., Burton Hollifield, Francisco J. Palomino, and Stanley E. Zin (2007), "Arbitrage-Free Bond Pricing with Dynamic Macroeconomic Models," Federal Reserve Bank of St. Louis *Review* 89, July/August, 305–26.

Greenspan, Alan (2005a), Federal Reserve Board's semiannual Monetary Policy Report to the Congress (Testimony Before the Committee on Financial Services, U.S. House of Representatives, Washington, D.C., July 20).

——(2005b), Federal Reserve Board's semiannual Monetary Policy Report to the Congress (Testimony Before the Committee on Banking, Housing, and Urban Affairs, U.S. Senate, Washington, D.C., February 16).

Gürkaynak, Refet S., Brian Sack, and Eric Swanson (2005), "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models," *The American Economic Review* 95 (1): 425–36.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright (2006), "The U.S. Treasury Yield Curve: 1961 to the Present," Finance and Economics Discussion Series, no. 2006–28 (Washington, D.C., Federal Reserve Board).

——(2008), "The TIPS Yield Curve and Inflation Compensation," Finance and Economics Discussion Series, no. 2008–05 (Washington, D.C., Federal Reserve Board).

Kohn, Donald L. (2005), "Monetary Policy Perspectives on Risk Premiums in Financial Markets" (Speech at the Financial Market Risk Premiums Conference, Federal Reserve Board, Washington, D.C., July 21), www.federalreserve.gov/boarddocs/speeches/2005/20050721/default.htm.

Kydland, Finn E., and Edward C. Prescott (1977), "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85 (3): 473–91.

Macaulay, Frederick R. (1938), *Some Theoretical Problems Suggested by the Movements of Interest Rates, Bond Yields, and Stock Prices in the United States Since 1856* (Cambridge, Mass.: National Bureau of Economic Research).

McCulloch, J. Huston, and Heon-Chul Kwon (1993), "U.S. Term Structure Data, 1947–1991," Ohio State University Working Paper, no. 93-6 (March), www.econ.ohio-state.edu/jhm/ts/mcckwon/mccull.htm.

Mehra, Rajnish, and Edward C. Prescott (1985), "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15 (2): 145–61.

Merton, Robert C. (1973), "An Intertemporal Capital Asset Pricing Model," *Econometrica* 41 (5): 867–87.

Mishkin, Frederic S. (1990), "What Does the Term Structure Tell Us About Future Inflation?" *Journal of Monetary Economics* 25 (1): 77–95.

Phillips, Peter C. B. (1998), "Econometric Analysis of Fisher's Equation," Cowles Foundation Discussion Papers no. 1180 (New Haven, Conn., Yale University, Cowles Foundation for Research in Economics, May).

Piazzesi, Monika, and Martin Schneider (2006), "Equilibrium Yield Curves," NBER Working Paper Series, no. 12609 (Cambridge, Mass., National Bureau of Economic Research, October).

Ross, Stephen A. (1976), "The Arbitrage Theory of Capital Asset Pricing," *Journal of Economic Theory* 13 (3): 341–60.

Rudebusch, Glenn D. (1993), "Propagation of Federal Funds Rate Changes to Longer-Term Rates Under Alternative Policy Regimes," in *Operating Procedures and the Conduct of Monetary Policy: Conference Proceedings*, ed. Marvin Goodfriend and David H. Small (Washington, D.C.: Federal Reserve Board).

——(1998), "Do Measures of Monetary Policy in a VAR Make Sense?" *International Economic Review* 39 (4): 907–31.

Rudebusch, Glenn D., Brian P. Sack, and Eric T. Swanson (2007), "Macroeconomic Implications of Changes in the Term Premium," Federal Reserve Bank of St. Louis *Review* 89, July/August, 241–69.

Rudebusch, Glenn D., and Eric T. Swanson (2009), "The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks," Federal Reserve Bank of San Francisco Working Paper Series, no. 2008–31 (March).

Stambaugh, Robert F. (1988), "The Information in Forward Rates: Implications for Models of the Term Structure," *Journal of Financial Economics* 21 (1): 41–70.