Volatility and the Treasury yield curve[1]

Christian Gilles

Introduction

The topic for this year's autumn meeting is the measurement, causes and consequences of financial market volatility. For this paper, I limit the scope of analysis to the market for US Treasury securities, and I examine how the volatility of interest rates affects the shape of the yield curve. I consider explicitly two types of measurement issues: since yields of different maturities have different volatilities, which maturity to focus on; and how to detect a change in volatility. Although understanding what *causes* the volatility of financial markets to flare up or subside is perhaps the most important issue, I will have nothing to say about it; like much of contemporaneous finance theory, I treat interest rate volatility as exogenous. To provide context for the analysis, I discuss the reasons that led to work currently going on at the Federal Reserve Board, which is to estimate a particular three-factor model of the yield curve. That work is still preliminary, and I have no results to report. Current efforts are devoted to resolving tricky econometric and computational issues which are beyond the scope of this paper. What I want to do here is to explain the theoretical and empirical reasons for estimating a model in this particular class.

This project's objective is to interpret the nominal yield curve to find out what market participants think will happen to future short-term nominal interest rates. It would be even better to obtain a market-based measure of expected inflation, but this goal would require data not merely on the value of nominal debt but also on the value of indexed debt, which the Treasury does not yet issue. In any event, the expectation of future nominal rates is a necessary first step on the road toward a measure of expected inflation, and it also has some independent value for a central bank because it shows how markets interpret the current stance of monetary policy.

Until recently, economists frequently assumed that rational expectations demanded that current forward rates be unbiased predictors of future spot rates - following Cox, Ingersoll and Ross (1981), we call this version of the expectations hypothesis the (strong or pure) *yield-to-maturity expectations hypothesis* (YTM-EH). But the overwhelming evidence is that forward rates are biased predictors of future spot rates, leading macroeconomists to theorize about the presence of term premiums, often informally justified by appeal to behavioral assumptions such as market segmentation, preferred habitat and so on. In an attempt to deal with the empirical failure of the strong YTM-EH, econometricians later tested a weaker version, which postulated that changes in forward rates signal changes of equal magnitude in the expectations of future rates, or, equivalently, that the term premium at each maturity is constant through time. As reported by Campbell and Shiller (1991), for example, empirical results strongly reject even this weak YTM-EH.

Underlying the Board's work on the expected path of interest rates, there is a continuous-time, arbitrage-free, three-factor model of the yield curve. Arbitrage-free models of asset prices in finance, in which rational expectations are always hard-wired, shed much light on the behavior of term premiums. In explaining these premiums, and therefore in understanding how to get from a forward rate to an expected future spot rate, the volatility of interest rates plays a star role. More precisely, volatility plays two roles: in its first role it acts alone to produce a *convexity premium,* and in its second role it interacts with investors' preferences to produce a *risk premium.*

1 The views expressed in this paper should not be construed as reflecting those of the Federal Reserve Board or other members of its staff. I gratefully acknowledge countless conversations with Mark Fisher, which have shaped my understanding of the relation between volatility and yields.

The next two sections explain the concept of a convexity premium and the following section focuses on the risk premium. Section 4 contains a discussion of the properties of well-known diffusion models of the yield curve, and Section 5 introduces the three-factor model that the Federal Reserve Board is currently trying to estimate. The paper concludes with a brief overview of some of the econometric issues.

1. The convexity premium in a static setting[2]

Term premiums always embed a factor which pulls forward rates below the expected future spot rates at corresponding horizons. This downward pull stems from the convex relationship between the price and the yield of zero-coupon securities. Perhaps the best way to understand how this effect works is to examine it in a simple setting, in which investors are risk-neutral and there are no dynamic complexities.

To fix the terminology, let $P(t,\tau)$ denote the time-t price of a zero-coupon, default-free bond that matures at time τ. Then the yield to maturity on this bond is

$$y(t,\tau) := -\frac{\log[P(t,\tau)]}{\tau - t};$$

the instantaneous forward rate at time t for horizon τ is

$$f(t,\tau) := -\frac{d\log[P(t,\tau)]}{d\tau};$$

and the spot rate at time t is

$$r(t) := \lim_{\tau\to t} f(t,\tau).$$

Unless otherwise specified, the term "yield curve" in this paper refers to the graph of the yield to maturity on zero-coupon bonds, $y(t,\tau)$ (also called a zero rate), as a function of the time to maturity τ-t. The graph of $f(t,\tau)$ as a function of τ-t will be called the "forward rate curve" or some similar expression.

If the future path of the rate of interest were known with certainty, then the current forward rate would have to equal the future spot rate to avoid an obvious arbitrage opportunity. That is, if $r(t)$ is known at time zero for all $t > 0$, then

$$P(t,\tau) = \exp\left(-\int_t^\tau r(s)ds\right), \tag{1}$$

so that $y(t,\tau) = (\int_t^\tau r(s)ds)/(\tau - t)$ and $f(t,\tau) = r(\tau)$ for all $t<\tau$. In particular, if the spot rate is constant, say $r(t) = r$ for all t, then $y(0,\tau) = f(0,\tau) = r$ for all τ, so that the yield curve is flat.[3]

Now, imagine an economy in which a coin flip at date 0 determines which of two deterministic paths interest rates will follow. To be specific, suppose that if the coin comes up heads the spot rate stays constant at $r_1 = 5\%$ per year, and if it comes up tails the spot rate stays constant at $r_2 = 15\%$ per year. Note that, before the coin is flipped, $r(t)$ is a random variable with a constant expected value of 10%, independent of t. Whatever the outcome, then, the post-flip yield curve is flat; but what is the shape of the pre-flip yield curve?

---

2 This section follows the exposition in Fisher and Gilles (1995a).

3 In general, the yield curve must be distinguished from the forward rate curve, but here the two curves are identical.

The pure YTM-EH offers to this question the answer that the pre-flip yield curve is also flat at the expected rate of 10%. This answer may appear to be the natural outcome of the assumption of risk neutrality, and as such may be intuitively appealing. Despite the appearance, though, risk neutrality does not imply that the yield curve will be flat; in fact, no rational behavior would imply a yield curve flat at the expected rate because this would give rise to the following arbitrage opportunity. Post-flip bond prices have a yield of either 5% or 15%. A portfolio consisting of 8.906 units of the thirty-year bond and 4.9251 units (a short position) of the ten-year bond is worth \$1 at either rate; that is, unwinding the position immediately after the flip will cost \$1 whether the coin comes up heads or tails. If the pre-flip yields on these bonds were both equal to 10%, establishing the position before the flip would generate \$1.37. In other words, a pre-flip yield curve flat at 10% allows an investor to sell for \$1.37 a promise to pay \$1 unconditionally after the flip. This arbitrage opportunity proves the claim that the yield curve cannot be flat at 10%.

It is then natural to wonder what shape the yield curve would have if investors were indeed risk-neutral. By risk neutrality, we mean that the value of a random pay-off to be received in the next instant equals its expected value. Applying this asset pricing formula to a discount bond of maturity τ, the pre-flip price, $P_0(0,\tau)$, of this bond equals its average post-flip price, $[P_1(0,\tau) + P_2(0,\tau)]/2$, where P_i is the price when $r(t) = r_i$, for $i = 1, 2$. Figure 1 below illustrates the result. The curved line in the left panel is the graph of the convex function $y = \log[P]/10$, which is the relationship between the yield to maturity y on a ten-year bond and its price P. The two post-flip prices are $P_1 = \$0.61$ and $P_2 = \$0.22$, corresponding to the two yields $r_1 = 0.05$ and $r_2 = 0.15$. The pre-flip price is therefore the average $P_0 = 0.41$, which corresponds to a yield of 8.80%, 120 basis points below the average yield of 10%.

Figure 1

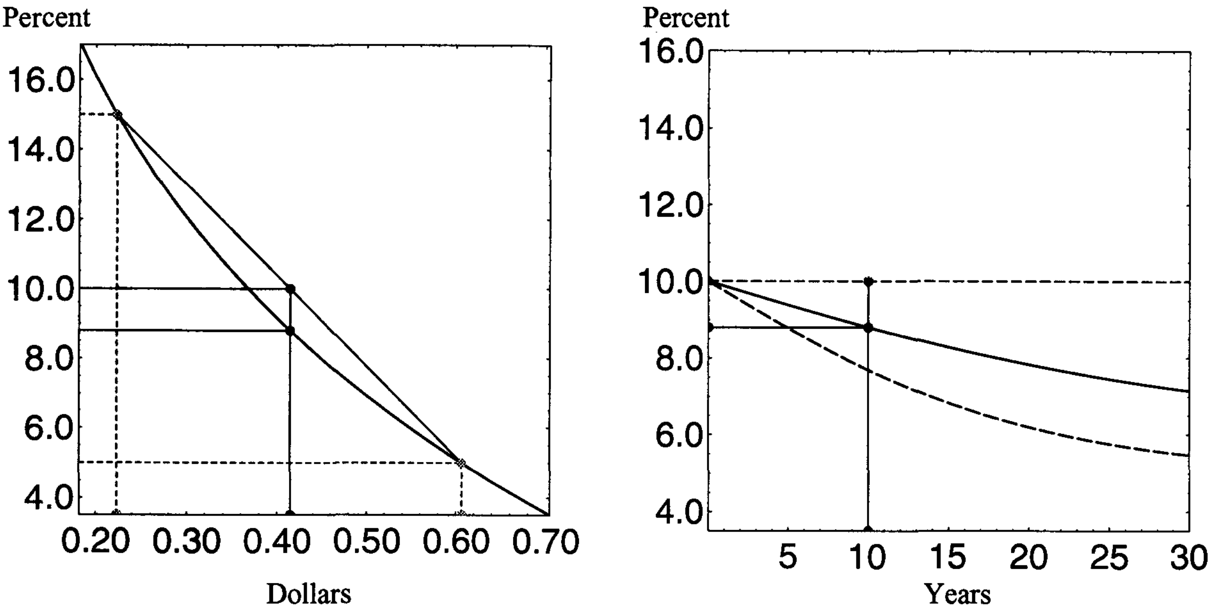

By repeating the same procedure for different values of the maturity τ, we trace out the yield curve shown as the solid line in the right panel of Figure 1. The downward-sloping dashed line shows the corresponding forward rates. Because the functional relation between price and yield is more convex the higher the maturity of the bond, the spread between the expected future spot rate and the forward rate increases with maturity. The result is that, when the spot rate is not expected to change, the yield curve is downward-sloping and the current forward rate underpredicts the future spot rate, with the bias an increasing function of maturity.

Observe the crucial role that volatility plays in this scenario. A decrease in rate volatility that keeps future expected rates unchanged can be modeled by a less extreme distribution of rates around its mean of 10%, say 7% and 13% instead of 5% and 15%. With such a decrease, the straight line joining the two outcomes in the left panel of Figure 1 would move toward the origin, thus decreasing the spread between expected rate and yield. In other words, the yield curve in the right panel would be flatter, and the bias in forward rates would be reduced.

In summary, the convexity of the relationship between price and yield leads forward rates to underpredict future spot rates. The bias increases with the horizon at which we predict and with the volatility of the future spot rate.

2. The convexity premium in a dynamic setting[4]

Although it is ideally suited to developing and testing one's intuition, the static model we have just used is not rich enough to develop an understanding of the dynamic evolution of yield curves and the role that volatility plays in this evolution. We thus introduce a class of richer models that have more potential to capture the essential empirical regularities, but we will see that the simple mechanism exposed in the simple static setting survives basically intact.

In the previous setting, the spot rate of interest changed randomly only once (and by a large amount); all subsequent changes, if any, were supposed to be deterministic. In the new setting, the spot rate changes randomly all the time, each time by a small amount. Formally, we suppose that the change in the spot rate at time t over the next interval of time dt is

$$dr(t) = \mu_r(t)dt + \sigma_r(t) \cdot dW(t). \tag{2}$$

Here, $\mu_r(t)$ is the expected change (called the "drift") of the spot rate; $\sigma_r(t)$ is the volatility (also called the "diffusion") of that change; and $W(t)$ *is* a Wiener process, that is, a continuous random process such that $W(t+s)$-$W(t)$, the increment over the interval of time s, is normally distributed with zero mean and variance s, and is independent of other increments (over non-overlapping intervals).[5] A process with a zero drift, like $W(t)$ itself, is called a "martingale". The drift $\mu_r(t)$ and the diffusion $\sigma_r(t)$ may be constant or deterministic functions of time, but they may also be random processes.

The process for the spot rate is a crucial factor explaining the shape of the yield curve at any date. In fact, maintaining the assumption of risk neutrality for the time being, the short rate process is the *only* factor that matters for the yield curve. But this statement makes sense only if the use of the term "risk neutrality" is clarified in this dynamic setting. Here, it refers to what Cox, Ingersoll and Ross (1981) call the *Local Expectations Hypothesis* (LEH), which postulates that the instantaneous expected rate of return on any asset equals the spot rate $r(t)$. In other words, writing the process for a bond maturing at time τ as

$$\frac{dP(t,\tau)}{P(t,\tau)} = \mu p(t,\tau)\, dt + \sigma p(t,\tau) \cdot dW(t), \tag{3}$$

the LEH requires $\mu p(t,\tau) = r(t)$. It turns out that in this case the bond price is a natural generalization of the formula (1), which applies in the deterministic case:

$$P(t,\tau) = E_t\left[\exp\left(-\int_t^\tau r(s)ds\right)\right], \tag{4}$$

where E_t denotes expectations taken with time-t information. The problem now is to find out what the pricing equation (4) implies for the relationship between the forward rate curve and the path of expected future spot rates.

4 This section and the next one borrow heavily from Fisher and Gilles (1995b).

5 With more than one source of risk, $W(t)$ is a vector of independent Wiener processes, σ_r is a commensurate vector, and $\sigma_r(t) \cdot dW(t)$ stands for the inner product $\Sigma_i \sigma_{ri} dW_i(t)$.

It is clear from the form of the equation that the process for the short rate is the only factor entering bond prices. It is also clear that the so-called Jensen's inequality for convex function comes into play, driving a convexity wedge between forward rates and expected future spot rate. Recall that, by definition, $y(t,\tau) = -\log[P(t,\tau)]/(\tau-t)$ and $f(t,\tau) = -d\log[P(t,\tau)]/dt$. Thus, passing the $-\log[\cdot]$ operator through the expectations operator in (4) produces $y(t,\tau) = E_t\left[\int_t^\tau r(s)ds\right]/(\tau-t)$ and $f(t,\tau) = E_t[r_\tau]$ as postulated by YTM-EH. But this procedure is illegitimate, of course: $-\log[x]$ is a convex function of x, so that Jensen's inequality requires both $f(t,\tau) < E_t[r_\tau]$ and $y(t,\tau) < E_t\left[\int_t^\tau r(s)ds\right]/(\tau-t)$.

Results in Heath, Jarrow and Morton (1992) allow a deeper analysis of convexity premiums. To fix notation, write the yield to maturity on the bond maturing at τ as the sum of the average expected spot rate from t to τ and a *term premium* $\Lambda_y(t,\tau)$: of

$$y(t,\tau) = \frac{1}{\tau - t} E_t\left[\int_t^\tau r(s)\,ds\right] + \Lambda_y(t,\tau). \tag{5}$$

Then, the term premium is a pure convexity premium and can be written as[6]

$$\Lambda_y(t,\tau) = \frac{1}{\tau - t} E_t\left[\int_{v=t}^{\tau} -\frac{1}{2}\left|\sigma p(v,\tau)\right|^2 dv\right], \tag{6}$$

Similarly, write the forward rate as the sum of the expected future spot rate and a term premium $\Lambda_f(t,\tau)$.

$$f(t,\tau) = E_t([r_\tau] + \Lambda_f(t,\tau).$$

The term premium is a pure convexity premium, given by

$$\Lambda_f(t,\tau) = \frac{\partial}{\partial\tau}((\tau - t)\Lambda_y(t,\tau)) = E_t\left[\int_{v=t}^{\tau} \sigma p(v,\tau)\cdot\frac{\partial}{\partial\tau}\sigma p(v,\tau)dv\right]. \tag{7}$$

The expression (6) underscores the importance of volatilities (or diffusions) for the convexity premium in yields, and it focuses on the diffusions that are important: those of the bond maturing at τ at all future dates until maturity. Of course, in the light of expression (4), these diffusions all depend on the process for the short rate. But this is not the same as saying that only the diffusion of the spot rate, $\sigma_r(t)$, matters, because both this diffusion and the drift $\mu_r(t)$ may themselves be random and the drift and diffusions of the processes for $\sigma_r(t)$ and $\mu_r(t)$ contribute to σp in complicated ways just as much as $\sigma_r(t)$ itself.

It is also clear from (6) that the yield on a bond maturing at date τ always underpredicts the average level of the spot rate over the life of the bond. It might seem from the form of that expression that the absolute value of $(\tau-t)\Lambda(t-\tau)$ is increasing in τ, but this is strictly necessary only when $|\sigma p(t,\tau)|$ increases in τ. The term $|\sigma p(t,\tau)|$ may decrease in τ over some range in such a way that $\Lambda_f(t,\tau)$ is positive for some values of τ and the forward rate overpredicts the expected future spot rate in that range. But such examples are artificial; virtually all natural assumptions about the spot rate process (2) imply that $|\sigma p(t,\tau)|$ increases monotonically in τ, so that $\Lambda_f(t,\tau)$ is negative and the forward rate underpredicts the future spot rate at all maturities.[7]

---

6 With a unique source of risk, σp is a scalar and $|\sigma p|$ is its absolute value; with several sources of risk, σp is a vector and $|\sigma p|$ is its norm.

7 Note that if τ is sufficiently close to t (that is, at sufficiently short maturities), $|\sigma p(t,\tau)|$ must be increasing in τ, as $\sigma p(t,t) = 0$.

3. The risk premium

There is a voluminous empirical literature on the performance of forward rates as predictors of future spot rates, typically using Treasury bill rates and thus focusing attention on the short end of the maturity spectrum (see MacDonald and Hein (1989), for one random example). The evidence of a bias in forward rates is conclusive, but the bias is positive, contrary to what would result from a convexity premium alone, *i.e.* under the L-EH. Reconciling theory with this evidence requires dropping the L-EH, thus appealing to risk premiums.

The shape of the average yield curve provides an immediate clue as to the direction and size of the forward rate bias. Assuming the spot rate is stationary, the average spot rate (in a sufficiently long sample) would equal the mean of its stationary distribution θ; then, if forward rates were unbiased predictors of future spot rates, the average forward rate would equal θ, so that the average yield curve would be flat at θ. Figure 2 shows the short end of the average end-of-month yield curve. Although the sample period, December 1987 to September 1995, may be too short for the average short rate to provide a good estimate of the mean of its stationary distribution, the main feature of the curve - namely, its positive slope - would show up in any longer sample. How does theory explain an upward-sloping average yield curve?

Figure 2

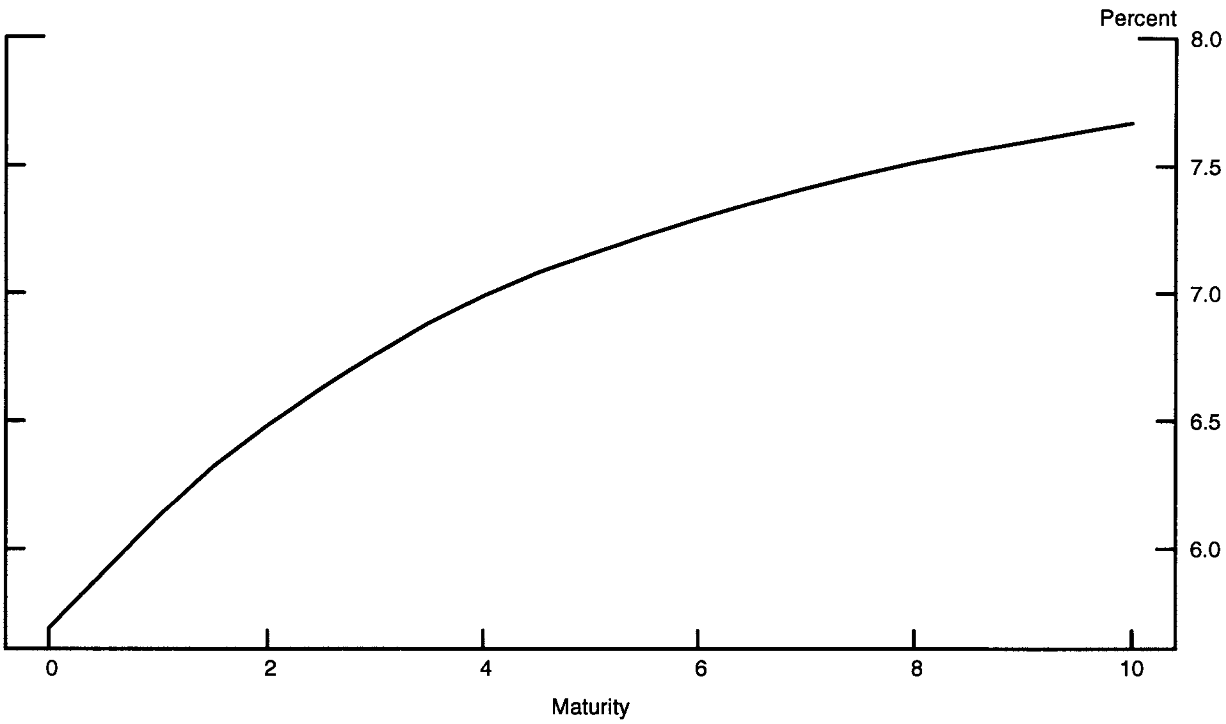

Given the interest rate process in (2), absence of arbitrage is equivalent to the existence of $\lambda(t)$, which may be a stochastic process but is the same for all securities, such that the rate of return on the bond maturing at τ - whose process is described in (3) - is

$$\mu p(t,\tau) = r(t) + \lambda(t) \cdot \sigma p(t,\tau). \tag{8}$$

The variable $\lambda(t)$ is called the "market price of risk". It has the same dimension as $W(t)$ - and as σp - so there really is one price per source of risk. The no-arbitrage condition thus says that expected return may exceed the riskless rate $r(t)$ as compensation for the risk in the security, where the diffusion σp measures the exposure to risk and $\lambda(t)$ is the compensation per unit of exposure. The L-EH thus amounts to assuming a zero market price of risk.

A non-zero market price of risk leads to the presence of a risk premium in yields in addition to the convexity premium. The general representation of the term premium $\Lambda_y(t,\tau)$ in (5), valid even when $\lambda(t) \neq 0$, is

$$\Lambda_y(t,\tau) = \frac{1}{\tau - t} E_t \left[\int_{v=t}^{\tau} \left(\lambda(t) \cdot \sigma p(v,\tau) - \frac{1}{2} |\sigma p(v,\tau)|^2 dv \right) \right]. \tag{9}$$

The first term in the integrand gives rise to the *risk premium,* while the second is the now-familiar convexity premium. The representation for the forward term premium that we introduced in (7), $\Lambda_f(t,\tau) = \frac{\partial}{\partial \tau}((\tau - t)\Lambda_y(t,\tau))$, which remains valid in the general case, transmits to that term premium the same decomposition into a risk premium and a convexity premium.

The average positive slope of the yield curve, then, is evidence in favor of the empirical importance of risk premiums. Under the L-EH, risk premiums are zero, so that convexity premiums would impart a negative slope to the yield curve. The YTM-EH, which asserts that forward rates are unbiased predictors of the future spot rates, amounts to postulating that risk premiums and convexity premiums cancel each other out at all possible maturities. Although this outcome is not a logical impossibility, it requires some peculiar circumstances: as expression (9) makes clear, the risk premium is linear in σp, while the convexity premium is quadratic in $|\sigma p|$. Equality for all maturities is impossible (except in the trivial deterministic case) when there is only one source of risk, so that $\sigma p(t,\tau)$ is a scalar. When $\sigma p(t,\tau)$ is a vector, equality is possible but unlikely.

Far more likely when $\sigma p(t,\tau)$ increases with τ is that the convexity premium gains relative to the risk premium as τ-t increases. In that case, the average yield curve would be hump-shaped, which is exactly what we observe, as Figure 3 illustrates. Figure 3 serves to underscore the fact that the long end of the yield curve contains much information about rate volatility, because it amplifies convexity premiums which are driven by variances. The short end, by contrast, in particular the slope of the yield curve near the zero maturity, contains much information about risk premiums, and therefore about the market price of risk. This information is little affected by convexity premiums because for short maturities (i.e. for small τ-t), $|\sigma p(t,\tau)|$ is small and $|\sigma p(t,\tau)|^2$, which drives the convexity premium, is of second-order magnitude.

Figure 3

Mean zero coupon yield

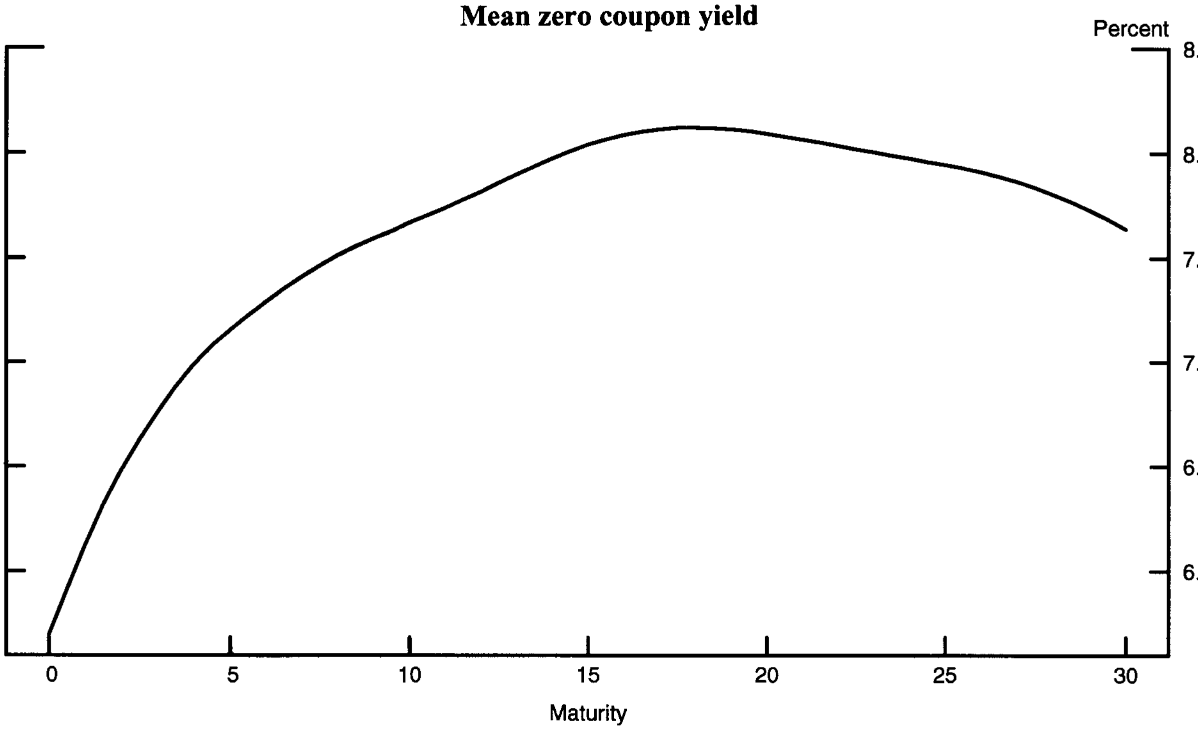

4. Models of the yield curve

4.1 The Vasicek model

A condition *sine qua non* for the empirical relevance of a model is that it be capable of accounting for the hump in the average yield curve that appears in Figure 3. The simplest possible diffusion model of the yield curve is the so-called Merton model, in which the drift and the diffusion of the short rate are both constant:

$$dr(t) = \mu dt + \sigma dW(t).$$

This model does not generate a satisfactory average yield curve because the spot rate is not stationary, so that the average yield curve is not even well defined.

In one of the earliest attempts to use continuous-time diffusions to model the yield curve, Vasicek (1977) proposed to postulate that the spot rate follows an Ornstein-Uhlenbeck process:

$$dr(t) = k[\theta - r(t)]\, dt + \sigma dW(t). \tag{10}$$

With this process, the spot rate is stationary with unconditional mean θ and the drift, which in absolute value is proportional to the deviation between $r(t)$ and θ, always drives the rate toward θ with a constant diffusion σ. The constant k is the speed at which the spot rate reverts to its mean. It is then possible to choose a constant market price of risk λ and generate an average yield curve that looks somewhat like Figure 3.[8]

Despite its ability to generate a realistic average yield curve, the Vasicek model has three substantial defects from an empirical standpoint. First, there is only one source of risk in the model, and thus all bond prices are perfectly correlated, a condition obviously violated by real-world prices of Treasury securities. Second, interest rates can become negative in the model. This feature is acceptable in a model of real rates, but is undesirable in a model of nominal rates. Third, the constancy of the diffusion and that of the market price of risk imply that the volatility of $P(t,\tau)$ is a deterministic function of the maturity τ-t. Under these conditions, the term premiums $\Lambda_y(\tau-t) \equiv \Lambda_y(t,\tau)$ and $\Lambda_f(\tau-t) \equiv \Lambda_f(t,\tau)$ are also deterministic functions of the maturity; that is, for a fixed maturity, term premiums do not change through time. This property of the model does not match the overwhelming evidence that term premiums do vary. We discuss later how to generalize the Vasicek model to overcome these problems, but first we turn to the evidence on varying term premiums.

If term premiums were constant, then $f(t,t+s)$, the forward rate at time t for maturity $t+s$, would predict $r(t+s)$ with a fixed bias equal to $\Lambda_f(s)$, independent of time t. Therefore, fixing s and running either of the two regressions:

$$r(t+s) = \alpha_1 + \beta_1 f(t,t+s) + \varepsilon_1(t),$$

or (subtracting $r(t)$ from both sides)

$$r(t+s) - r(t) = \alpha_2 + \beta_2[f(t,t+s) - r(t)] + \varepsilon_2(t),$$

the same coefficients, namely $\alpha_1 = \alpha_2 = \Lambda_f(s)$ and $\beta_1 = \beta_2 = 1$, should result. But Campbell and Shiller (1991) showed that such regressions typically produce $\beta_1 = 1$ and $\beta_2 < 0$. This second result occurs because $\Lambda f(t,t+s)$ is not equal to a constant $\Lambda f(s)$, but instead varies and is (positively) correlated with $r(t)$. This gives rise to a standard omitted variable problem and the estimate of β_2 is biased. This effect has been well described by Frachot and Lesne (1994) in the present context, which

---

8 The sign of σ is arbitrary because $dW(t)$ is symmetrically distributed around zero. But σ and σp always have opposite signs, because a shock that drives the interest rate up drives the price of discount bonds down. The usual convention assigns a positive sign to σ, so that σp is negative. With this convention, (8) requires $\lambda(t) < 0$ for the yield curve to be upward-sloping.

is testing the weak version of the YTM-EH, as well as by Frachot (1994) in the context of testing the uncovered interest rate parity hypothesis.

In the first regression, the omitted variable does not bias the result, so we get $\beta_1 = 1$, as expected. This is so because both the short rate and the forward rate behave almost like martingales. In addition, the two series are cointegrated, so that in a regression in level the regression coefficient is uniquely determined by the cointegrating vector: it is the coefficient that produces residuals with a small variance, even when there is an omitted variable problem. In fact, one reason to run the regression in the second form is to correct for the unit root problem in the first form. The correction is successful, since both $r(t+s) - r(t)$ and $f(t,t+s) - r(t)$, for fixed $s > 0$, are stationary, but the very success of the correction for cointegration allows the omitted-variable problem to bias the regression coefficient β_2.

As previously noted, whereas in the Vasicek model nominal interest rates can turn negative, in the real world they never do. This feature of the data is itself evidence that volatilities change. When the spot rate is low, its volatility must be sufficiently low to prevent any risk of the rate turning negative. Hence diffusions must somehow depend on the level of rates, and therefore must be stochastic. In other words, a stochastic volatility model can take care simultaneously of two of the three major counterfactual features of the Vasicek model: that rates can turn negative and that term premiums are constant. The remaining defect, namely the perfect correlation among all bonds, requires using more than one source of risk, which is what multifactor models of the yield curve are designed to do.

4.2 The one-factor Cox-Ingersoll-Ross model

Before turning our attention to multifactor models, we discuss the famous one-factor model proposed by Cox, Ingersoll and Ross (CIR, 1985), in which the spot rate is always positive and its volatility changes; as a result, term premiums also change through time. In the CIR model, the short rate has the same drift as in the Vasicek model, but the diffusion is proportional to the square root of the spot rate:

$$dr(t) = k[\theta - r(t)]\,dt + \sigma\sqrt{r(t)}dW(t). \tag{11}$$

This model easily passes the test of being consistent with the average yield curve shown in Figure 3. Its main advantage is that, when the stochastic market price takes the form

$$\lambda(t) = \gamma\sqrt{r(t)}, \tag{12}$$

for some constant γ,[9] then bond prices have a closed-form solution of the form:

$$P(t,\tau) = A(\tau - t)\exp[-B(\tau - t)r(t)], \tag{13}$$

for some specific (but complicated) functions $A(\tau - t)$ and $B(\tau - t)$ that involve all the parameters of the model. Therefore, both the yield and the term premium for the zero-coupon bond maturing at τ are proportional to the current spot rate. As explained by Frachot and Lesne (1994), the CIR model can generate term premiums that vary with the spot rate in a manner roughly consistent with the regression coefficients that Campbell and Shiller obtained.

In the model generated by (11) and (12), the spot rate is the single "factor" affecting the yield curve. All bond prices are perfectly correlated, and they depend on the current level of the spot rate but not on how the spot rate reached this level. For most empirical purposes, both of these properties are too restrictive. Fortunately, it is possible to generalize the model to more than one factor, relaxing the offending restrictions while preserving some of the tractability of the setting.

9 Again, if $\sigma > 0$, then a positive slope of the average yield curve requires $\gamma < 0$.

4.3 Multifactor models

When bond prices are perfectly correlated, their variance covariance matrix is degenerate, with rank equal to 1. Decomposition of an estimated variance-covariance matrix by principal component analysis can then provide some information about the number of factors affecting bond prices and what they correspond to. Such analysis - see for example Litterman, Scheinkman and Weiss (1991) - always finds that there are at least two factors and seldom rejects that there are three. These three factors are commonly associated with the level, the slope, and the curvature of the yield curve, but such factors can be captured by many different sets of measures. What set to choose and how to model their evolution depend in great part on mathematical tractability.

It is possible to generalize the Vasicek model to a multifactor setting without complicating the analytical expressions too much. This can be done by postulating that each of several factors obeys an independent Ornstein-Uhlenbeck process like (10), and the factors add up to the spot rate. A popular way to implement this suggestion in a two-factor model is to take a long-term yield as the first factor and the spread between the spot rate and that yield as the other factor, as in Brown and Schaefer (1993). With this choice, one factor - the yield - represents the level of the yield curve and the other - the spread - represents its slope, which agrees with the evidence from principal component analysis. This procedure overcomes the problem of perfect correlation between bond prices, but introduces new difficulties, which we discuss in more detail below, in the context of the CIR model. The other two problems remain: interest rates can turn negative and term premiums are constant.

The CIR model can accommodate several factors by decomposition of the short rate, in much the same way as the Vasicek model. In fact, Cox, Ingersoll and Ross (1985) were the first to discuss this method. It turns out that, if each factor is driven by an independent square-root process of the form (11), and the price of each risk is also proportional to the square root of the corresponding factor, then the yields on all zero-coupon bonds are linear in the factors, preserving an important feature of (13); models with that property have been called *exponential affine* by Duffie and Kan (1993). The decomposition that Cox, Ingersoll and Ross suggested was that of the nominal spot interest rate into a real spot interest rate and a spot inflation rate. Longstaff and Schwartz (1992) and Chen and Scott (1993) implemented that suggestion, while the latter paper showed how to estimate such a model by maximum likelihood methods.

In Brown and Schaefer (1993), a multifactor generalization of the Vasicek model, the diffusions are assumed to be constant. This simplifies the pricing formulas but does not agree with the data; in particular, interest rates can become negative and term premiums are constant. Square-root diffusions, as in Longstaff and Schwartz (1992) or Chen and Scott (1993), allow term premiums to vary and, because factors cannot switch sign (with appropriate parameter values), the interest rate can stay positive. And yet, matching the data is still a problem. To see why, suppose for example that the factors consist of a long-term rate and a spread, as in Brown and Schaefer. Then neither one factor can switch sign, which (assuming that the sign of the spread is compatible with an upward slope) implies that the yield curve cannot become inverted. In reality, of course, the yield curve sometimes slopes down. Similarly, if one factor is the real rate and the other the spot inflation rate, then the real rate cannot become negative, which is also counterfactual. We are thus inevitably led to models in which the factors do not add up to the spot rate.

Duffie and Kan (1991) count among the set of exponential affine models some in which the spot rate does not equal the sum of the factors. Note, however, that the spot rate is itself a yield, and thus must equal some linear combination of the factors as all yields do. One possibility is for the spot rate itself to be included among the set of factors (in this case, the weight on each of the other factors in the linear combination that gives the spot rate is zero). In one such model, the spot rate follows the stochastic process (11), but the mean θ is not a constant. It is itself random, and follows a similar process

$$d\theta(t) = \xi\left[\overline{\theta} - \theta(t)\right]dt + s\sqrt{\theta(t)}dW_2(t). \tag{14}$$

According to this model, $\theta(t)$ is a mean to which the spot rate tends to revert in the short run. This mean moves through time in a random fashion, but it is stationary and has a mean of $\overline{\theta}$ (ξ and s are fixed parameters). For a fixed level of the spot rate, a higher level of $\theta(t)$ steepens the yield curve, so again it is possible to associate the first factor with a level effect and the second with a slope effect.

Das (1993) examined a CIR model with stochastic short mean (although Das did not restrict the short mean to follow the process (14)). He found that this type of model was able to capture features of the yield curve that previous two-factor models could not capture. In particular, an increase in the short mean of the spot rate produces a steepening of the yield curve, resulting in a new curve uniformly above the old one. On the other hand, the yield curve often becomes more humped, with yields in the middle range moving up while yields at both the short end and the long end of the maturity spectrum move less or even decline. It is this kind of movement that explains why principal component analysis finds that a third factor, curvature, helps to explain yield curve dynamics. This phenomenon seems to call for a model where the diffusion of the spot rate varies somewhat independently of the other factors.

4.4 Chen's model

Chen (1995) has analyzed a three-factor model that incorporates erratic movements in volatility as well as changes in the short-term mean. In that model, the short rate process again obeys (11) and the short-term mean process again obeys (14). In addition, the diffusion parameter σ is now time-varying, so that the complete model is

$$\begin{aligned} dr(t) &= k[\theta(t) - r(t)]\,dt + \sigma(t)\sqrt{r(t)}dW_1(t) \\ d\theta(t) &= \xi[\overline{\theta} - \theta(t)]\,dt + s\sqrt{\theta(t)}dW_2(t) \\ d\sigma(t) &= \varsigma[\overline{\sigma} - \sigma(t)]\,dt + v\sqrt{\sigma(t)}dW_3(t), \end{aligned} \tag{15}$$

where ζ and v are fixed parameters, and $\overline{\sigma}$ is the mean volatility. In addition, the three Brownian motions may be correlated.

Chen was unable to produce a closed-form solution for his general model, but he found one for a closely related special case, and he also showed how to correct the formulas based on the special dynamics to fit the general dynamics to any desired degree of accuracy (the corrections are computationally intensive, however). For the special dynamics, the spot rate process is replaced by

$$dr(t) = k[\theta(t) - r(t)]\,dt + \sigma(t)\,dW_1(t),$$

so that the volatility is not directly related to the level of the short rate; in addition, the three Brownian motions are uncorrelated. One of the drawbacks of this special case is that the spot rate can become negative. But the probability of rates turning negative may be sufficiently low as not to cause problems in practice, and estimating the parameters of the model with the special dynamics is a necessary first step toward analyzing the general dynamics, in which the spot rate is always positive. In any event, the Board work on the yield curve that I referred to earlier is devoted to estimating the parameters of the Chen special dynamics model.

It is approximately correct to associate the stochastic mean and volatility in the Chen model as the slope and curvature factors identified by principal component analysis. Figures 4 and 5

Figure 4

Effect of increase in short mean

(solid line: theta = 0.04; dashed line: theta = 0.06)

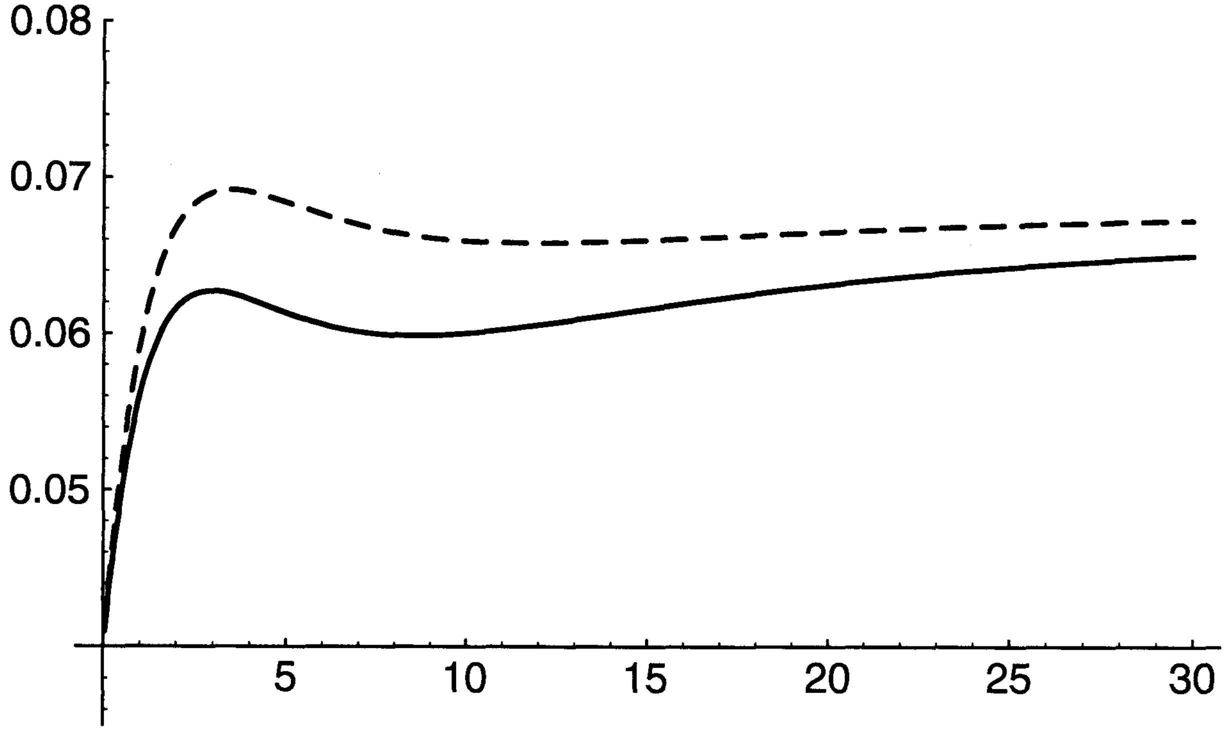

Figure 5

Effect of increase in volatility

(solid line: sigma = 0.05; dashed line: sigma = 0.1)

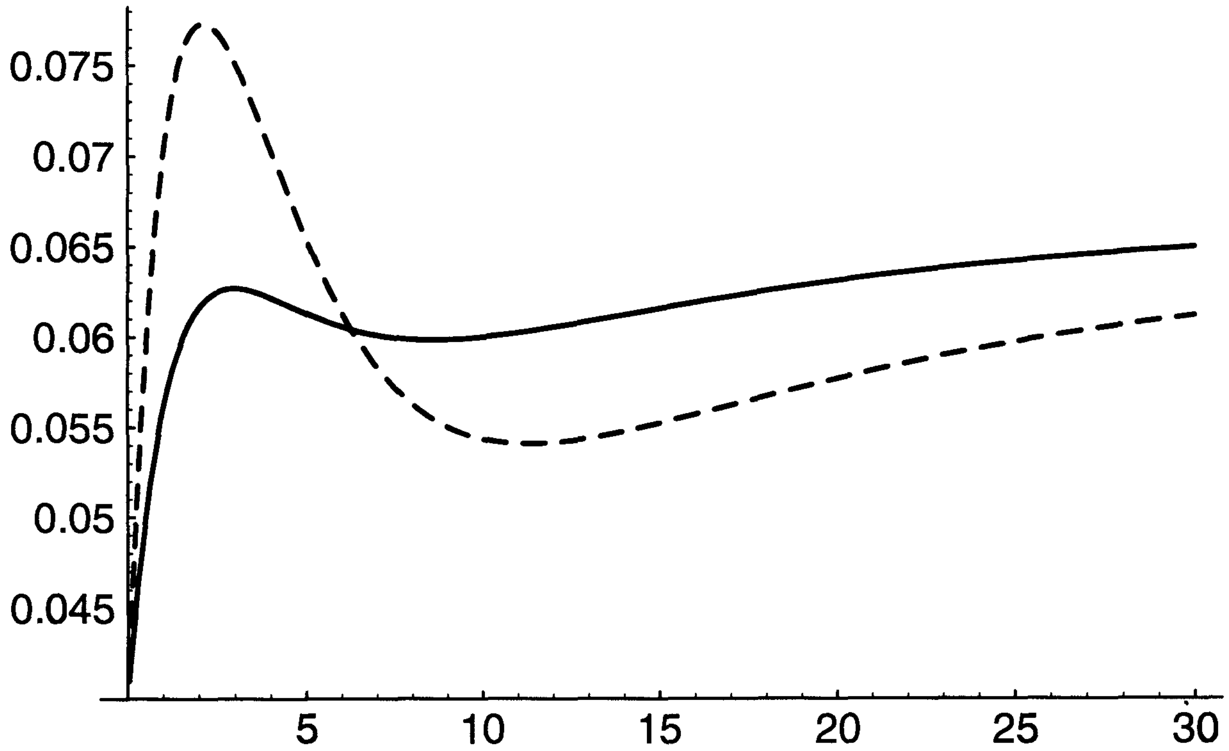

are intended to support this claim. These figures show yield curves generated by the model with arbitrary parameters.[10]

In Figure 4, the spot rate equals 0.04 and the volatility equals 0.05. Then, an increase in the short mean from 0.04 (its unconditional mean) to 0.06 increases the slope of the yield curve at the short end of the maturity spectrum, and the new yield curve is uniformly above the original one.

In Figure 5, the spot rate equals 0.04 and the short mean is also equal to its unconditional mean of 0.04. Then an increase in the current level of the volatility increases curvature, with the short rate unaffected, rates in the medium range above their original level and rates at the long end below their original level. This illustrates our earlier observation that increases in volatility affect yields through both a risk premium and a convexity premium. The risk premium tends to dominate at the short end, where it pulls yields up, while the convexity tends to dominate at the long end, where it pulls them down. To answer an important question, within the context of a structural model which allows volatility changes (such as the Chen model), detecting a change in volatility is relatively straightforward. In principle, it suffices to watch for these telltale changes in the shape of the yield curve. Formal econometric procedures make this possible. I will thus conclude with some remarks on the econometrics.

5. Econometric issues

In any one-factor diffusion model, a time series of observations on the spot rate is in principle sufficient to estimate the parameters of the spot rate process, but not to determine the shape of the yield curve. This is because the market price of risk, which affects the yield curve, has no influence on the spot rate dynamics and therefore cannot be estimated from a time series of the spot rate. On the other hand, the yield curve on any given day contains much information about the market price of risk and other parameters of the model, but is not sufficient to reveal fully the dynamics of the short rate.

In the Vasicek model, for example, bond prices depend on the market price of risk λ; but each time the parameter λ occurs in a price formula, it is added to the mean spot rate θ. It is therefore impossible to separate λ from θ with a single day of data, although the sum $\theta + \lambda$ might be sharply estimated. A time series of the spot rate is necessary to provide an estimate of θ alone, which can then be used to isolate λ. Similarly, in the CIR model bond prices involve λ only through the sum $\kappa + \lambda$. In that model, moreover, it is extremely difficult to estimate the price of risk sharply, because estimates of κ, the speed of mean reversion, are often severely biased, even in large samples, when the true speed is close to zero. A speed of mean reversion equal to zero corresponds to a martingale (a unit root process), and it is well known that it is difficult to reject the hypothesis of a unit root in interest rates.

In multifactor models, the identification and bias problems are compounded by the fact that at least some of the factors are typically not observable. In Das (1993), for example, the short mean is a latent variable; that is also the case in the Chen model, where in addition the diffusion $\sigma(t)$ is also a latent factor. This problem is not as difficult to overcome as it may appear. Pearson and Sun (1994) noted that in exponential affine models, since yields are linear in the factors, n points on the yield curve can in principle be used to uncover the current values of all of the factors in an n-factor model. They used this fact to estimate a two-factor CIR model. In the Chen model, that means that three points on the yield curve (given knowledge of the parameters) are sufficient to uncover the values of the three factors. Therefore, although the instantaneous short mean and volatility are not directly observable, they are easily detectable.

10 These were chosen, for no particular reason, as: $k = 0.4$, $\xi = 0.7$, $\overline{\theta} = 0.04$, $s = 0.1$, $\zeta = 0.1$, $\overline{\sigma} = 0.1$, $v = 0.1$. In addition, the parameters for the market prices of factor risk were -1 for the spot rate, and -2 for both the short mean and the volatility.

But identification issues are as important in a multifactor model as in a one-factor model. With time series of two points on the yield curve, it is possible to estimate the parameters of the random processes that the factors obey in a two-factor model, but it may not possible to estimate the two prices of risk. Pearson and Sun resolved this issue by assuming that these prices were zero. This procedure can deliver correct prices of bonds and derivative securities, but it biases any forecast of future spot rates. For forecasting purposes, it is necessary to use more points on the yield curve than there are factors to uncover, in order to estimate the prices of risk.

In terms of estimation method, some Monte Carlo studies suggest that maximum likelihood based on the correct density dominates more naive methods (see Gourieroux, Monfort and Renault (1993)). But more work needs to be done to confirm this finding in the case of the Chen model. This work is complicated by the fact that the density is a very complicated function of the parameters, so that it is extremely difficult to find the region where likelihood is maximized.

References

Brown, R. and S. Schaefer (1993): "Interest rate volatility and the shape of the term structure", IFA Working Paper 177, Institute of Finance and Accounting, London Business School.

Campbell, J.Y. and R.J. Shiller (1991): "Yield spreads and interest rate movements: a bird's eye view", *Review of Economic Studies* 58, pp. 495-514.

Chen, L. (1995): "Interest rate dynamics, derivative pricing, and risk management", Working Paper, Board of Governors of the Federal Reserve System, January.

Chen, R.-R. and L. Scott (1993): "Maximum likelihood estimation for a multifactor equilibrium model for the term structure of interest rates", *Journal of Fixed Income* 3, pp. 14-31, December.

Cox, J., J. Ingersoll and S. Ross (1981): "A reexamination of traditional hypotheses about the term structure of interest rates", *Journal of Finance* 36, pp. 321-346.

Cox, J., J. Ingersoll and S. Ross (1985): "A theory of the term structure of interest rates", *Econometrica* 53, pp. 363-384.

Das, S.R. (1993): "Mean rate shifts and alternative models of the interest rate: theory and evidence", Working Paper, Stern School of Business, New York University, December.

Duffie, D. and R. Kan (1993): "A yield-factor model of interest rates", Working Paper, Graduate School of Business, Stanford University.

Fisher, M. and C. Gilles (1995a): "Risk, convexity and the shape of the yield curve", Working Paper, Board of Governors of the Federal Reserve System, April.

Fisher, M. and C. Gilles (1995b): "Expectations hypotheses, uncovered interest-rate parity, and the Fisher equation: a unified framework", Working Paper, Board of Governors of the Federal Reserve System, June.

Frachot, A. (1994): "A reexamination of the uncovered interest rate parity hypothesis", Working Paper, Banque de France, January.

Frachot, A. and J.-P. Lesne (1994): "Expectations hypotheses and stochastic volatilities", Working Paper, Banque de France, March.

Gourieroux, C., A. Monfort and E. Renault (1993): "Indirect inference", *Journal of Applied Econometrics* 8, S85-S118.

Heath, D.R., R. Jarrow and A. Morton (1992): "Bond pricing and the term structure of interest rates: a new methodology", *Econometrica* 60, pp. 77-105.

Litterman, R., J. Scheinkman and L. Weiss (1991): "Volatility and the yield curve", *Journal of Fixed Income* 1, pp. 49-53, June.

Longstaff, F. and E. Schwartz (1992): "Interest rate volatility and the term structure: a two-factor general equilibrium model", *Journal of Finance* 47, pp. 1259-1282.

MacDonald, S.S. and S. Hein (1989): "Futures rates and forward rates as predictors of near-term Treasury bill rates", *The Journal of Futures Market* 9, pp. 249-262.

Vasicek, O. (1977): "An equilibrium characterization of the term structure", *Journal of Financial Economics* 5, pp. 177-188.