Implied Interest Rate Skew, Term Premiums, and the "Conundrum"

J. BENSON DURHAM

J. BENSON DURHAM
is an assistant director of the Division of Monetary Affairs and the chief of the Monetary and Financial Market Analysis Section of the Board of Governors of the Federal Reserve System in Washington, DC.
j.benson.durham@frb.gov

Alan Greenspan, then Federal Reserve Chairman, lodged "conundrum" into the lexicon of fixed-income analysts with his February 2005 testimony to Congress, as the substantial drop in global bond yields amid increases in the federal funds rate puzzled many financial market participants. Subsequent analysis, particularly the application of latent-factor affine term structure models (ATSMs) (Kim and Wright [2005]),[1] suggests that the drop in longer-term yields during the last tightening cycle appears to reflect declines in term premiums, as opposed to expected short rates, which actually increased somewhat over the episode. In addition to the general motivation for linking the output of latent factor models to observable variables, the conundrum has unleashed considerable interest in explaining movements in term premiums, defined as the required yield on a bond above the expected path of interest rates through its maturity.

This article extends previous research (Kim and Wright [2005]; Backus and Wright [2007], Durham [2006a]) that addressed time-varying ATSM-based term premiums by focusing on a possible correlate absent from the literature, option-implied interest rate skew. As outlined later in the article, the skew, irrespective of the mean and variance, of investors' interest rate expectations potentially affects required bond yields over the anticipated short-rate path. Briefly, if the risks to rates are skewed to the downside (upside), borrowers will pay a lower (higher) premium to lock in borrowing costs over a given horizon. From the perspective of investors, long-term bonds are more (less) attractive if the distribution of rates is skewed to the downside (upside), because expected capital appreciation is therefore skewed to the upside (downside), all else equal. To invoke informally a consumption-based CAPM approach, skew may capture investors' propensity to pay more for assets that are expected to provide compensation during broader economic downturns.

Some evidence strongly supports this hypothesis—the correlation between 6-month-ahead skew derived from options on Eurodollar futures and the ATSM-based 10-year zero-coupon term premium is positive and statistically significant over the monthly sample from May 1989 through May 2007, as well as alternative periods. In particular, the conundrum does not seem to be particularly problematic for the complete regression model, as the residuals are actually lower on average during this episode, which raises the issue of just why this particular period is so puzzling. Moreover, variation in near-term skew helps account for more of the decline in term premiums during the conundrum period than any other correlate. Specifically, an out-of-sample prediction based on a reduced-form model estimated before the most recent tightening cycle (i.e., using data only through May 2004) suggests that the subsequent variation in skew would correspond to an approximate 35 basis point (bp) decline in the premium between June 2004 and June 2006. The actual net decline was about 70 bps.

There are some strong caveats. For example, the horizons of implied skew and term premiums should largely correspond, but possible liquidity issues and limited data aside, estimates of skew derived from longer-dated interest rate caps generally do not corroborate the result. Also, the estimate of the term premium is based on a Gaussian ATSM in which the skew of the disturbances is necessarily zero, and therefore the correlation perhaps fundamentally contradicts the underlying model. And non-model-based estimates of the term premium do not similarly seem to correlate with skew estimates. These caveats aside, no other variable is robust to such sensitivity analysis and thus the findings with respect to skew do not seem comparatively fragile. Indeed, there seems to be considerable specification bias on both the left and right sides of the equation in reduced-form models of the term premium.

This article next discusses other possible determinates of the term premium and then outlines the arguments regarding implied skew and some baseline empirical results. The conundrum is then discussed in more detail with particular respect to developments in skew during the episode, followed by a description of the sensitivity analyses regarding sample selection, the term structures of skew and term premiums, and alternative measures of the term premium.

WHAT DRIVES TERM PREMIUMS?

Latent factor arbitrage-free models of the term structure are understandably unsatisfying to some, insofar as the underlying factors are by construction unobservable. ATSMs may fit the data well, but they provide no direct economic interpretation or intuition. However, one can get a sense of what economic and financial variables are relevant to yield curve movements by examining the empirical correlates of the output of the model—namely, the expected path of the short rate and term premiums.[2] With respect to the latter, term premiums derived from some ATSMs are a function of the volatility and market price of risk parameters. Therefore, empirically speaking, time variation in the term premium should be related to observable proxies for perceived uncertainty and attitudes toward risk.

To guide broadly the specification search, consider the closed-form expression of the term premium from a (single-factor) Vasicek [1977] model in which the market price of risk is a function of the underlying latent factors. The underlying stochastic process is

$$dX_t = \kappa(\theta - X_t)dt + \sigma dZ_t \tag{1}$$

where X is the underlying factor, κ is the mean-reversion parameter, θ is the long-run mean of X, σ is the volatility parameter, and dZ is the Brownian increment. The market price of risk follows:

$$\lambda_t = \lambda_0 + \lambda_1 X_t \tag{2}$$

And finally, given the closed-form solution to the model, the expression for the term premium—in this case the difference between the forward rate at some horizon, $f(\tau)$, and the expected (instantaneous) short rate,[3] $E[r_T | t]$—is

$$f(\tau) - E[r_T \mid t] = -\frac{\sigma^2}{2\kappa^2}(e^{-\kappa\tau} - 1)^2 + \frac{\sigma}{\kappa}\lambda_t(e^{-\kappa\tau} - 1) \tag{3}$$

Therefore, the term premium is a function of κ, τ, σ, and λ. Leaving aside the fact that in practice some of these are strictly parameters, observable correlates could relate to risk and the market price of risk.[4]

For example, previous research reports an empirical link between interest rate option-implied volatility and ATSM-based term premium estimates (Kim and Wright [2005]; Backus and Wright [2007], Durham [2006a]). Therefore, the specifications of 10-year zero coupon term premiums include the implied volatility derived from options on 10-year Treasury note futures. Notably, this captures the implied volatility of the underlying over the life of the option and not the horizon of the term premium,[5] but nonetheless, the coefficient should be positive.[6]

The use of observed implied volatility measures necessitates two caveats. First, implied volatility can be thought of as expected volatility and a premium for volatility risk and therefore with respect to Equation (3) can be related to either σ or λ. Second, although there are alternatives, the primary measure of the term premium is based on an ATSM in which σ does not vary. Therefore, any correlation between the term premium and implied volatility may seem somewhat coincidental, although again the premium for volatility risk embedded in implied volatility and λ is still time varying.

Also under the general volatility rubric, the model includes the variance of responses from the Michigan survey of short-run inflation expectations. This is strictly a measure of disagreement among respondents and not a measure of individual uncertainty. However, disagreement and uncertainty tend to be positively correlated in survey data. Presumably, the inflation risk premium component of the nominal term premium should be positively correlated with uncertainty about future inflation, and the estimated coefficient should be positive (Durham [2006a]). Of course, uncertainty regarding the outlook for the real economy is also relevant on this score. However, inflation and output dispersion among survey respondents is quite closely correlated, for example, using Blue Chip data.[7]

A number of other variables might track attitudes toward risk or relative preferences for Treasury securities. For example, the implied volatility on the S&P 500 might affect the demand for safe assets, insofar as increased (decreased) uncertainty in the stock market prompts investors to allocate from equities (Treasuries) into Treasuries (equities). Assuming no change in the expected path of short rates, term premiums should decline (increase) and therefore the estimated parameter should be negative. Also, foreign demand for U.S. assets, specifically Treasury securities, might potentially have an effect.[8] There are no obvious choices for proxies, but the correlation between foreign custody holdings of U.S. Treasury securities (as a ratio of Treasury securities outstanding) and the term premium should be negative.

Finally, there is some evidence that term premiums are countercyclical (Backus and Wright [2007]) as uncertainty and/or the market price of risk could conceivably be greater during economic downturns. There are a number of possible proxies—Backus and Wright [2007] document the positive correlation between unemployment and distant-horizon forward rates—but the following regressions include the contemporaneous value of capacity utilization in the regressions[9] and the expected sign is negative.[10]

Turning to some initial results, Regression 1 in Exhibit 1, which covers average monthly data from May 1989 through May 2007, largely reconfirms some of these previous findings. For example, 10-year Treasury note implied volatility and the (dispersion) variance of Michigan survey respondents are both positively correlated with the ATSM-based 10-year zero-coupon term premium, as expected, although the parameter on the survey data is only significant with 10% confidence.[11] The volatility on the S&P 500 and foreign custody holdings of U.S. Treasuries enter with negative coefficients, again as anticipated, that are statistically significant. Finally, the coefficient for capacity utilization has the expected negative sign, but the estimate is statistically insignificant and the R-squared for the model is 0.726.

Before turning to skew, note that this overall result—from the individual coefficients to the overall fit of the model—is somewhat reassuring in the context of a latent factor model; that is, the underlying factors are of course not observable, but the output of the model, at least with respect to term premiums, seems to be consistent with some economic interpretation.[12]

WHY MIGHT SKEW AFFECT TERM PREMIUMS?

The intuition regarding skew and term premiums can be expressed from different perspectives. Consider the perspective of the borrower and think of the term premium as the relative price of locking in borrowing costs over some fixed horizon. If the implied distribution of expected short rates is skewed to the downside, then irrespective of the mean or variance of that density borrowers are more inclined to roll over short-term obligations because the odds are tilted toward lower future funding costs, and therefore the relative price of securing rates over longer horizons drops. Conversely, borrowers would pay more to avoid rolling over debt if the risks were skewed toward higher borrowing costs over the horizon. This implies a positive correlation between skew and term premiums.

The implications are identical, but now consider the perspective of the lender or investor. Longer-term bonds are more attractive, all else equal, if the distribution of expected rates over a given horizon is skewed to the downside, because, simply given the inverse relation between yields and bond prices, capital appreciation is skewed to the upside, all else equal. With negative interest rate skew, investors are willing to pay more to hold longer-dated bonds, bond prices are comparatively higher, and the term premium component of yields is lower. One could also add some consumption-based CAPM flavor to this logic if downside skew in rate

EXHIBIT 1
Regression Models of Term Premium Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Sample:	5/1989–5/2007	5/1989–5/2007	5/1998–5/2007	5/1998–5/2007	5/1998–5/2007	5/1998–5/2007	5/1998–5/2007	5/1998–5/2007	5/1998–5/2007	5/1998–5/2007
Dependent Variables (Term Premium Estimates):	*10-year Zero-Coupon Nominal Term Premium*	*10-year Zero-Coupon Nominal Term Premium*	*10-year Zero-Coupon Nominal Term Premium*	*10-year Zero-Coupon Nominal Term Premium*	*10-year Zero-Coupon Nominal Term Premium*	*1-year Instantaneous Forward Nominal Term Premium*	*FRA Term Premium, 1 Year Ahead*	*Eurodollar Futures Slope, 4 to 5 Years Ahead*	*Treasury Forward Curve Slope, 9 to 10 Years Ahead*	*Excess 5- vs. 1-year bond return, Cochrane and Piazzesi [2005]*
Independent Variables										
10-year Treasury Note Implied Volatility	0.1207	0.0434	0.0029	0.1217	0.1677	-0.0322	0.0636	0.0375	0.0075	-0.0015
	(0.009)**	(0.450)	(0.952)	(0.172)	(0.071)+	(0.488)	(0.300)	(0.002)**	(0.514)	(0.716)
Michigan Inflation Survey Variance	0.0099	0.0122	-0.0011	-0.0054	-0.0081	0.0061	0.0148	-0.0082	-0.0039	-0.0022
	(0.051)+	(0.022)*	(0.837)	(0.507)	(0.322)	(0.297)	(0.233)	(0.003)**	(0.028)*	(0.006)**
S&P 500 Implied Volatility	-0.0169	-0.0079	0.0108	-0.0041	-0.0074	0.0045	0.0160	0.0012	0.0003	0.0008
	(0.045)*	(0.266)	(0.190)	(0.614)	(0.340)	(0.603)	(0.073)+	(0.669)	(0.833)	(0.123)
Foreign Custody Holdings of U.S. Treasuries	-0.0761	-0.0693	-0.0156	-0.0259	-0.0251	-0.0235	-0.0221	0.0075	0.0032	-0.0006
	(0.001)**	(0.001)**	(0.465)	(0.317)	(0.386)	(0.271)	(0.176)	(0.304)	(0.345)	(0.491)
Capacity Utilization	-0.0073	-0.0111	-0.0197	-0.0082	-0.0062	0.0507	0.0080	-0.0283	-0.0071	-0.0076
	(0.813)	(0.689)	(0.548)	(0.861)	(0.895)	(0.172)	(0.763)	(0.016)*	(0.388)	(0.002)**
Eurodollar Implied Skew, 6 months ahead		0.5437	0.7400			0.5330	-0.2069	-0.0030	-0.0357	-0.0088
		(0.007)**	(0.002)**			(0.018)*	(0.343)	(0.954)	(0.281)	(0.515)
Eurodollar Implied Skew, 9 months ahead				0.1968						
				(0.443)						
Eurodollar Implied Skew, 12 months ahead					-0.0016					
					(0.990)					
Constant	2.6197	2.9595	2.4308	1.4258	1.0931	-3.2831	-0.6449	2.1587	0.5868	0.6435
	(0.024)*	(0.006)**	(0.076)+	(0.376)	(0.499)	(0.021)*	(0.686)	(0.000)**	(0.027)*	(0.000)**
Observations	217	217	109	109	109	109	109	109	109	109
Adjusted R-squared	0.726	0.764	0.614	0.459	0.445	0.341	0.424	0.641	0.162	0.461
Durbin-Watson	0.310	0.379	0.496	0.281	0.285	0.346	0.537	0.334	0.282	0.610

p values in parentheses; + significant at 10%; * significant at 5%; ** significant at 1% (based on Newey-West standard errors).

expectations is associated with deleterious macroeconomic outcomes. The investor conceivably pays more for an asset (e.g., a longer-term bond) that hedges (by way of capital appreciation) against unfavorable states of the world.[13]

Turning to measurement, similar to implied volatilities, the skew of the distribution of expected interest rates is derived from options prices, but notably using quotes both at the money and otherwise. The value of an option can be expressed as its discounted expected value under the risk-neutral measure, as in

$$C(F, t; K, T) = \underbrace{e^{-r(T-t)}}_{discount} \int_K^{\infty} \underbrace{p(F, t; F', T)}_{density} \underbrace{(F' - K)}_{payoff} dF' \quad (4)$$

where C is, say, the value of a call option on an underlying futures rate F with strike K, t is the current date, T is the maturity date, r is the risk-free rate, and p is the (risk-neutral) density. Differentiating twice under the integral sign recovers the desired density, following

$$p(F, t; K, T) = e^{r(T-t)} \frac{\partial^2}{\partial K^2} C(F, t; K, T) \quad (5)$$

One can of course recover moments given the density, and this article uses two measures of skew that are closely correlated in practice.[14] Perhaps the most standard calculation, S, follows

$$S = \frac{\mu_3}{\sqrt{\mu_2^3}} \quad (6)$$

where μ_2 and μ_3 are the second and third moments of the distribution, respectively, centered around the mean. The second type of measure resembles octile skew, OS,[15] as in

$$OS = \frac{[Q(0.875) - Q(0.5)] - [Q(0.5) - Q(0.125)]}{Q(0.875) - Q(0.125)} \quad (7)$$

where $Q(0.5)$, for example, denotes the median of expected rates.[16] Note that implied skew is derived from a risk-neutral distribution and that the measure includes information about both expected skew and the market price of skew, neither of which are disentangled in the analysis that follows.

This article uses two derivative instruments. First, options on Eurodollar futures measure near-term skew at the 6-, 9-, and 12-month horizons. To reflect the fact that the U.S. Federal Reserve Bank's Federal Open Market Committee (FOMC) moves in 25 bp increments, the underlying densities are actually point distributions in 25 bp intervals.[17] Exhibit 2 shows the time series of skew at the 6-month horizon, along with the ATSM-based estimate of the term premium. Second, interest rate cap(let)s provide a read on skew at longer horizons, and the underlying distributions are continuous. Exhibit 3 illustrates the implied probability density surface based on interest rate caps on June 28, 2007.[18]

Turning to empirical results, Regression 2 in Exhibit 1 adds the implied skew from the 6-month-ahead Eurodollar implied probability density function (PDF). Notably, the coefficient is safely statistically significant, and the magnitude of the coefficient is considerable. The R-squared values increase from 0.726 to 0.764.

WHERE'S THE CONUNDRUM?

As noted in Exhibit 4, the model that includes skew plugs some of the gap between the predicted and actual values during the so-called conundrum period. Visual inspection suggests that the model does not seem to fit the conundrum period any worse than any other part of the sample. Indeed, the model actually performs comparatively better during this period—the average absolute value of the size of the residual in the complete sample is about 26 bps,[19] but the average absolute value of the residual from July 2004 through June 2006—the period of the most recent tightening cycle—was about 16 bps. Moreover, the average residual (1 bp) is not even negative over this period, and therefore term premiums do not appear consistently to be "too low."

The coefficients from Regression 2 in Exhibit 1 and variation in the underlying independent variables suggest that skew may have played a notable role during this particular period, even though, in general, the contribution of skew to variance explained over the complete sample is perhaps somewhat modest. From the day before the FOMC first raised rates on June 30, 2004, through the day of the last tightening on June 29, 2006, the estimated zero-coupon term premium declined a net 71 bps (from 128 to 57 bps). The estimated parameter suggests that the

Exhibit 2

Implied Skew, 6 Months Ahead (Options on Eurodollar Futures) May 1989–May 2007

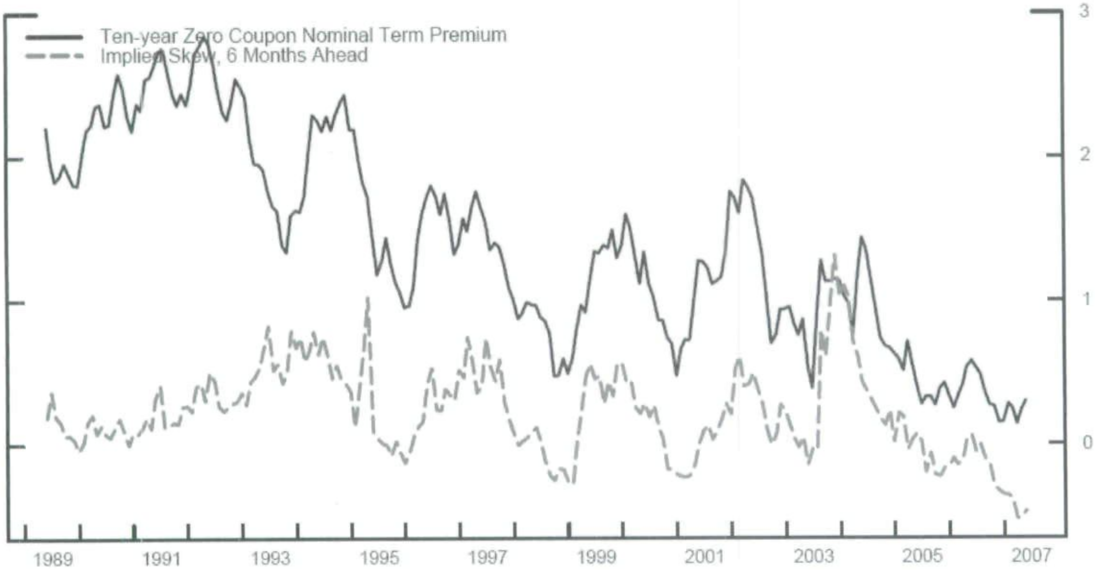

move in skew over the same period implies a 38 bp drop in the term premium,[20] greater than half the total net decline. Also, skew accounts for much more of the decline in the term premium during the period than any other variable. The drop in 10-year Treasury note implied volatility and the increase in capacity utilization during the conundrum period suggest 12 and 5 bp declines in the term premium, respectively, and the underlying movements in Michigan inflation survey variance, S&P 500 implied volatility, and foreign custody holdings of U.S. Treasuries all contrarily suggest very modest increases in term premiums.[21]

One natural suspicion regarding these results is that the coefficients are in-sample and derived using data from the conundrum period. However, estimation of the coefficients using the May 1989 to May 2004 sample, notably before the first tightening of the most recent cycle, nonetheless suggests that the model performs well out of sample from June 2004 through May 2007. The average absolute value of the residual over the last tightening cycle, at 21 bps, is actually lower than the total sample average of 27 bps. Also, the predicted change in the term premium given the variation in skew over the cycle, at 34 bps, is only slightly lower than the in-sample estimate. Therefore, this simple model would have performed reasonably well before the onset of the last tightening cycle and the ensuing fall in long rates.

To digress somewhat, if the conundrum is simply about why the term premium embedded in longer-dated yields declined over the most recent tightening cycle, then strong caveats regarding econometric correlations and true causation aside, the model goes some way toward addressing the puzzle, in general, and skew seems to play an important role, in particular. In fact, judging again from the relative size of the residuals, any conundrum over this particular period, even on an out-of-sample basis, is not readily apparent.

Perhaps the move in term premiums per se is not regarded as particularly curious. Is the conundrum simply that long-term rates fell as the FOMC raised short rates (or removed policy accommodation), which contrasts with the rise in long rates during the 1986, 1994, and 1999 tightening episodes? But, of course, the starting and ending points of tightening episodes are inherently different. And in principle, longer rates could move either up or down or if policy firming is greater or less than expected, respectively.[22] Put more broadly, that movements along the yield curve are not perfectly correlated hardly constitutes a new puzzle for fixed-income analysts.

Could the conundrum also regard the particular constellation of movements in expected short rates (up only slightly) and term premiums (down considerably) over the period? However, at least the (multi-factor)

EXHIBIT 3

Implied Probability Density Function Surface Based on Interest Rate Caps

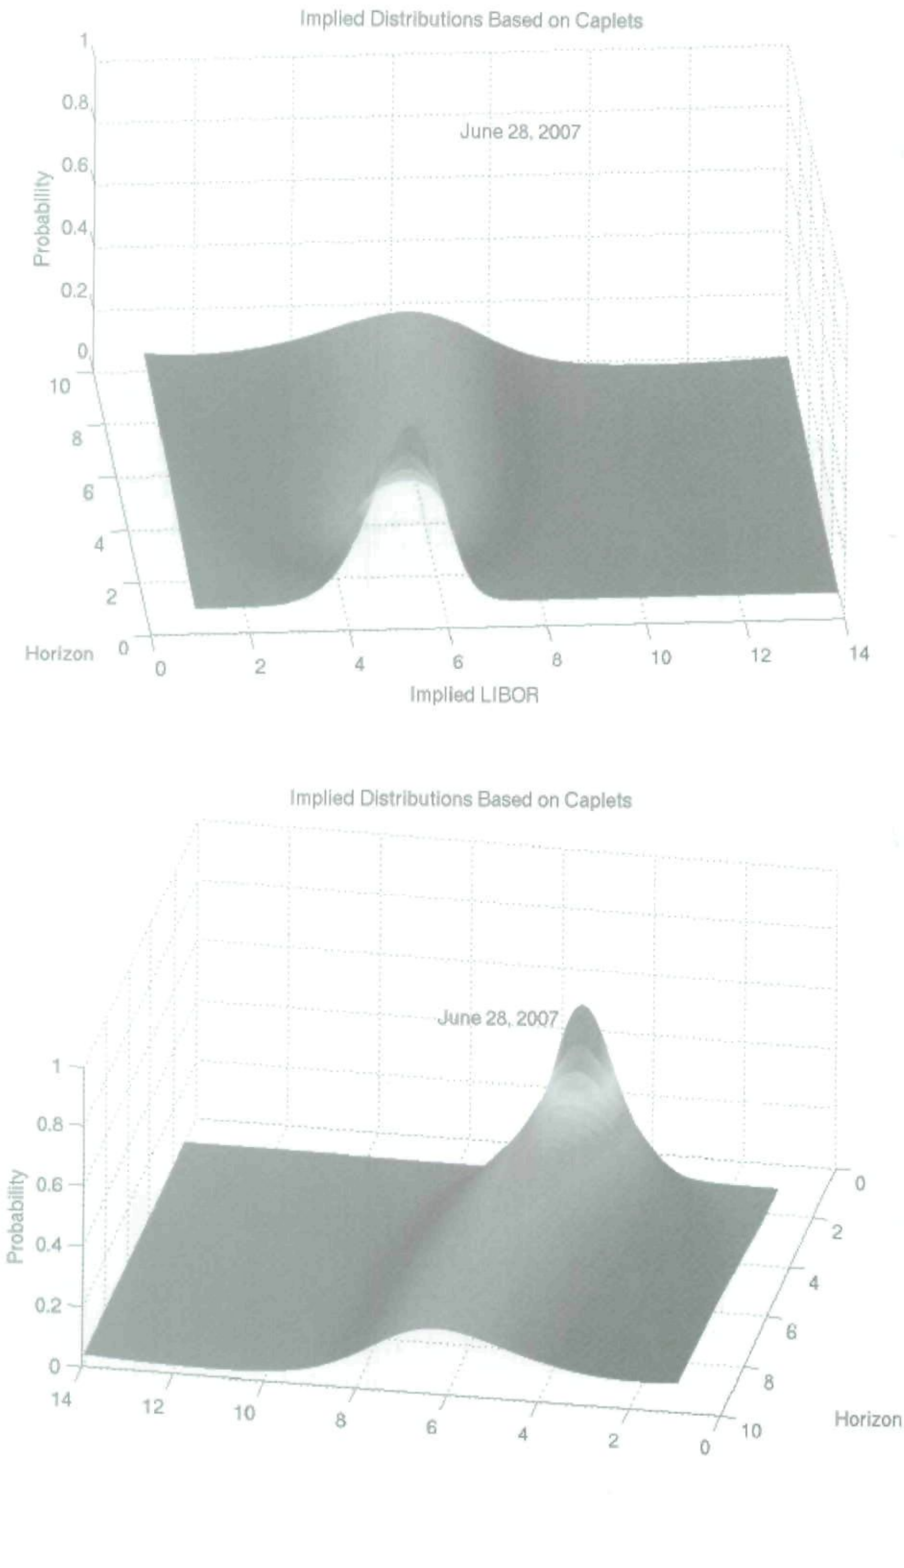

ATSM employed here does not restrict expected short rates and term premiums from moving in opposite directions, which does not seem inconsistent with theory and too uncommon in practice, even at a daily frequency. As on February 27, 2007, for example, investors can revise down sharply their outlook for policy amid greater uncertainty and higher term premiums. Conversely, market participants may revise the path for policy rates upward slightly amid decreased uncertainty and less aversion to risk.

Perhaps analysts should be more precise about where exactly the conundrum lies.[23] To claim that interest rates (or term premiums) are too low is a judgment that ultimately requires some underlying model, reduced-form or otherwise. But if variation in term premiums generally accords with developments in other factors, then possibly the conundrum actually regards the expected path of the short rate; that is, perhaps the puzzle regards the "dogs that did not bark" as opposed to those that did. Whether the expected path of short rates implied by ATSMs or survey data was too low vis-à-vis investors' forecasts for inflation and output during the period is an unanswered question beyond the scope of this article.

SOME SENSITIVITY ANALYSIS AND CAVEATS

Before inferences get too far afield, some sensitivity analysis and caveats regarding the relation between skew and term premiums would be instructive. These regard parameter stability, the horizon of skew, and alternative measures of the term premium.

E X H I B I T 4

Reduced-form Models of Term Premium Estimates May 1989–May 2007

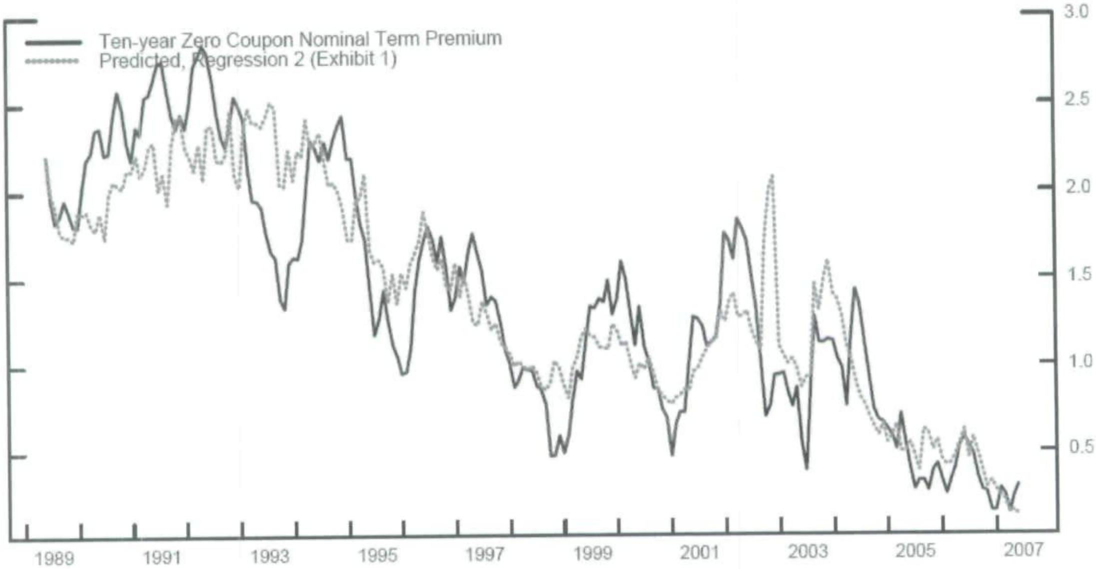

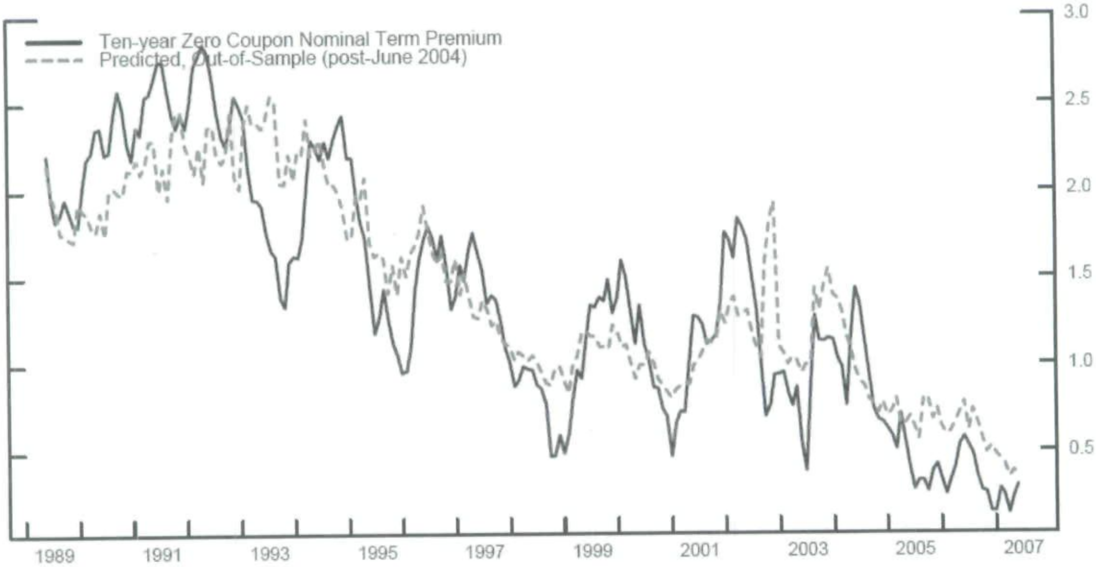

Sample Selection

The relation seems fairly robust to sample selection, particularly with respect to more recent periods. Regression 3 in Exhibit 1 includes data only from May 1998 through May 2007,[24] and the safely significant coefficient on 6-month-ahead skew is actually larger than in the complete sample (0.740 versus 0.545). None of the other possible correlates of the term premium are statistically significant in the regression—only skew seems to matter. In addition, as noted in Row 1 in Exhibit 5, given daily data from May 10, 2005, through June 21, 2007, the (univarite) correlation between the term premium and skew is still significant and highly positive.

EXHIBIT 5

Correlations with 10-Year Zero-Coupon Term Premium May 10, 2005–June 21, 2007 (Daily)

	Skew Horizon	Correlation Coefficient
(1)	6 months	0.6335*
(2)	9 months	0.3704*
(3)	12 months	0.2326*
(4)	5 years	-0.2573*
(5)	10 years	-0.1962*

* *Denotes significance with 5% confidence.*

Horizons of Skew and Term Premiums

The results outlined previously document specifically the relation between skew at the 6-month horizon and 10-year zero-coupon term premiums. True, options at shorter horizons are generally much more liquid, but ideally distant-horizon skew is more relevant for longer-dated term premiums, and near-term skew is more germane to shorter-dated term premiums. To address this issue, some alternative specifications examine whether the result is robust when the horizons for skew and the term premium more closely correspond. For example, the left side of Regressions 4 and 5 in Exhibit 1, which cover the May 1998 to May 2007 sample, include the implied skew from distributions derived from options on Eurodollar futures that expire 9 and 12 months ahead.[25]

Even though these horizons still hardly match that of the term premium, the results cast some doubt on the general relation. Regression 4 produces a positive coefficient for 9-month-ahead skew that narrowly misses the 10% confidence level, and the parameter for 12-month-ahead skew from Regression 5 is statistically insignificant. Also, as Rows 2 and 3 of Exhibit 5 indicate, the daily univariate correlations between skew at these horizons and the term premium are statistically significant, but the magnitude is notably less than that for the 6-month horizon.

Although liquidity may be an issue and data are limited, a more direct assessment uses skew from implied distributions derived from interest rate caplets at distant horizons. Data are only available from May 10, 2005, through June 21, 2007, but as Rows 4 and 5 in Exhibit 5 suggest, the correlation between the term premium and implied skew 5 and 10 years ahead is actually slightly negative.

The negative correlation, although based on few data, is somewhat curious, perhaps especially given that 6-month-ahead skew from options on Eurodollar futures is still positively correlated with ten-year-ahead estimates from caplets, with a pairwise coefficient of about 0.38 over the sample. There are also differing views among practitioners as to whether one can extract fundamental signals from far-dated out-of-the-money interest rate caps, as flows might be more related to transient technical hedging demands as opposed to views on the distribution of expected policy rates in the very distant future. Indeed, at least over this sample, caplet-based skew at more distant horizons is just as volatile as skew in the near term,[26] which presents a "sensitivity puzzle" under the assumption that expectations for the path for policy are more stable than estimates of the equilibrium policy rate well beyond the current business cycle. Nonetheless, taken at face value, and for these immediate purposes, limited available data suggest that skew at longer horizons supports the hypothesis.

Although longer-dated term premiums may be of primary interest in the context of the conundrum, one can also more closely match horizons by considering near-term skew and term premiums at the shorter end of the curve. Toward that end, the left side of Regression 6 includes the ATSM-based estimate of the instantaneous forward (Treasury) term premium one year ahead. The coefficient on 6-month-ahead skew is statistically significant, but interestingly the magnitude is somewhat less than that for the 10-year zero coupon term premium (0.533 versus 0.740). Moreover, perhaps more germane to the expected path for monetary policy, Regression 7 includes on the left side the year-ahead forward term premium from an ATSM calibrated exclusively to near-term forward rate agreements.[27] This model produces a statistically insignificant coefficient for skew. Therefore, these models curiously suggest that longer-dated term premiums respond more than shorter-dated premiums (if at all) to near-term skew.

Alternative Measures of the Term Premium

Another critical caveat regards the particular measure of the term premium under consideration. In closed form, the term premium from the Gaussian 3-factor ATSM is a function of volatility and the market price of volatility risk, notably not skew. All disturbances are normally distributed and have zero skew by definition, and therefore the positive correlation between 6-month-ahead Eurodollar option-implied skew and 10-year zero-coupon term premiums, while perhaps intuitive and robust econometrically, is formally perplexing within the confines of this particular term premium estimate.

Therefore, one key question is whether alternative non-model-based measures of the term premium also produce a statistically significant relation. These measures are perhaps generally not preferable to the ATSM-based estimate, largely given the absence of the no-arbitrage constraint, but at least they do not rest on some of the key (Gaussian) assumptions that are particularly restrictive in the case of skew. There are surely other worthy metrics, but such alternatives include the slope of the Eurodollar futures curve between 4 and 5 years ahead (Regression 8 in Exhibit 1), the slope of the instantaneous Treasury forward curve between 9 and 10 years ahead (Regression 9),[28] and the predicted excess return of 5-year over 1-year Treasury securities from Cochrane and Piazzesi [2007].[29]

Even though each measure notably declined during the conundrum period,[30] the covariance matrix of these various estimates over the May 1998 to May 2007 sample suggests considerable uncertainty about measures of the term premium in general because the (pairwise) correlation coefficients between the ATSM 10-year zero-coupon measure and these alternatives is 0.220, –0.307, and 0.273, respectively. Perhaps unsurprisingly, therefore, Regressions 8, 9, and 10 produce statistically insignificant coefficients for 6-month-ahead skew. Thus, the result does in fact seem sensitive to the specification of the term premium on the left side, and the model that explicitly suggests that skew is not priced actually produces the robust result.

CONCLUSION

Insofar as the conundrum in interest rates over the most recent Federal Reserve tightening cycle reduces to a puzzle about the level of term premiums, one measure of (option-implied) interest rate skew goes some way in addressing the issue. This is not to suggest that skew "explains" the conundrum; the query might simply shift to an examination of why expected skew and the market price of skew moved considerably over the period. This article makes no progress in delineating changing expectations from changing risk tolerances with respect to skew, although the same can be said for implied volatility. Also, the general performance of the reduced-form model does not imply that accounting for movements in term premiums is trivial—explaining movements in yields over any period is hardly a precise science.

The empirical relation between skew and term premiums seems sensitive both to the horizon of skew and to the precise measure of the term premium. But before jettisoning the notion that skew matters, note that few other potential correlates of the term premium are likely to be too robust under this analysis. For example, as Exhibit 1 indicates, the proxy for implied volatility, which is also subject to the same issues related to matching horizons, is only statistically significant in 3 of the 10 models. In addition, there are both statistically significant positive and negative coefficients across the 10 regressions for inflation survey variance, and S&P 500 implied volatility, foreign custody holdings of U.S. Treasuries, and capacity utilization are not consistently robust. Therefore, comparatively speaking, the results for skew are not exceptionally fragile. The fact that term premium estimates are not very closely correlated is also noteworthy on this score. Unlike, say, empirical studies of excess stock returns or economic growth, concepts for which there is little disagreement regarding measurement, there should also be a healthy dose of suspicion about the left side of the equation when regressing term premiums on possible correlates. To say the least, there is room for empirical progress, and meanwhile inferences should tread carefully.

With respect to future research and analysis, a more ideal measure of the term premium might be derived from a no-arbitrage model of the term structure that somehow includes a time-varying skew parameter, which in turn relates in an intuitive way to term premiums. Such a model is unlikely to be very tractable, but meanwhile perhaps some non-Gaussian models might be somewhat better suited for these purposes.[31]

ENDNOTES

Without implication, the author thanks Jim Clouse, Dan Covitz, Spencer Dale, Wesley Phoa, and Jonathan Wright for helpful comments. The views presented are solely those of the author and do not necessarily represent those of the Federal Reserve Board or its staff.

[1]The model follows Kim and Orphanides [2005] and, in turn, Dai and Singleton [2002]. The model is a 3-factor Vasicek model in which the market price of risk is a linear function of the underlying latent factors. Survey data on interest rate expectations inform the estimates of the parameters. The model is calibrated to zero-coupon Treasury yields across the term structure, in turn, derived from a smoothed Svensson yield curve.

[2]In addition, one can examine the empirical observable correlates of the underlying factors, which in this case seem closely related to the level, slope, and curvature of the term structure.

[3]Given that the term premium is a function of the underlying factors, this formulation relaxes the expectations hypothesis assumption that term premiums are constant.

[4]There is no extensive attempt in this analysis to net out the convexity premium from the term premium and therefore the impact of skew on the convexity component. However, previous empirical results suggest that the first term in Equation (3), the so-called convexity term or Jensen's inequality term, is generally smaller than the second. Moreover, at least given limited daily data from January 2, 2004, though June 29, 2007, the correlation between skew (and implied volatility) and the term premium is robust to conditioning on a proxy for the convexity premium—the spread between an estimated (constant-horizon) Eurodollar futures curve and a corresponding forward LIBOR curve derived from forward rate agreements (FRAs) and swap rates.

[5]A more appropriate measure for a zero-coupon or instantaneous forward term premium would be the implied volatility on an underlying asset that matches the horizon of the term premium.

[6]There is no effort here to classify observable variables as related exclusively to risk, on the one hand, and the price of risk, on the other. Rather, Equation (3) simply provides some background for the specification search. One could attempt to parse more precisely a given move in term premiums into its risk and market price of risk components by using a stochastic volatility model in which both volatility and the market price of risk were functions of the underlying factors. The fixed parameters and the factors would determine the relative contributions of σ and λ to a given movement in term premiums.

[7]The Blue Chip dispersion measure might also be preferable particularly given the longer horizon of the forecasts. However, these data are only available twice a year.

[8]See Warnock and Warnock [2006].

[9]The unemployment rate is not statistically significant with the expected sign in models that correspond to those in Exhibit 1 (i.e., that replace capacity utilization with the unemployment rate).

[10]This list of possible correlates is not necessarily exhaustive. For example, pension reform has received considerable attention among analysts because the imperative to match marked-to-market assets with liabilities would presumably increase demand for long-dated assets and thereby compress term premiums. However, although there are proxies for pension solvency, the prospects for reform are difficult to measure empirically. Also, some analysts have suggested that mortgage convexity hedging has the potential to amplify, but notably not generate, movements in interest rates.

[11]Statistical significance is based on Newey-West standard errors.

[12]Analysts and researchers face a choice between the so-called macro-finance models that include observable factors and fit the yield curve comparatively poorly versus latent factor models that (within sample) fit the term structure well. Latent factors models are perhaps preferable if the objective is to decompose yields into expected rates and term premiums without substantial fitting errors.

[13]In the case of stocks, Harvey and Siddique [2000] argue that expected returns should reward systematic skew and formulate an asset pricing model that incorporates conditional skew. They report that systematic skew helps explain the cross section of expected U.S. stock returns conditioned on other stock market anomalies such as book-to-market, size, and momentum.

[14]Without continuously quoted strikes, one must interpolate between prices or implied volatilities. For example, one method follows Shimko [1993] and involves estimating a quadratic relation between implied volatility and strike. The resulting function is substituted for σ in Black's formula (for futures or caplets) and differentiated twice, as in Equation (5), to obtain the density.

[15]More precisely, skew from the Eurodollar option-implied densities follows:

$$\frac{[Q(0.9)-Q(0.5)]-[Q(0.5)-Q(0.1)]}{Q(0.9)-Q(0.1)}-1$$

[16]Both measures have their advantages, but octile skew is usefully bounded between +1 and −1.

[17]One can easily construct a point distribution from a continuous distribution. The underlying Eurodollar distributions actually map into expectations for the federal funds rate by using LIBOR versus federal funds basis swaps. The distributions also represent an interpolation between the relevant contracts to obtain constant horizons of 6, 9, and 12 months ahead. I make no adjustment for the fact these options are American rather than European. Note that interest rate caps are sequences of European call options.

[18]The steps required to extract implied densities from interest rate caps are as follows. "Flat" implied volatilities on caps, with strikes from 1% to 14% that expire between 1 and 10 years ahead, are routinely quoted. A continuous flat volatility term structure is interpolated using a Nelson-Siegel functional form. The corresponding spot implied volatilities for each strike level and each (quarterly) horizon for caplets that expire from 6 months to 10 years ahead are derived using a simple bootstrapping technique. Given the spot implied volatility surface (across strikes) for each horizon, the estimation of the PDF largely follows Shimko [1993], but with no assumptions about the functional form of the density outside the range of quoted strikes.

[19]This compares with the average term premium of about 139 bps and a range from 14 to 284 bps from May 1989 through May 2007.

[20]The implied distribution of expected short rates was generally skewed to the upside prior to the first tightening and became more symmetrical, on net, by the time of the last firming. This pattern might be expected given the very low level of interest rates and the "considerable period" of monetary policy accommodation that preceded the cycle. Of course,

the zero bound on nominal rates makes the implied distribution of expectations truncated toward the left tail.

[21]The precise period to make comparisons during the last tightening cycle is somewhat arbitrary. As an instructive alternative, consider the last day before the first tightening (June 29, 2004) and the day on which the term premium hit its lowest level, 16 bps, during the period (August 31, 2005). The coefficient from Regression 2 and the change in skew corresponds with a 41 bp drop in the term premium compared to the total net decline of 112 bps between those two days. The move in 10-year Treasury note implied volatility over that period corresponds with an 8 bp decline in term premiums, and the other variables on the right side imply very small effects.

[22]Even if expected short rates increased in the near term but declined further out the term structure, one could spin a story in which market participants anticipated the FOMC to overshoot.

[23]A broader etymology of the word *conundrum* might be relevant in this context. According to *The American Heritage Dictionary of the English Language* Fourth Edition [2000], conundrum implies "a dilemma" as opposed to a puzzle per se. Indeed, the decline in longer-term yields as the FOMC increased the funds rate created a policy dilemma in that, on the one hand, if the decline reflected perceptions of economic weakness (i.e., lower expected short rates), then the FOMC would be disposed toward an easier policy stance, all else equal. On the other hand, if the drop owed mainly to diminished term premiums, then less restrictive financial conditions would auger for tighter monetary policy. A reliable decomposition of yields into expected rates and term premiums would "solve" this conundrum and prescribe the appropriate policy action.

[24]This particular sample is dictated by data availability on other skew horizons and term premium estimates.

[25]These estimates are positively but certainly not perfectly correlated with skew at the 6-month horizon, as the correlations are 0.873 and 0.784, respectively.

[26]Over the daily sample period, the standard deviation and the range of skew are greater for caplets that pay in 10 years compared to those that pay in 1 year.

[27]The parameter estimation and calibration of the model on the left side of Regression 7 differs from Kim and Orphanides [2005] in that no survey data are used and the model is calibrated to FRAs instead of zero-coupon Treasury yields (in turn, derived from a Svensson yield curve). Calibration to FRAs is perhaps preferable with respect to near-term policy expectations given that FRAs correlate more closely with the policy rate (i.e., the federal funds rate) than Treasury bills.

[28]The simple assumption underlying these slope-based measures is that expected policy at comparatively long horizons beyond the current business cycle is assumed to be constant, and therefore the spread between nearly adjacent, distant-horizon forward/futures rates is largely due to term premiums.

[29]The author thanks Jonathan Wright for providing updated estimates from Cochrane and Piazzesi [2007]. Not that this alternative term premium is germane to the 5- and not the 10-year horizon. Therefore, any discrepancy in results could also relate to this difference in maturity.

[30]This suggests more confidence in the stylized fact that a substantial portion of the decline in yields during the conundrum owed to term premiums.

[31]For example, although skew is not priced, under a jump-diffusion model in which jumps are drawn from a mixture of Gaussian distributions (Durham [2006b]), the term premium may be a function of the weight between positive and negative jump distributions or the relative variances of those densities.

REFERENCES

Backus, David K., and Jonathan H. Wright. "Cracking the Conundrum." Unpublished manuscript, 2007.

Cochrane, John H., and Monika Piazzesi. "Bond Risk Premia." *American Economic Review*, 95 (2007), pp. 138–160.

Dai, Q., and K.J. Singleton. "Expectations Puzzles, Time-Varying Risk Premia, and Dynamic Models of the Term Structure." *Journal of Financial Economics*, 63 (2002), pp. 415–441.

Durham, J. Benson. "An Estimate of the Inflation Risk Premium Using a Three-Factor Affine Term Structure Model." Federal Reserve Board Finance and Economics Discussion Series No. 42, 2006a.

——. "Additional Analytical Approximations of the Term Structure and Distributional Assumptions for Jump-Diffusion Processes." *Journal of Fixed Income*, 15 (2006b), pp. 61–73.

Harvey, Campbell R., and Akhtar Siddique. "Conditional Skewness in Asset Pricing Tests." *Journal of Finance*, 55 (2000), pp. 1263–1295.

Kim, Don, and Athanasios Orphanides. "Term Structure Estimation with Survey Data on Interest Rate Forecasts." Federal Reserve Board Finance and Economics Discussion Series No. 48, 2005.

Kim, Don, and Jonathan Wright. "An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates." Federal Reserve Board Finance and Economics Discussion Series No. 33, 2005.

Shimko, David. "Bounds of Probability." *Risk*, 6 (1993), pp. 33–37.

Vasicek, O. "An Equilibrium Characterization of the Term Structure of Interest Rates." *Journal of Financial Economics*, 5 (1977), pp. 177–188.

Warnock, F., and V. Warnock. "International Capital Flows and U.S. Interest Rates." NBER Working paper Series No. 12560, 2006.