Commentary

John H. Cochrane

The paper by Glenn Rudebusch, Brian Sack, and Eric Swanson (2007) is an impressive survey of several literatures concerned with monetary economics and interest rates. It is well done, so I think the best thing for me to do is to highlight what I think are the central points and to give my views on those points.

WHAT "CONUNDRUM"?

Figure 1 presents the federal funds rate and 1- to 15-year forward rates through the past two recessions. This comparison lets us easily consider to what extent the recent behavior of long-term forward rates represents an unusual experience or not.

My first reaction to Figure 1 is that the patterns are strikingly similar. Short-term yields and forwards decline, spreads widen, and then yields and forwards recover as spreads tighten again. In both episodes there is a little blip on the way down in which long-term yields and forwards rise much more than short-term ones, despite no movement in the funds rate (late 1992 and 2002). In both episodes there is an event on the way up in which all yields and forwards increase sharply (late 1994 and 2004).

The main difference between the two episodes is that the rise in the federal funds rate in 2004-06 is much smoother and more predictable and long-term forward rates, in particular the 10-year rate, falls while the funds rate is rising. Though long-term forwards decline overall in both recoveries (1994-96 and 2004-06), the earlier experience includes a blip up in all rates through 1995, which is later reversed. This experience is missing in the second period. This remaining difference is Greenspan's "conundrum."

The difference is already small. Furthermore, because the rise in the funds rate was much steadier in the later episode, the behavior of *market* rates relative to the funds rate (which reflects different behavior by the Fed) is even less different across the two episodes than the overall behavior of interest rates. If one regards long-term rates as dynamically driven by the federal funds rate, it's not obvious that there is *any* difference in the behavior of markets.

Even if the later period is different, why is it puzzling? First, long forwards *should* fall when the Fed tightens. This is exactly how the world is *supposed* to work. Tighter policy now means lower inflation later, and thus *lower* nominal rates in 10 years. There is no model or estimate anywhere in which the Fed can raise real rates for 10 years without reducing inflation. Prices are not that sticky! In 1994, the opposite nearly one-for-one *rise* of long forwards with rises in the federal funds rate was viewed as a conundrum for just this reason. The main, somewhat convoluted, story used to explain the 1994 events is that an interest rate rise communicates bad news about inflation from the Fed to the markets—information

John H. Cochrane is a professor of finance at the Graduate School of Business at the University of Chicago and a research associate at the National Bureau of Economic Research. The author acknowledges research support from the Center for Research and Security Prices and from a National Science Foundation grant administered by the National Bureau of Economic Research. A draft of this paper with color graphics is available at http://faculty.chicagogsb.edu/john.cochrane/research/Papers/.

Federal Reserve Bank of St. Louis *Review*, July/August 2007, *89*(4), pp. 271-82.

Figure 1

Federal Funds Rate and Selected 1- to 15-Year Forward Rates Through Two Recessions

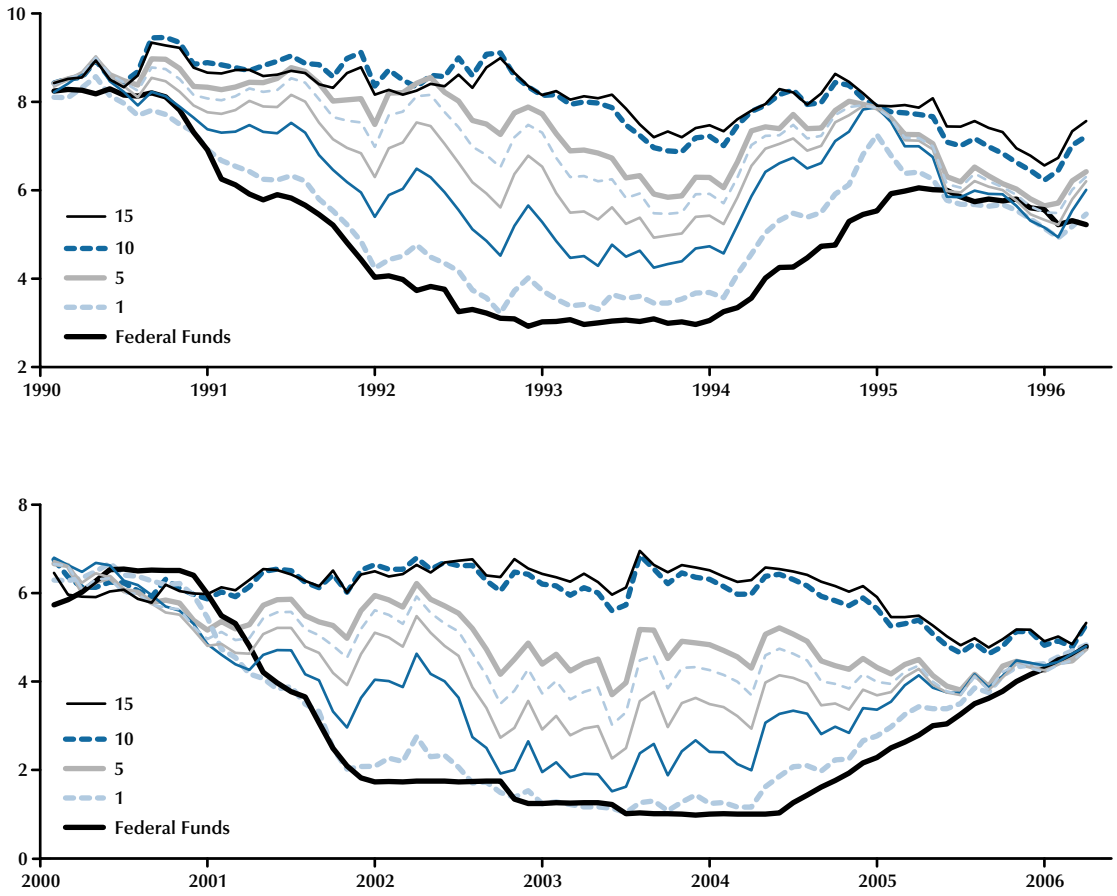

SOURCE: Forward-rate data are from Gürkaynak, Sack, and Wright (2006).

that for some reason the markets did not already have, of course. Greenspan himself echoed this view in 1994[1]:

> In early February, we thought long-term rates would move a little higher as we tightened. The sharp jump in [long] rates that occurred appeared to reflect the dramatic rise in market expectations of economic growth and associated concerns about possible inflation pressures.

Of course, this is a simplistic discussion. A tightening has to be unanticipated in order for it to lower forward rates through this channel, and evidence from other sources, such as the Treasury inflation-protected securities mentioned by Chairman Greenspan or foreign interest rates, also bears on the issue. Still, where did anyone get the idea that monetary policy should control long-term rates and that it is puzzling if long-term rates do not "respond" positively to tightening? The natural benchmark predicts exactly the opposite, if any, effect.

Second, to the extent that the decline in forward rates represents a cyclical or secular decline in term premia, that decline also is perfectly natural. Term premia, like all risk premia, *should*

[1] I owe the quote to Gallmeyer et al. (2007).

decline as we come out of recessions, and have done so in every past recession. Even negative term premia are not a puzzle—they *should* be negative. In a world with stable inflation, interest rate variation comes from variation in *real* rates; and, in such a world, long-term bonds are safer investments for long-term investors. Rolling over *short-term* bonds runs the "reinvestment risk" that short-term (real) rates will change, so short-term bonds should bear the burden of any bond risk premia. We expect only a *positive* term premium in a world with unstable inflation and relatively constant real rates, such as the 1970s. Because short-term rates adapt quickly to inflation changes, rolling over short-term bonds has less risk to a long-term investor than does buying only long-term bonds in this environment.

In sum, were I a Fed Chairman testifying to Congress with the plots of Figure 1 in hand, I would be tempted to point out that, far from a "conundrum," the world is finally behaving exactly the way it should—and so is the central bank. The increased transparency and predictability of operating procedures, seen in the steadiness of the rise in funds rates in 2004-06 verses the less predictable rise in 1994-96, has communicated to the markets the Fed's steadfastness in controlling inflation. We are moving to the sensible world of negative risk premia, which is exactly what we should see once markets understand that inflation is vanquished forever. The conquest of inflation has removed an unnecessary risk premium for long-run investors and issuers of long-dated nominal bonds. I don't necessarily believe all this, of course, but it would be awfully tempting to make this argument were I defending the Fed's actions before a congressional committee. The "conundrum" is Greenspan: Why did he say anything else?

Finally, it is academics' job to remind policy debaters of basic economics, so I think we should pounce anytime somebody says something like "[the] decline in the term premium…is financially stimulative and argues for greater monetary policy restraint." Every price reflects both supply and demand. Low interest rates can reflect a lack of good investment projects as easily as they can reflect an abundance of savings. To take a local example, low housing prices in East St. Louis do not seem to be particularly "stimulative."

DECOMPOSING THE YIELD CURVE

My grumpy comments about "conundrum" and the "stimulative" effects of low prices notwithstanding, this episode *does* highlight the importance of splitting the yield curve into expected future rates and risk premia and of understanding the dynamic structure of risk premia and their macroeconomic underpinnings. Here Rudebusch, Sack, and Swanson provide a very nice summary of the state of the art.

I think the bottom line is that we know less than we think about this decomposition and far less than the pronouncements in policymakers' quotes imply. The paper can be read as a comprehensive survey of one failure after another. Here, let me give two quick, and I hope memorable, points in this litany of ignorance.

Levels, Differences, and Standard Errors

I learned two important lessons while Monika Piazzesi and I (2006) investigated this kind of decomposition. First, *how you specify trends, cointegration, etc.—which the data say very little about—is overwhelmingly the most important issue in driving the decomposition of the long-maturity end of the yield curve.* Second, *the standard errors are very large.* For these reasons alone, any statements decomposing the recent experience of forward rates into changes in expected interest rates versus declining term premia are subject to huge uncertainty.

To see this point, let's try the simplest approach to decomposing the yield curve. I run a vector autoregression (VAR) of five forward rates on their lags (I use the Fama-Bliss data available from the Center for Research and Security Prices, and I use a three-month moving average of forward rates on the right-hand side, which Piazzesi and I (2005) find improves forecasts by mitigating measurement error):

Figure 2

March 2006 Forward Curve and Expected One-Year Rates from a VAR in Levels and from a Cointegrated VAR

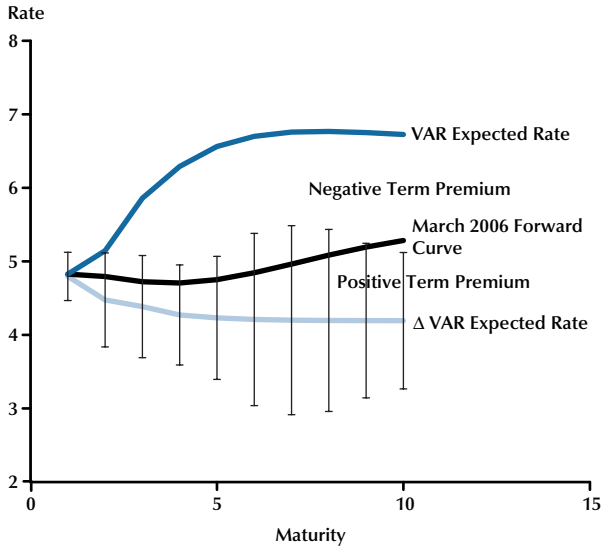

NOTE: The one-standard-error bars are computed from a direct regression forecast, $X_t' \text{cov}(\hat{\beta}) X_t$, using Hansen-Hodrick correction for serial correlation due to overlap.

$$\begin{bmatrix} y_{t+1}^{(1)} \\ f_{t+1}^{(2)} \\ \vdots \\ f_{t+1}^{(5)} \end{bmatrix} = A + B \begin{bmatrix} y_t^{(1)} \\ f_t^{(2)} \\ \vdots \\ f_t^{(5)} \end{bmatrix} + \varepsilon_{t+1},$$

where

$f_t^{(n)}$ = forward at time t for loans from $t + n - 1$ to $t + n$

$y_t^{(1)}$ = one-year rate at time t.

We can use this VAR to generate forecasts at each date of future one-year rates, leaving ("estimating") the term premium as a residual,

$$f_t^{(n)} = E_t\left(y_{t+n}^{(1)}\right) + rp_t^{(n)}.$$

You don't *have* to estimate fancy term-structure models to decompose the yield curve into expected interest rates and a risk premium.

Figure 2 presents the results, evaluated on March 2006 (the last point in my data sample), in the line labeled "VAR expected rate." The line captures a lot of common opinion: It says that interest rates are expected to rise gently over the next few years, leaving a negative term premium, which is puzzling until you think through the economics of long-term bond investing in a low-inflation world. This kind of decomposition also says that much of the recent decline in forward rates comes from the term premium rather than changes in expected long-term rates.

This all seems very sensible. However, Figure 3 examines the same calculation over a longer time interval. The lines represent, at each date, expected one-year rates one, two, three, etc., years in the future: that is, for $E_t(y_{t+k}^{(1)})$ for $k = 1,2,3,\ldots$ at each t. The graph dramatically makes the point that long-horizon expected one-year rates calculated by this method simply reflect reversion to the mean. The 6.25 percent asymptote in Figure 2 represented no specially sophisticated regression forecast; it was simply the sample interest rate.

There is nothing logically or econometrically wrong with this conclusion, but do we really believe it? For example, in 1980, this decomposition says that everyone knew interest rates would decline from 16 percent back to an unconditional mean of a bit over 6 percent, and rather rapidly, so the then-flat yield curves represented very large risk premia for holding long-term bonds. But did people really believe inflation would be tamed, or did perhaps the flat yield curves of the time really represent a good chance that inflation would re-emerge? Similarly, perhaps the sample mean is now too high an estimate. Our data come from inflation and its conquest. Perhaps it is sensible now to think a "structural shift" has happened, so the long-run mean should be a good deal less than 6.25 percent.

As an alternative, let us try a forecast that ignores this "level" information. On a statistical basis, forward rates are clearly best modeled by a single common trend that has a root that is near if not equal to 1 and stationary spreads around

Figure 3

Multiperiod One-Year-Rate Forecasts from a VAR in Levels

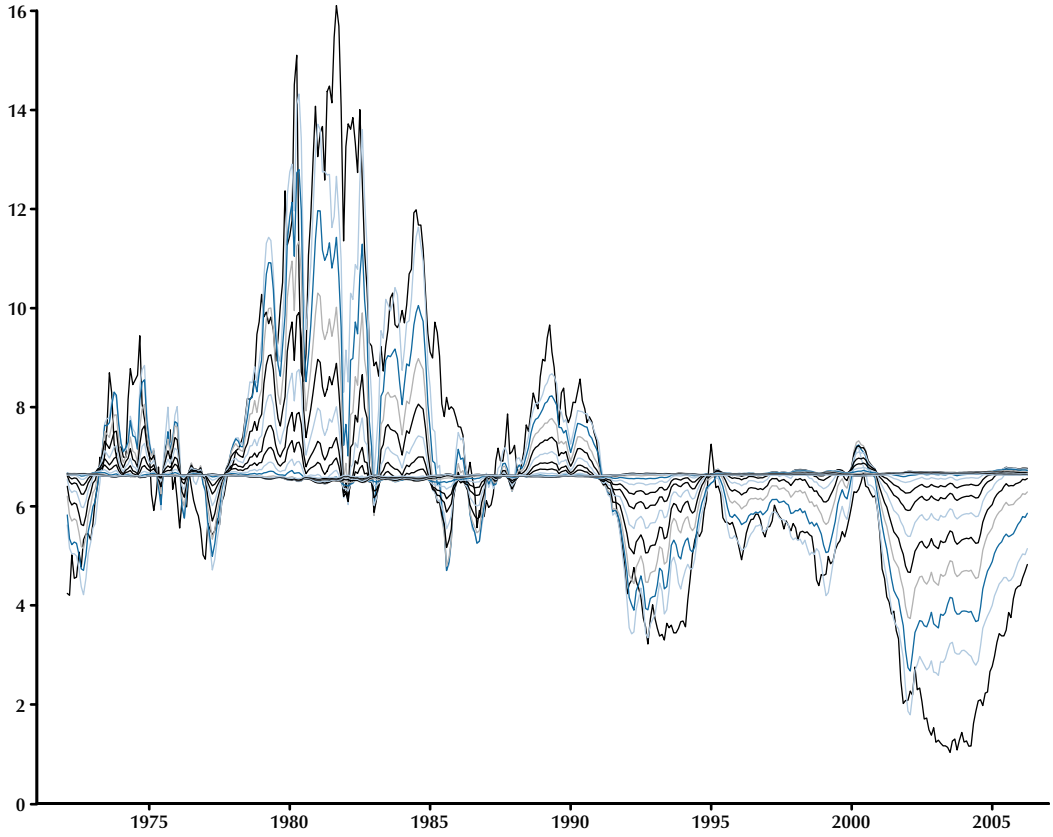

NOTE: Current one-year rate and expectations of one-year rates one, two, three, etc., years in the future, calculated by a simple VAR.

that trend. I estimate a VAR imposing that restriction:

$$\begin{bmatrix} y_{t+1}^{(1)} - y_t^{(1)} \\ f_{t+1}^{(2)} - y_{t+1}^{(1)} \\ \vdots \\ f_{t+1}^{(5)} - y_{t+1}^{(1)} \end{bmatrix} = A + B \begin{bmatrix} f_t^{(2)} - y_t^{(1)} \\ \vdots \\ f_t^{(5)} - y_t^{(1)} \end{bmatrix} + \varepsilon_{t+1}.$$

This is equivalent to simply running forecasting regressions that set to zero a coefficient on the level of interest rates:

$$\text{Before}: f_{t+1}^{(n)} = a^{(n)} + B_{n,1} \times y_t^{(1)} + B_{n,2}\left[f_t^{(2)} - y_t^{(1)}\right] + B_{n,3}\left[f_t^{(3)} - y_t^{(1)}\right] + ;$$

$$\text{now}: f_{t+1}^{(n)} = a^{(n)} + 0 \times y_t^{(1)} + B_{n,2}\left[f_t^{(2)} - y_t^{(1)}\right] + B_{n,3}\left[f_t^{(3)} - y_t^{(1)}\right] + \ldots$$

Intuitively, we still allow information such as "the yield curve is upward sloping" to forecast interest rate changes. We ignore information such as "interest rates are low" to tell us interest rates will rise.

Figure 4 presents expected one-year rates over time by this method. You can see the *huge* difference. One-year rates are certainly not being forecast to revert to a constant unconditional mean! In particular, the flat yield curves of the 1980s are not now interpreted to reveal huge risk premia plus expected declines in interest rates. This is

Figure 4

Multiperiod One-Year-Rate Forecasts from a VAR in Levels

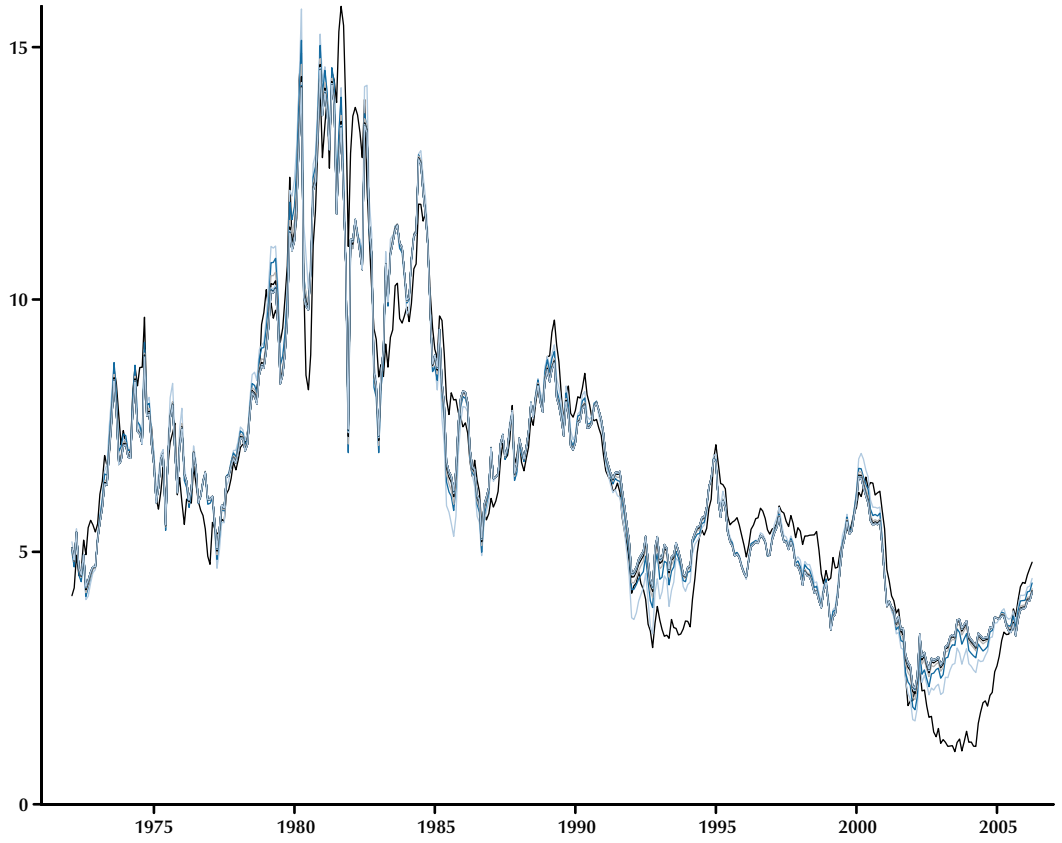

NOTE: Current one-year rate and expectations of one-year rates one, two, three, etc., years in the future, calculated by a VAR that imposes a single common trend in forward rates.

not a pure random-walk model, and there is still some forecastability left. For example, the steeply upward-sloping yield curves of 2003-04 do forecast substantial rises in short rates.

Figure 2 includes the March 2006 one-year rate forecast from this method, in the line labeled "Δ VAR Expected Rate." This is also a sensible forecast. Because we no longer use the information that the current one-year rate is slightly below its sample mean, we are left only with slope information. The unusually flat slope of the forward curve means, in this forecast, that interest rates will *decline* somewhat, so that the term premium is still somewhat positive. However, Figure 2 shows that long-term interest rate forecasts by this method have been rising in recent years; so, the decline in forward rates since 2004 is attributed even more to declining term premia by this method than by the VAR in levels.

Can statistics help us? Alas, no. Testing for unit roots, cointegration, etc., and imposing the resulting structure on the analysis is not fruitful. One naturally wants to think about "structural shifts," changing means, and so forth, and these will be even more imprecisely estimated in now-shorter samples. It is certainly true that the dominant root of a persistent set of variables is estimated with downward bias, so the actual reversion to the mean is slower than the VAR in levels indicates, but whether that mean makes

any sense in the first place is not something statistics can really help us with.

Can fancier models help us? In particular, most of the term-structure literature does not look at simple VAR forecasts. Instead, it estimates the parameters of "affine models." To think about what these can do, it's useful to have a specific example in front of us, so here is the Ang and Piazzesi (2003) model that we use in Cochrane and Piazzesi (2005 and 2006). A vector of state variables X_t follows an AR(1) process; the stochastic discount factor is an exponential function of the state variables, with "market prices of risk" (loadings of M on shocks to X) that also depend on the state variables,

$$X_t = \mu + \phi X_{t-1} + \Sigma \varepsilon_t \tag{1}$$

$$\begin{aligned} M_t &= \exp\left(-\delta_0' X_t - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\varepsilon_t\right) \\ \lambda_t &= \lambda_0 + \lambda_1' X_t. \end{aligned} \tag{2}$$

Assuming the shocks ε are i.i.d. normal with unit variance, we can find this model's prediction for bond prices,

$$P_t^{(n)} = E_t\left(M_{t+1}M_{t+2}...M_{t+n}\right) \tag{3}$$

$$p_t^{(n)} = A_n + B_n X_t, \tag{4}$$

and then yields and forward rates. Inverting (4), we can reveal the "state variables" from bond prices, yields, or forward rates. Thus, this model becomes a structured factor model in which a large collection of prices, yields, or forward rates are described in terms of a few linear combinations of those same prices, yields, or forward rates.

But, underlying the whole thing, we see a VAR(1) in yields, prices, or forward rates—just as we have been estimating all along! Thus the *only* way the affine model can give us *any* different answers from those of the ordinary least squares–estimated VARs above is if the structure of market prices of risk means that we use information in the cross-section of bond prices to infer something about the dynamics. In general, this is not the case. In Cochrane and Piazzesi (2005) we show how to *construct* market prices of risk, λ, from a given discrete-time VAR(1) to turn it into an affine model. Thus, in general, the affine model lives on top of a VAR estimate of long-term forward rates and adds *nothing* to it. (There remains the possibility that by *restricting* or *modeling* market prices of risk, λ, in sensible ways, one obtains information about the VAR (1), and this is the point of our 2006 paper. But this is not [yet] a common idea, and its success lies in the believability of a priori restrictions on λ.)

In addition, once we have settled on a specification, we have to wonder how much sampling uncertainty in estimating the parameters translates into uncertainty about the forecasts. To address this question in a simple and transparent way, I run direct forecasting regressions,

$$\begin{aligned} &y_{t+k}^{(1)} - y_t^{(1)} = \\ &\left[1 \quad f_t^{(2)} - y_t^{(1)} \quad f_t^{(3)} - y_t^{(1)} \quad ... \quad f_t^{(5)} - y_t^{(1)}\right]\beta_k + \varepsilon_{t+k}, \\ &y_{t+k}^{(1)} - y_t^{(1)} = X_t'\beta_k + \varepsilon_{t+k}, \end{aligned}$$

where the second equation defines notation. I find the covariance matrix of $\hat{\beta}_k$, including a Hansen-Hodrick correction for serial correlation due to overlapping data, and then I calculate the error as

$$\sigma_t^2\left[\hat{E}_t\left(y_{t+k}^{(1)} - y_t^{(1)}\right)\right] = X_t' cov(\hat{\beta}_k, \hat{\beta}_k')X_t.$$

This is the error in the measurement of expected interest rates due to sampling uncertainty in the coefficients that comprise the regression forecast. It is not the forecast error—that is, it is not a measure of how large $\sigma^2(\varepsilon_{t+k})$ is. The one-standard-error bars in Figure 2 present this calculation. The term premium is not statistically significant, and the large difference between the two specifications is barely two standard errors. The Hansen-Hodrick correction for serial correlation is undoubtedly optimistic—at the right-hand end of the graph we're forecasting interest rates 10 years ahead in 45 years of data—so the true sampling uncertainty is undoubtedly a good deal larger.

Now, understanding that large roots and common trends, which often must be specified a priori, are crucial to long-term forecasts and that long-run forecasts are subject to enormous sampling uncertainty is not news. However, as I read it, this sensitivity is not at all considered by the

literature that uses affine models to compute long-term yield-curve decompositions. We are usually treated only to one estimate based on one a priori specification, usually in levels, and usually with no measure of the huge sampling uncertainty. Needless to say, the usual habit of estimating 10-year interest rate forecasts by extrapolating models fit to *weekly* or monthly data is no help, and possibly a hindrance. The 520th power of a matrix is a difficult object to estimate.

In sum, when a policymaker says something that sounds definite, such as "long-run forward rates have declined, while interest rate expectations have remained constant, so risk premia have declined," he is really guessing, and we really have no idea whether this is a fact.

Measuring Risk Premia

We also know a good deal less about long-term risk premia than we think we do. Quotes such as those at the beginning of the paper suggest that risk premia are well *measured* if perhaps poorly understood. Nothing of the sort is true. We may have a decent handle on *one*-year risk premia, as surveyed in the paper and the subject of my next set of comments, but the 10-year forward-rate premium reflects not only *this* year's expected excess bond returns, but this year's expectations of *next* year's expected returns, and so on and so forth. If you like equations, an easy one in which to see this point is

(5)
$$y_t^{(n)} = \frac{1}{n}\left[y_t^{(1)} + E_t\left(y_{t+1}^{(1)}\right) + E_t\left(y_{t+2}^{(1)}\right) + \ldots + E_t\left(y_{t+n-1}^{(1)}\right)\right] + \left[E_t\left(rx_{t+1}^{(n)}\right) + E_t\left(rx_{t+2}^{(n-1)}\right) + \ldots + E_t\left(rx_{t+n-1}^{(2)}\right)\right],$$

where $y^{(1)}$ = one-year yield and $rx_{t+1}^{(n)} = r_{t+1}^{(n)} - y_t^{(1)}$ = excess returns. The first term is the expectations hypothesis. The second term is the risk premium, and you see that the risk premium depends on *future* expected excess returns, not just on *current* expected excess returns.

Now, if expected excess returns lived off in their own space, moving away in response to shocks and then recovering without relation to the rest of the yield curve, then, yes, there would be one "risk premium" that accounts for expected excess returns, as well as long-horizon forward and yield-curve risk premia. Alas, this is not the case. Today's level, slope, and curvature have strong power to forecast next year's expected excess returns. (Characterizing these dynamics is a major point of Cochrane and Piazzesi, 2006.) We can easily be in a situation that this year's expected excess return, $E_t rx_{t+1}$, is large and positive, while future expected excess returns, $E_t(rx_{t+k})$, are strongly negative, so the risk premium in the yield curve can be negative as well. The one-year expected excess return can be positive while the 10-year forward rate is below its corresponding expected one-year rate. It is precisely by such differences in expected *future* risk premia that the two decompositions shown in Figure 2 can produce forward rate premia of different signs, despite the same initial return risk premium.

In sum, there is no single "risk premium." There is a full-term structure of return risk premia, which moves over time in interesting and still poorly measured ways. Sure statements that risk premia have moved down over time do not reflect any solid and independent measurement.

FORECASTING, TERM PREMIA, AND MACROECONOMICS

One of the major contributions of the Rudebusch, Sack, and Swanson paper is the empirical work linking bond risk premia and macroeconomics. By restating the points in my own way and slightly disagreeing with some conclusions, I think I can usefully highlight this important part of the paper.

Naturally, I like the Cochrane and Piazzesi (2005) measurement of the risk premium, so I'll focus my comments on that paper. Briefly, we noticed that regressions of excess returns on ex ante forward rates follow a nearly exact one-factor structure: That is, that regressions

$$rx_{t+1}^{(n)} = \alpha_n + \beta_{n,1} y_t^{(1)} + \beta_{n,2} f_t^{(2)} + \ldots + \beta_{n,5} f_t^{(5)} + \varepsilon_{t+1}^{(n)}$$

almost exactly follow

$$rx_{t+1}^{(n)} = b_n\left(\gamma_0 + \gamma_1 y_t^{(1)} + \gamma_2 f_t^{(2)} + ... + \gamma_5 f_t^{(5)}\right) + \varepsilon_{t+1}^{(n)}$$
$$= b_n\left(\gamma' f_t\right) + \varepsilon_{t+1}^{(n)},$$

where the last equality defines notation. A single "return-forecasting factor," $\gamma' f_t$, describes expected excess returns of all maturities. Longer-maturity bonds' expected excess returns move more, and shorter-maturity bonds' expected excess returns move less, but they all move in lockstep. Thus, we estimate the common "return-forecasting factor" by running a single regression of average (across maturity) returns on all forward rates:

$$\overline{rx}_{t+1} = \frac{1}{4}\sum_{n=2}^{5} rx_{t+1}^{(n)}$$
$$= \gamma_0 + \gamma_1 y_t^{(1)} + \gamma_2 f_t^{(2)} + ... + \gamma_5 f_t^{(5)} + \varepsilon_{t+1},$$

where the first equality defines notation. Sensitive to "levels" issues, we obtain nearly identical results by ruling out a level effect:

$$\overline{rx}_{t+1}^{n} = \gamma_0 + \gamma_2(f_t^{(2)} - y_t^{(1)})$$
$$+\gamma_3(f_t^{(3)} - y_t^{(1)}) + ... + \gamma_5\left(f_t^{(5)} - y_t^{(1)}\right) + \varepsilon_{t+1}.$$

The coefficients γ have a pretty tent shape. This measure of bond risk premia values *curvature* in the forward curve, not *slope* in the forward curve.

Figure 5 shows how this works and the connection between macroeconomics and bond risk premia: In January 2002 (shown by the first vertical lines in panels A, B, and D), the recession and interest rates have just finished their stage of steep decline, as seen also in the unemployment rate (panel D). The forward curve is upward sloping, but it is also very curved (panel C). The *curved* forward rate, interacting with the tent-shaped γ, is the sign of risk premia. This means (statistically) that the upward slope will not be soon matched by rises in interest rates, so the greater yields on long-term bonds are (risky) profit for investors. The risk premium (panel B) is very high. In fact, this prediction is borne out: Interest rates do *not* rise for several years, so investors who bought long-term bonds in January 2002 profited handsomely for a few years.

By contrast, consider January 2004. Now, the forward curve still *slopes up* substantially (panel C), but it is no longer particularly *curved*, so the tent-shaped γ coefficients no longer predict much of a risk premium (panel B). Now, the upward-sloping yield curve *does* signal rises in interest rates; the expectations hypothesis is working; returns on long-term bonds will be no higher (on average) than those on short-term bonds. Again, this prediction is borne out. This time, interest rates do rise. This is a repeated statistical pattern, working the same way in many previous recessions.

Having digested what term-premium forecasts are and how they work, we see that the graphs show several patterns seen in more-formal regressions. First, *the term premium* ($\gamma' f$ here, as well as other measures in the paper) *drives out slope variables for forecasting bond excess returns*. Previously, Fama and Bliss (1987), Campbell and Shiller (1991), and others found that measures of the term-structure slope forecast excess returns. Yes, we see the slope is high in 2002, when long-term bond holders turn out to make money. But it is also high in 2004 when they don't. When you put the slope and the curvature of the forward rate together in a multiple regression, the curvature measured by γ wins out. The slope seemed to forecast bond returns because it was correlated with the curve measure. (See, for example, Table A3 of the appendix to Cochrane and Piazzesi, 2005.)

Second, *the term premium is high in the depths of a recession*. In Figure 5, this is measured by the association of the term premium (panel B) with unemployment (panel D). The association is even stronger in previous recessions. In macroeconomic terms (that's why we're here), this is natural. The risk premium is high at the early stage of a recession, a time in which investors don't want to hold risk of any kind. Stock prices are low, predicting higher-than-average stock returns; interest rates are low relative to foreign interest rates, predicting high returns for holding exchange-rate risk. By January 2004, however, the recession is over, the period of growth and rising interest rates has set in, and everybody knows it. It's not a surprise that the premium for holding risk during recessions has vanished.

Figure 5

Macroeconomics and the Yield Curve

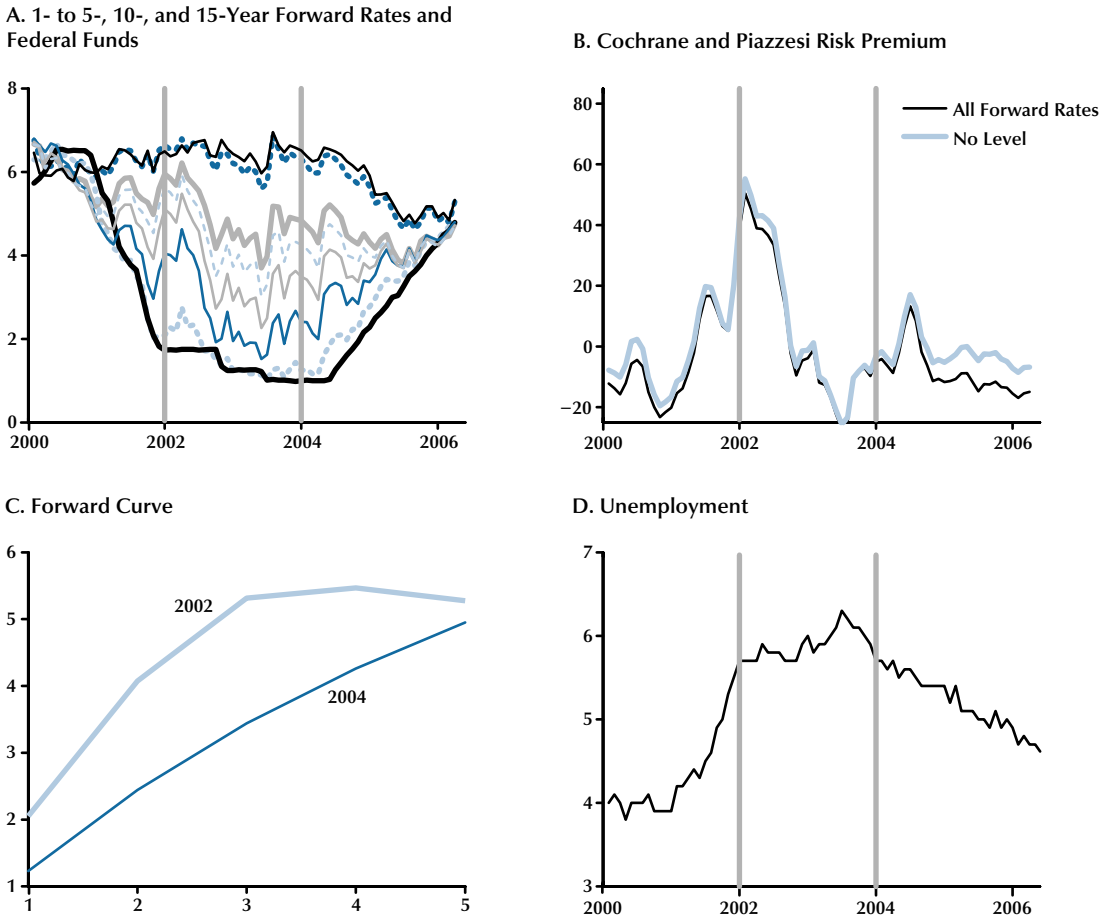

NOTE: Vertical lines mark interesting dates. Panel A: federal funds rate and 1- to 15-year forward rates through the previous recession. Panel B: Cochrane and Piazzesi (2005) measures of the bond risk premium. Panel C: forward curves on the two indicated dates. Panel D: unemployment rate.

SOURCE: Data for yields and forward rates past a five-year maturity are from Gürkaynak, Sack, and Wright (2006).

Third, and the major point of the authors' paper (as I see it), the *slope* of the yield curve drives out the *risk premium* for forecasting 1-year gross domestic product (GDP) growth. We see this in panel B of Figure 5: The *risk premium* is high in 2002, when GDP is not about to grow. The risk premium is low in 2004, when GDP is about to grow. The slope is high in both times. Thus, the *slope* carries GDP forecast power, and the slope, purged of its correlation with the risk premium, forecasts GDP even better.

This point is made in their paper in the regression of Table 2, last column, which I take to be the central result:

$$y_{t+4} - y_t = 0.38(4.22) + 0.96(5.62)(exsp_t - exsp_{t-4}) - 0.59(-1.93)(tp_t - tp_{t-4}) + \varepsilon_{t+4},$$

where y = GDP; *exsp* is the expectations-hypothesis component of the 10-year rate; *tp* is the term premium component of the 10-year rate, as in (5); *t*-statistics are in parentheses; and the sample is from 1962 to 2005.

The authors (p. 261) say of this regression, "The coefficient on the risk-neutral expectations component of the yield curve slope [0.96] is now larger and more statistically significant than in any of the earlier specifications," which is the same interpretation I gave in discussing the figure.

The authors also say (p. 261), "More importantly for this paper, we find that the estimated coefficient on the term premium is now negative and (marginally) statistically significant. According to these results, a decline in the term premium tends to be followed by *faster* GDP growth—the opposite sign of the relationship uncovered by previous empirical studies." I read the evidence differently: Rather than accept a marginally significant coefficient with the wrong sign, it seems to me the right lesson is that the second coefficient is zero. The *slope* of the yield curve forecasts GDP growth, but not risk premia. The *curvature* of the forward curve measures risk premia, but not GDP growth. Risk premia are high precisely when we are not sure whether the recession is over.

Table 8 of Ang, Piazzesi, and Wei (2006) runs the same sort of regression. At a four-quarter horizon, they find

$$y_{t+4} - y_t = a + 1.15(5.00)EH_t - 0.47(0.30)RP_t + \varepsilon_{t+4},$$

where EH = expectations hypothesis; RP = risk premium in the 20-quarter term spread (i.e., the terms in (5), estimated from a macro-affine model); and t-statistics are in parentheses. In this slightly different specification, they confirm the huge significance of the expectations-hypothesis term, but find an insignificant contribution due to the risk-premium term.

This view dovetails with the other side of risk premia that Monika Piazzesi and I (2006) have recently started investigating. From the basic asset-pricing relation, $1 = E_t(M_{t+1}R_{t+1})$, and (2), we can write

$$(6) \qquad E_t(rx_{t+1}^{(n)}) + \frac{1}{2}\sigma_t^2(rx_{t+1}^{(n)}) = cov(rx_{t+1}^{(n)}, \varepsilon_{t+1}')\lambda_t.$$

Note the absence of an n index on λ. The point of this equation is that expected excess returns on each bond must be earned in compensation for, and in proportion to, the covariance of that bond's return with macroeconomic shocks, ε. So far, we have been talking about the left-hand side: What models or state variables drive variation over time in expected excess returns? Now, it's time to start working on the right-hand side: What are the shocks? Piazzesi and I find that the term premium is almost exactly earned entirely in compensation for shocks to the level of the term structure. The prices of risk, λ, corresponding to other term-structure shocks are essentially zero.

In particular, expected returns seem *not* to be earned in compensation for "slope" shocks. This finding lets us start to think about macroeconomics. Whatever the macroeconomic determinants of bond risk premia, they must be variables with "level" effects on the term structure. This observation quickly rules out monetary policy, whose shocks typically raise short rates while lowering or leaving unchanged long rates—a "slope" shock.

CONCLUDING COMMENT

In sum, I think we are headed to a view that slope movements in the yield curve, which are related to monetary policy, are also related to expected GDP growth, as seen in the usual impulse-response functions. But slope movements do not signal risk premia (left-hand side of (6)), nor does covariance with monetary policy shocks generate real risk premia (right-hand side of (6)). Term premia are large in the early phases of recessions, when it's not clear how long the recession will last; they are revealed by the curvature of the forward rate, and they are earned in compensation for macroeconomic risks that correspond to shocks in the level of the yield curve.

Of course, we have no economic models of any of these fascinating statistical regularities. This point is made clear in the brilliant survey of total failures that occupies a large part of the paper. First, what are the *macroeconomic* state variables that drive variation in expected returns? What exactly are the times and states of nature in which expected returns are high? Second, expected returns are generated by the covariance of returns with *macroeconomic* shocks. What are these macroeconomic shocks? These questions are, as ever, the Holy Grail of macro-finance.

REFERENCES

Ang, Andrew and Piazzesi, Monika. "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables." *Journal of Monetary Economics*, July 2003, *50*(5), pp. 745-87.

Ang, Andrew; Piazzesi, Monika and Wei, Min. "What Does the Yield Curve Tell Us about GDP Growth?" *Journal of Econometrics*, March/April 2006, *131*(1-2), pp. 359-403.

Campbell, John Y. and Shiller, Robert J. "Yield Spreads and Interest Rate Movements: A Bird's Eye View." *Review of Economic Studies*, May 1991, *58*(3), pp. 495-514.

Cochrane, John H. and Piazzesi, Monika. "The Fed and Interest Rates—A High Frequency Identification." *American Economic Review*, May 2002, *92*(2), pp. 90-95.

Cochrane, John H. and Piazzesi, Monika. "Bond Risk Premia." *American Economic Review*, March 2005, *95*(1), pp. 138-60.

Cochrane, John H. and Piazzesi, Monika. "Decomposing the Yield Curve." Unpublished manuscript, University of Chicago, 2006.

Fama, Eugene F. and Bliss, Robert R. "The Information in Long-Maturity Forward Rates." *American Economic Review*, September 1987, *77*(4), pp. 680-92.

Gallmeyer, Michael F.; Hollifield, Burton; Palomino, Francisco and Zin, Stanley F. "Arbitrage-Free Pricing with Dynamic Macroeconomics Models." Federal Reserve Bank of St. Louis *Review*, July/August 2007, *89*(4), pp. 305-26.

Greenspan, Alan. "Statement by Alan Greenspan, Chairman, Board of Governors of the Federal Reserve System, before the Committee on Banking, Housing, and Urban Affairs, May 27, 1994." *Federal Reserve Bulletin*, July 1994.

Gürkaynak, Refet S.; Sack, Brian and Wright, Jonathan H. "The U.S. Treasury Yield Curve: 1961 to the Present." Unpublished manuscript, Board of Governors of the Federal Reserve System, 2006.

Rudebusch, Glenn D.; Sack, Brian P. and Swanson, Eric T. "Macroeconomic Implications of Changes in the Term Premium." Federal Reserve Bank of St. Louis *Review*, July/August 2007, *89*(4), pp. 241-69.