Discussion

Brian Sack*

Introduction

It is a pleasure to comment on this paper by Professor Kuttner, which presents some intriguing evidence on a very important topic. The question at hand is whether a central bank can influence long-term interest rates through "nonconventional" policy measures that is, using tools other than setting of the short-term interest rate. This paper adds importantly to the growing set of evidence on this topic.

This topic is of broad interest to central banks around the world. Economic activity is importantly tied to long-term interest rates perhaps to a greater extent than it is tied to short-term interest rates. As a result, monetary policy operates on the economy largely through its influence on long-term interest rates and broader financial conditions. Understanding how a central bank's decisions affect the entire yield curve is therefore a key issue in the day-to-day conduct of monetary policy.

The issue may be particularly important, however, in a low inflation environment, because low inflation brings with it the challenges of the zero lower bound (ZLB) on nominal interest rates. If emerging circumstances were to cause the central bank to push the short-term rate to near zero, then the central bank would lose its ability to "ease" through conventional monetary policy. The question arises, then, whether the

* Vice President, Macroeconomic Advisers, LLC

central bank could still provide stimulus to the economy by lowering long-term interest rates. The answer to this question is critical for assessing the cost associated with the ZLB. In turn, it might affect a central bank's decision of just how low of an inflation rate it wants to target.

These issues are real-world challenges that confront major central banks around the world. Recent years have presented clear examples in both Japan and, to a lesser extent, the United States. The Bank of Japan, of course, kept its policy rate at or very close to zero percent for an extended period from 1999 to 2006. And the Federal Reserve had its own deflation scare in 2003, in which concerns about the zero lower bound clearly affected both monetary policy decisions and fixed-income market dynamics. Assuming that central banks maintain a low inflation environment, the challenges posed by the ZLB are likely to emerge repeatedly, and the question posed in Kuttner's paper will be a key consideration in how policymakers should react in those instances.

How Can Central Banks Affect Long-term Rates?

Kuttner discusses two ways in which a central bank can affect long-term interest rates without changing its current policy setting. The bank can use its communications to affect the future path of short-term interest rates expected by market participants what he refers to as "open mouth operations." Or the central bank, through outright purchases and sales of its portfolio, can affect the relative supply of outstanding securities. If the demand curve for a given debt security is downward-sloping, changes in the supply of that security will affect its price.[1)]

1) Kuttner considers a third approach, which is imposing explicit pegs or ceilings on bond yields. However, since the peg would presumably be enforced by the central bank standing ready to purchase securities if the yield were to rise to that level, this approach is simply an extreme version of operating through changes in the supply of outstanding securities.

The paper leaves aside open mouth operations, citing other research suggesting that central banks have some scope for managing policy expectations and operating through this channel. It instead focuses explicitly on the second possibility whether the central bank can have market effects by changing the supply of long-term debt instruments.

Kuttner is particularly interested in market effects that operate through the term premium, or the excess return that investors expect to earn (over the short-term risk-free interest rate) by holding the longer-term instrument. Implicitly, he assumes that the central bank is not using these changes in supply to send signals about the future path of short-term interest rates. In that case, the purchases have no effect on the return that an investor would expect from rolling over short-term debt securities. Thus, if the yield on the longer-term security is affected by central bank purchases, it must be because the excess return that investors expect to earn is reacting. This observation is the key link to the empirical approach that Kuttner takes, which I discuss below.

Before turning to the empirical evidence, Kuttner provides a very interesting discussion of historical episodes in which the supply of U.S. Treasury securities changed considerably. In doing so, he lumps together both changes in the Federal Reserve's portfolio strategy and changes in the Treasury's debt management strategy, arguing that both entities should be thought of together when considering this strategy. That seems reasonable what should really drive prices in fixed-income markets is the supply of securities available to the private sector, regardless of which government entity is determining that supply.

The historical tour in the paper is a rich description of various events and policies that at least had the potential to significantly influence the Treasury yield curve. It is highly recommended reading, although the historical record does not definitively answer the question at hand. Instead, the tour very much leaves the reader with a mixed feeling about whether these events did or did not affect long-term interest rates. Thus, we have to turn to the empirical evidence. I will concentrate my comments on that evidence, since it is the main contribution of the paper.

Kuttner's Empirical Evidence

To assess whether the supply of Treasury securities affects their yields, Kuttner uses an empirical approach first implemented by Fama and Bliss (1987) and later used by Cochrane and Piazzesi (2005). The idea is simple: Take the one-year holding return on an *n*-period zero-coupon bond, and measure it as an excess return over the one-year risk-free rate. The realized excess return will depend on whatever unpredictable events take place over that year, as those will cause capital gains or losses on the longer-term instrument. However, those unexpected events should be uncorrelated with information known at the beginning of the year. Any portion of the excess return that can be predicted with variables known ex-ante must instead represent the *expected* excess return, or the term premium.

In their baseline set of results, Cochrane and Piazzesi estimate a regression in which these one-year holding returns are explained by a set of forward rates from the beginning of the one-year period. They find that a tent-shaped pattern of these forward rates, which I will call the CP factor, has exceptional predictive power, accounting for 34% of the variation in excess returns for a five-year zero-coupon security over the sample from 1964 to 2003. This evidence suggests that the term premium varies over time.

Kuttner wants to investigate whether changes in the supply of Treasury securities result in variation in the term premium. To that end, he adopts the Cochrane-Piazzesi approach but includes on the right-hand side of the equation two variables associated with changes in the supply of Treasury securities. The first variable is the fraction of the Fed's portfolio (known as the System Open Market Account, or SOMA) that is held in longer-term securities, those with remaining maturities beyond five years. The second variable is the fraction of outstanding, privately-held Treasury debt that has more than five years to maturity.

Kuttner's primary finding is that changes in SOMA holdings have very significant effects on term premiums. Indeed, his estimated coefficients suggest that a $15 billion shift in SOMA to longer-maturity issues would push down the expected return on the five-year note by 2.2 percentage points (using the full sample estimates with the CP factor included) or by 7.0 percentage points (using the post-1972 sample).

The results should immediately raise some eyebrows. These are massive effects from what would be a fairly trivial operation by the Federal Reserve. Indeed, if this were the case, the Fed would have absolutely no difficulty setting longer-term yields to whatever level it wants. Of course, one reason the effects appear so large is that these are changes in one-year returns, not changes in the yield of the securities (a five-year return). However, even under the assumption that the excess returns are unaffected after the first year, the effects are still surprisingly large. The $15 billion portfolio shift in this case would have reduced the five-year yield by 44 or by 140 basis points according to the two estimates.

Effects of this magnitude seem puzzling. After all, markets absorb much larger shifts in supply all the time. The most extreme example comes from the mortgage-backed securities (MBS) market. Increases in interest rates add to the duration of outstanding MBS, because the likelihood that households will repay their mortgages declines. The market must absorb these changes in duration, just as it would a change in duration from a shift in the SOMA. But the flows in duration from the MBS market are immense. From the beginning of 2005 to the middle of 2006, for example, the duration created in the MBS market by the rise in interest rates was equivalent to the issuance of over $700 billion of ten-year notes.[2)] The market seemed to have little problem absorbing this change, which raises the question of why such trivial swings in SOMA holdings would instead have such sizable effects.

Why does Kuttner find such large effects from changes in the SOMA holdings? The problem may be with the Cochrane-Piazzesi methodology upon which his analysis is based, as discussed next.

---

2) Note that this is an "exogenous" shift in the supply of debt securities, one that does not contain any signals about future short-term rates. Thus, it is comparable to shifts in the SOMA.

Concerns about the Cochrane-Piazzesi Approach

Cochrane and Piazzesi (2005) take a very reasonable approach that yields very unreasonable results. This paper has been widely-cited and has had considerable influence both within academic circles and among financial market participants. Thus, it is important to look closely at their findings.

The baseline Cochrane-Piazzesi results considered here are those in which the excess return on a five-year zero-coupon Treasury security is regressed on the CP forward-rate factor. (Cochrane and Piazzesi present a much broader set of results in their paper, including some evidence on how to interpret their forward-rate factor in macroeconomic terms.) Cochrane and Piazzesi estimate this equation using the Fama-Bliss dataset on yields, which provides annual observations over the period from 1964 to 2003. Their estimated coefficient is shown in the first row of Table 1.

Table 1. Predicting the Excess Return on Treasury Securities (Cochrane-Piazzesi Exercise)

Maturity	Sample	Coefficient on C-P Factor	Predicted 1Y Excess Return (2006Q2)
5Y	1964-03	1.43*	-2.6 pp
5Y	1972-06	1.66(0.41)	-3.0 pp
10Y	1972-06	3.07(0.84)	-5.5 pp
15Y	1972-06	4.39(1.25)	-7.8 pp

* Coefficient reported by Cochrane and Piazzesi (2005). All other coefficients are estimated.

Note: Regression predicts the excess returns to holding a zero-coupon Treasury security over the subsequent year, based on information known at the beginning of that year. The independent variable is the forward-rate factor from Cochrane and Piazzesi (2005). Standard errors obtained under the Newey-West procedure are shown in parentheses.

To dig deeper into these findings, I re-estimate this equation using a new yield curve database that has become available from Gurkaynak, Sack, and Wright (2006). The analysis here uses quarterly data over a slightly shorter sample, from 1972 to 2006. The results are shown in the second row of Table 1.[3] As can be seen, the estimated coefficient is strongly significant, with a magnitude that is similar to (though slightly larger than) that reported by Cochrane and Piazzesi. Thus, I am able to largely replicate their results.

The far right column in the table shows the predicted excess return for the coming year obtained from the estimation results. These figures provide our first hint that the results appear somewhat suspicious. According to the regression, investors are giving up 3 percentage points of return this year for the benefit of holding the five-year note. That is, investors could earn about 5% by holding a one-year risk-free instrument, but they would rather assume the risk of holding the five-year note at an expected return of around 2%.

To be sure, there is nothing wrong with the idea of a negative risk premium. However, it is the magnitude of the negative premium that is surprising. And the puzzle intensifies if we look at longer maturities. An advantage of using the GSW dataset is that it includes longer maturities. Whereas the Fama-Bliss dataset only extends to five-year maturities, the GSS dataset reports zero-coupon yields out to the maturity of the longest Treasury security outstanding. Over the sample used here, that allows the analysis to be extended out to a 15 year maturity point.

The same regression is estimated for ten-year and fifteen-year zero-coupon Treasury securities, with the results reported in Table 1. For the ten-year security, the coefficient nearly doubles, with the results suggesting that investors are currently giving up 5.5 percentage points of return to hold that issue. For the fifteen-year security, the estimated coefficient is even bigger, and the regression predicts that investors today are giving up a whopping 7.8 percentage points of return. Thus, in

3) The standard errors are computed using the Newey-West technique. The error terms form the regressions would be expected to have serial correlation, given our use of overlapping quarterly observations. This is the same approach as used by Kuttner.

both cases investors are holding these securities with the expectation of not only earning a return well below the risk-free rate, but of actually taking outright losses over the next year.

Negative term premiums of this magnitude are simply implausible. This view is certainly supported by anecdotal evidence, as it is impossible to find any investor who is willingly holding the ten-year note today if he is expecting to earn a negative return over the next year. It is also supported by survey evidence, which does not show expectations of sharp capital losses on longer-term debt securities. The Blue Chip Economic Indicators survey released in July, for example, suggested that investors expected the ten-year yield to increase only 10 basis points over the subsequent year.

The recent predictions from the Cochrane-Piazzesi are not atypical, either. The exercise produces remarkable variation in the term premium throughout the sample. The expected excess return on the ten-year Treasury security varies over a range from -11 percentage points to +15 percentage points. And it changes rapidly from quarter to quarter, with a standard deviation of the quarterly change of 3 percentage points.

It is not clear *why* the Cochrane-Piazzesi approach produces such volatile results. Indeed, there is nothing conceptually wrong with the method that they employ. One possibility is that the predictive power just happens to be quite high in this sample, and that investors expected returns did not change as much as the regression suggests. In any case, perhaps we should think of the Cochrane and Piazzesi findings not so much as revealing an important characteristic of bond market pricing, but as revealing a puzzle in the behavior of the market over the past forty years.

But this, of course, becomes problematic when the method is used to answer an important empirical question of whether central bank purchases affect term premiums. I worry that Kuttner's finding of very large SOMA effects is capturing the same type of sensitivity that is apparent in the baseline Cochrane-Piazzesi results.

Another Reason to Be Cautious

Kuttner also finds that changes in the other supply variable, the one based on outstanding Treasury debt, have no effect on the excess returns of longer-term Treasury securities. This finding presents another important reason to be cautious about his results. As discussed above, Kuttner argues that the Fed and the Treasury should be grouped together when considering their policies that affect the supply of Treasury securities. Thus, it is worrisome that the empirical exercise produces sizeable effects from changes in the SOMA and nonexistent effects from the Treasury's debt management decisions.

This difference also raises a question about the SOMA variable used in the analysis. What should really matter for market pricing is the supply of longer-term securities relative to a fairly broad benchmarkthe size of all financial holdings of the private sector, or a measure of total wealth, or perhaps even the size of the economy. The SOMA variable is instead defined as the amount of long-term securities held by the Fed relative to *its own* total holdings. In is not clear that this is the right denominator. After all, if the SOMA shrunk to very small levels relative to the amount of outstanding Treasury debt, one would not expect a given percent change in its composition to have as large of an effect on the markets.

The Bottom Line

Given these arguments, I believe that we have to take Kuttner's findings with a grain of salt. His results are clearly interesting and provocative, and the paper is an important addition to the collection of evidence. But I continue to be skeptical of the magnitude of the effects that he finds. Admittedly, it is not exactly a compelling criticism to say that I find his method to be fine but just do not like his results. But, in the end, that is where I was left.

These concerns are supported by two observations noted above. The first is the volatility of the expected excess returns that arises in the baseline Cochrane-Piazzesi approach upon which Kuttner's exercise is based. The second is the apparent sensitivity of his results to the specific supply variable used in the exercise.

Where does this evidence, and the broader array of research that is available, leave us on the question of whether central banks can affect long-term interest rates? I would argue that the literature has not produced a clear answer in some areas. My own views are as follows.

First, central banks can surely affect long-term interest rates through the channel of the expected path of the short-term interest rate, at least if they are willing to commit themselves. Thus, it is not an issue of whether central banks *can* affect long rates, but whether they are *willing* to do so. We have considerable empirical evidence that Federal Reserve statements have considerable influence on longer-term Treasury yields (see, for example, Bernanke, Reinhart, and Sack (2004)). And those communications have, at best, provided only vague hints about policy prospects. Communications that were seen as explicit and credible promises about future policy actions would presumably have a much larger impact.[4)]

Second, it seems possible, though less definitively the case, that central banks can operate on long-term interest rates though the term premium. There are certainly episodes in which significant changes in the supply of Treasury securities seems to strongly influence the yields on those instruments. A primary example is the period that Kuttner discussed as the "end of the long bond." But here, the evidence is much less definitive, despite the efforts of Professor Kuttner and others. Hence, I would argue that the efficacy of this policy tool is still quite uncertain at this point.

---

4) One could argue that a central bank cannot credibly commit to future policy actions, but I am not convinced by that argument. Note that a central bank could always back up its commitments with current actions. Specifically, as part of its commitment device, it could do forward RP's at a zero interest rate, allowing private firms to go ahead and lock in a zero financing rate for some date in the future.

References

Bernanke, Ben S., Vincent R. Reinhart, and Brian P. Sack (2004), "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment," *Brookings Papers on Economic Activity* 2004:2, 1-78.

Cochrane, John H. and Monika Piazzesi (2005), "Bond Risk Premia," *American Economic Review* 95, 138-160.

Fama, Eugene F. and Robert R. Bliss (1987), "The Information in Long-Maturity Forward Rates," *American Economic Review* 77, 680-692.

Gurkaynak, Refet S., Brian Sack, and Jonathan H. Wright (2006), "The U.S. Treasury Yield Curve: 1961 to the Present," Finance and Economics Discussion Series working paper no. 2006-28, Federal Reserve Board of Governors.