

Discussion

Nicola Anderson

The Bank of England, like many other central banks, has long been interested in extracting useful information from the bond markets. One of Fung and Remolona's major contributions is to bring us one step closer to being able to do this in a manageable and meaningful way. I would like to start by summarizing what are, in my view, the three highlights of the paper.

First, they use the Kalman Filter to estimate the two-factor model of the term structure. This seems to me to be a very neat way of solving a rather complex problem.

The only potential drawback I can see is that, by fitting the model to estimated yields rather than directly to bond prices, the Kalman Filter estimates may be exposed to some sort of measurement error. But this will depend on the choice of technique used to estimate the yields. For example, yields estimated using the Nelson and Siegel approach may be biased as a result of restrictions imposed on the functional form of the yield curve. I am not so sure about the Canadian data, but since the U.S. yields are estimated using cubic splines, I would expect that the chances of these being similarly biased are very small.

Second, I would like to highlight the way in which the term structures of Canada and the United States are estimated jointly in the paper. It seems to me that there are many potential benefits from exploring this approach. The most obvious, as the paper implies, is to assess the extent to which countries have a common source of risk—either monetary or real—and, conversely, the extent to which monetary policies are perceived by the market to diverge. Applying this to the United Kingdom, for example, might

be particularly useful as the decision as to when to join the European Monetary Union draws closer.

But more generally—and perhaps on a more ambitious note—I would be interested to see whether this joint-estimation approach could be extended to tell us something about exchange-rate risk premiums.

Third, the model presented in the paper offers a very tractable approach to providing a full decomposition of the nominal yield curve. This must be of interest to any policy-maker who wants to find out what the market thinks about the prospects for future real growth and inflation, as well as how the market views the monetary stance. And it is this area that I would like to discuss in further detail.

As many of you know, in the United Kingdom we are luckier than most in being able to dismantle the nominal term structure. We have a relatively well-established market in index-linked bonds. So, after overcoming a few problems arising from the indexation lag, we can split nominal yields into real yields on one hand and into an inflation component—which we refer to as the inflation term structure—on the other.

Ideally, however, to extract the kind of information I mentioned before, we would like to go one step further. That is, we would like to decompose real yields into real rate expectations and risk premiums and similarly the inflation term structure into inflation expectations and inflation-risk premiums.

Fung and Remolona's paper demonstrates how we can achieve this decomposition in a single step, using only the nominal term structure. This does not mean, however, that information from the index-linked market is then redundant. On the contrary, I imagine that estimated yields on index-linked gilts could prove to be a very useful source of information, not only for estimating the model, but also for testing the specification of the underlying factors—in particular, the real component. And, in fact, I believe that Remolona has already done some work with Frank Gong and Mike Wickens on applying his model to the U.K. market.

What I want to concentrate on today, however, is the inflation-risk premium. As the authors point out, the source of risk they are pricing in the model has to do with revisions in expectations. What I want to focus on is what the model does not tell us about the risks of inflation—and why—and hence one way in which I believe it could be usefully extended. To do this, I would like to examine the relationship between the inflation-risk premium and the cost of government borrowing.

One of the main reasons commonly cited for issuing index-linked bonds is to reduce the cost of government borrowing; the idea being that, to invest in nominal debt, investors demand an inflation-risk premium. And

this raises the expected real cost of borrowing from issuing nominal bonds above the equivalent cost of issuing index-linked debt.

In terms of the returns to the investor, I can therefore define a measure of the inflation-risk premium as the difference between the expected real return on the nominal bond, denoted here by a \$ sign, and the expected real return on a real bond.

Now I will show how this definition relates to the model in the paper. For this purpose, I need only concentrate on two securities: a nominal bond with one period to maturity, and a real or perfectly indexed bond, also with one period to maturity.

We know that in real terms holding the real bond until its maturity is a riskless strategy. The real payoff at the end of the period is known with certainty to be equal to 1 unit of the consumption bundle. So the real return is guaranteed to equal the inverse of the current price:

$$R_{1t} = \frac{1}{P_{1t}}.$$

Conversely, while the nominal bond is riskless in nominal terms—offering a guaranteed \$1, say, at maturity—we know that it is risky in real terms because of inflation. In this case, denoting the rate of inflation by I , the real payoff on the bond is $1/I$, and the expected real return is the expected real payoff divided by the price. Thus,

$$E_t[R_{1t}^{\$}] = \frac{E_t[1/I_{t+1}]}{P_{1t}^{\$}}.$$

So, I will consider first the pricing of the nominal bond. In the paper, the nominal bond is priced using a nominal stochastic discount factor. Denoted by M , this values a nominal payment at the end of the period across different states of the world. Since in each state the nominal payoff on the nominal bond is always equal to \$1, the price of the bond is just given by the expected value of M :

$$P_{1t}^{\$} = E_t[M_{t+1}] = E_t[M_{t+1}^*/I_{t+1}].$$

An alternative approach would be to value the nominal bond in *real* terms. In this case, we denote the *real* stochastic discount factor by M^* ; this is used to discount the *real* payoff on the nominal bond, which we know is equal to $1/I$. So, taking the expected value of the discounted real payoff, we get the term on the far right of the equation.

Similarly, we can price the real bond as the expected discounted value of its real payoff. Since on the real bond this is just equal to 1 unit, the price is then just equal to the expected value of M^* :

$$P_{1t} = E_t[M_{t+1}^*].$$

Finally, we can derive an expression for the expected real returns on the two types of bond in terms of the real stochastic discount factor. To be consistent with the model in the paper, however, we do this in logs. Then the real return on the real bond is given by r , where m^* is the log of the real discount factor:

$$r_{1t} = -E_t[m_{t+1}^*] - \frac{1}{2}V_t[m_{t+1}^*].$$

The expected return on the nominal bond is:

$$E_t[r_{1t}^*] + \frac{1}{2}V_t[r_{1t}^{\$}] = -E_t[m_{t+1}^*] - \frac{1}{2}V_t[m_{t+1}^*] + CV_t[m_{t+1}^*, \pi_{t+1}],$$

where π is the log of the inflation rate.

Comparing the two returns, notice first that the second-order term on the left-hand side of the second equation arises only because we are working in logs. The inflation-risk premium, as defined by the difference in the real cost of borrowing between the two bonds, is therefore given by the covariance of the real stochastic discount factor with inflation.

Intuitively, this makes sense; the real stochastic discount factor is a measure of how investors value real payments in the future. If this is unexpectedly high at the same time as inflation, this means that the return on the nominal bond, which varies negatively with inflation, will be low precisely when the investor most values it. In other words, the nominal bond will increase the overall risk of future real wealth, for which investors will demand compensation by way of the inflation-risk premium.

The question is: How does this definition relate to the model presented in the paper? To see this, notice first that the log of the nominal stochastic discount factor is equal to the difference between the real stochastic discount factor and the log rate of inflation:

$$m_{t+1} = m_{t+1}^* - \pi_{t+1}.$$

So, in terms of the underlying factors, if log inflation is represented by terms in x_1 , the real stochastic discount function is represented by terms

in the real component, x_2 . The perceived processes for these two factors can therefore be defined as shown here:

$$-m_{t+1}^* = x_{2t} + \lambda_2 x_{2t}^{\frac{1}{2}} u_{2,t+1},$$

$$\pi_{t+1} = x_{1t} + \lambda_1 x_{1t}^{\frac{1}{2}} u_{1,t+1}.$$

The covariance between m^* and the rate of inflation is directly proportional to the covariance between the shocks u_1 and u_2 . But in the paper, this is assumed to equal zero. And so, the inflation-risk premium, as I have defined it here, is also assumed by the model to equal zero.

So what does all this mean in practice? If there *is* an inflation-risk premium on nominal bonds, the model presented in the paper may well be mis-specified. On the other hand, if the cost of borrowing on the two types of bonds is equal, there is no problem and everyone is happy—except for the growing collection of governments, including those of both Canada and the United States, who are issuing index-linked bonds in the belief that it is saving them money.

Either way, it would be useful to extend the model to allow for a non-zero inflation-risk premium so that we can let the data decide its value. Of course, if this risk premium did turn out to be non-zero, this would also provide us with valuable information about how the market expects shocks to inflation would affect real activity, to the extent that this is reflected by the real stochastic discount factor m^* .

To conclude, I would say that this is a very interesting and innovative paper that offers a useful framework within which to estimate a decomposition of the nominal term structure of interest rates.

One potential drawback is the assumption that the real stochastic discount factor and inflation are uncorrelated. Extending the model to include a non-zero correlation, however, may not be easy. Such a model may no longer belong to the affine class of models, models that we know are very easy to use.

However, given the large payoffs that policy-makers could reap from being able to extract this type of information from the yield curve, I can only encourage the authors to put such an extension of their model near the top of their research agendas, and so capitalize on the substantial progress they have already made.