Discussion

Alan White

The Gravelle-Muller-Stréliski paper addresses the question, To what extent is the forward rate an unbiased predictor of the realized spot rate?

The authors fit the observed behaviour of the spot and forward rates with a vector error-correction model (VECM). This captures all the relevant structure in the data. The error-correction term is interpreted as a time-varying risk premium. The realized change in interest rates is then related to the spread between the forward and spot rates, a constant risk premium, and the estimate of the time-varying risk premium. The constant portion of the risk premium thus estimated is about 100 basis points at the 9-month point. That is, the realized rate is on average about 100 basis points below the forward rate.

What is a plausible level for the risk premium in interest rates? Normally, risk premiums are measured in an asset-pricing framework by observing investments' realized risk characteristics. For instance, we could ask, What is the standard deviation of the returns earned by a 1-year investment in bonds of different maturities? This could then be related to the observed standard deviation of returns for other types of investments to determine the relative riskiness of debt investments.

For example, consider an n-month investment in an m-month t-bill. Let $P_{t,j}$ be the price at time t of a j-month t-bill, and let us assume that the expectations hypothesis holds. In this case, the expected future price of a t-bill is its forward price, $E_t[P_{t+j,k}] = P_{t,k+j} / P_{t,k}$. The realized return on the n-month investment in the m-month t-bill is $P_{t+n,m-n} / P_{t,m} - 1$ and the expected return is the yield on the current n-month t-bill. Using these

measures, we can assess the riskiness of this investment and estimate a risk premium. If the estimated annualized risk premium on this investment is s per year, it can be shown that the expected future spot rate is expected to be about $n\,s\,/\,(m-n)$ below the forward rate.

Historical data for the last decade show that the estimated annualized standard deviation of the return from a variety of investments in Government of Canada t-bills has been:

Life of investment, *n* (months)	2	3	6	
Life of t-bill, *m* (months)	3	6	12	TSE 300
Annual standard deviation of return (per cent)	0.15	0.52	1.03	15.17
Relative risk	0.00992	0.03419	0.06811	1
TSE risk premium (per cent)	5	5	5	
Interest rate risk premium (per cent)	0.05	0.17	0.34	
Forward premium (per cent)	0.10	0.17	0.34	

Presuming that the risk premium is proportional to the standard deviation of return (a common result in asset-pricing models) and that the risk premium on equities is 5 per cent, an ad hoc estimate of the term premium arising from risk is about 34 basis points at the 6-month point on the term structure. A more elaborate asset-pricing model might find different results. If we consider the correlation between returns on t-bill investments and returns on equity investments, we might find lower forward-risk premiums.

It is interesting to speculate on why the risk premium estimated on the basis of the average observed forward-risk premium is so much higher than that based on an admittedly ad hoc asset-pricing model. One possibility is that the asset-pricing model is incorrect. A second possibility is that the assumption that expectations are on average realized is flawed. This is a very common assumption in econometric work. It appears to be based on a stationary model of the world, in which the players learn the statistical properties over time.

The assumption that expectations are, on average, realized is a very strong one. For instance, let us explore how many trials it takes to accurately estimate the expected outcome of a binomial process such as a coin toss. The following chart shows the estimated mean and standard deviation of the process for 1,000 trials. A head is given a value of one and a tail is given a value of 0. The estimated properties of the process are based on all the trials

Figure 1

Number of Trials to Accurately Estimate Outcome of a Binomial Process

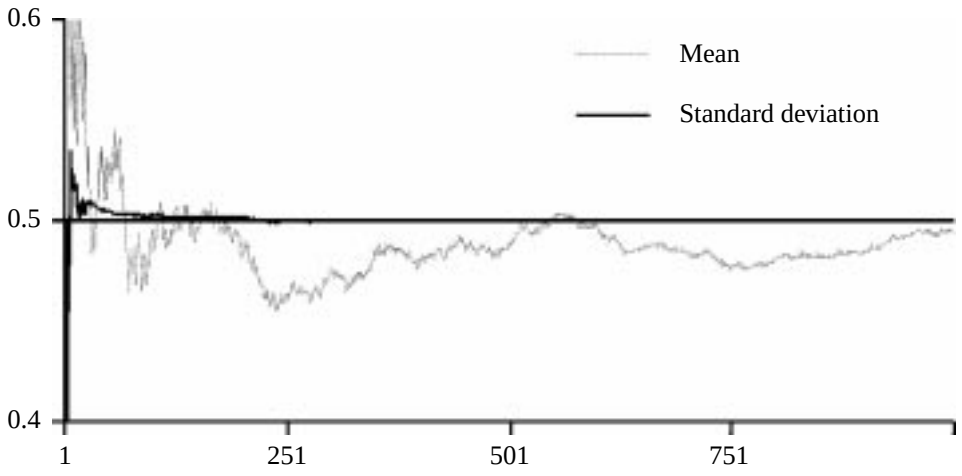

to date. The standard deviation is well known within 250 trials, but even after 1,000 trials the mean is not correctly estimated.[1]

This chart reveals that it is not easy to "learn" the statistical properties of a simple stationary stochastic process. If the process is not stationary, it may be impossible to ever learn its statistical properties, unless we can determine why it is not stationary. How easy is it, then, to learn the properties of an economy to the point that we can safely say that expectations are on average realized? If we cannot, with reasonable confidence, say that expectations are on average realized, what is the meaning of the observed econometric results?

One other possible source of difference in the estimated risk premium is the econometric techniques used. The procedure uses the VECM to extract the time-varying risk premium, and then uses this estimate in a new estimate of the expected future spot rates. This two-stage process introduces an errors-in-variables problem, and can compromise the quality of the ultimate statistical tests. It would be interesting to see how well the procedure works on simulated data. Several sets of simulated data could be constructed to match the various views of the world that the empirical tests are trying to distinguish. For example, one set of data would incorporate a

1. The standard deviation of the estimate of the mean is $\sigma/\sqrt{n}$, of course. When $\sigma = 0.5$ and $n = 1{,}000$, this is 0.0158. The width of the 95-per-cent confidence interval for our estimate of the mean after 1,000 trials is then about 0.062.

time-varying risk premium, one would not. It would be comforting to be able to show that the econometric techniques were able to extract the time-varying risk premium that was built into the data, and failed to extract the time-varying risk premium that was not present.

Notwithstanding the concerns I have raised about the interpretation of the observed term premium (the forward rate less the realized spot rate), the model the authors present may have considerable value as a forecasting tool. To verify this, some out-of-sample testing should be done. Econometric models that have many degrees of freedom (such as the VECM in this paper) by definition fit any sample of data better than models with few degrees of freedom. This does not make them better models. It may be that there is substantial "over-fitting" going on, and that the predictive power of the model for out-of-sample data is poor.