

Discussion

Arturo Estrella

Atta-Mensah and Yuan construct a connection between two concepts that are closely related, but have very different properties. Inflation expectations are intuitively clear and easily interpretable, but they are difficult to observe in practice. In contrast, forward interest rates are conceptually more complicated, especially if time-varying risk premiums of various types are allowed, but they are easy to observe in that they can be simply calculated from the term structure of current interest rates. The main point of the Atta-Mensah and Yuan paper is that if one could find a simple way to “translate” forward rates into inflation expectations, one could produce a construct that combines the attractive features of the two concepts: The construct would be easy to observe and calculate, and it (or its simple transform) would be easily interpreted and useful for policy analysis.

The authors use various analytical techniques:

- a theoretical consumption-based capital-asset-pricing (C-CAPM) model;
- statistical time-series modelling techniques (vector autoregression [VAR], autoregressive conditional heteroscedasticity [ARCH], and error correction);
- financial modelling (the Svensson model of forward rates);
- historical episodes (inflation-targeting announcements); and
- empirical evidence of various sorts.

They then ask a useful question from the point of view of economic modelling and policy analysis: How can we go from observed forward rates to unobserved inflation expectations?

The paper proceeds to use the technical results to examine three case studies associated with inflation targeting in Canada:

- the announced intentions to stress price stability in monetary policy (January 1988);
- the initial announcement of inflation targets (February 1991); and
- the announcement of the extension of the initial inflation targets (December 1993).

The results obtained from all of these case studies are quite plausible. First, the direction of the effects is as expected; the new information about low inflation targets leads to a decline in inflation expectations at almost every point along the maturity curve. Second, the exercise provides some reasonable magnitudes for the effects, which help quantify the theoretically expected results.

All in all, the authors deserve credit for developing and applying a procedure that leads to some useful policy-related results. However, a discussant is compelled to be critical, so I must ask, What is there in this paper that one may criticize? To find something about which to be critical, we have to delve more deeply into the details of the modelling, which are rich and complex and probably more complicated than they need to be to accomplish the purposes of the paper. Let us turn thus to the individual links in the paper's chain of reasoning, and to the possible weaknesses of those links.

One complicated feature of the paper is the use of five independent models. The following inventory lists the models in order of appearance in the paper:

1. The C-CAPM model used is based on intertemporal maximization of a utility function of the form

$$\sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\gamma}}{1-\gamma}.$$

2. VAR-ARCH. This is a bivariate VAR(1) with inflation and the log-difference of consumption. The errors are assumed to follow an ARCH(1) process.
3. A forward-rate rule in which expected inflation is modelled as the fitted value from a simple linear regression of expected inflation on the forward rate:

$$\pi_t = a + b f_t.$$

4. A Svensson forward-rate function. This is an extension of the Nelson–Siegel function with added flexibility. It is of the form

$$f_t(m, b),$$

where m is the settlement date and b is a vector of parameters.

5. A vector error-correction model (VECM), which has as endogenous variables M1, CPI, real GDP, and the 90-day commercial paper rate. The authors provide little detail in the paper, relying instead on references to their other work.

With this inventory in hand, we can then examine how these models fit in the paper’s overall argument. The structure of this argument can be outlined in these seven steps.

Step 1. Use the C-CAPM and the VAR-ARCH to express the variables of interest (forward rates, expected nominal and real rates, expected inflation, as well as forward and inflation premiums) as functions of the parameters β and γ of the utility function.

Step 2. Use the expression for the forward rate from step 1 plus forward rates derived from the Svensson forward-rate function to estimate β and γ .

Step 3. Use the results of steps 1 and 2 to argue that the forward and inflation premiums are small (≤ 2 bps) and may be ignored.

Step 4. With the conclusion from step 3 in hand, discard everything derived so far!

Step 5. Use the VECM to construct another estimate of expected inflation (and the breakdown of the forward rate into an expected real interest rate and expected inflation).

Step 6. Regress expected inflation from step 5 on the Svensson forward rates to estimate the forward-rate rule.

Step 7. Use the forward-rate rule to estimate the change in expected inflation in the three case studies.

At this level of detail, we can now identify three specific problematic features. The first issue is related to the paper’s modelling strategy and to the surprising decision associated with Step 4 above. More precisely, the paper shows at least four different ways to estimate expected inflation. There are other variables for which multiple methods are also presented, but expected inflation is the most important, given the key role it plays in the paper. The four estimates of expected inflation are given by:

- VAR-ARCH
- VAR-ARCH + C-CAPM + Svensson

- VECM
- VECM + Svensson + forward-rate rule.

Of these four alternatives, only results for the last are actually presented and used in the paper. The others are either used in intermediate steps in the calculations or are unused by-products of a given model.

Several questions arise as a result of this type of redundancy. For instance, are the different estimates empirically similar or do they differ in visible ways? What is the role of the VECM? Is it truly needed, or could the its role have been played by the joint model that was discarded in Step 4? None of this is made clear.

The second problematic feature of the paper is the extreme result regarding the size of the risk premiums (both forward and inflation). Are Canadian data so different from U.S. and U.K. data? If so, are there any intuitive explanations as to why they differ so dramatically? Another possible explanation may lie in the C-CAPM model, which is known to have spawned several well-known problems, such as:

- the equity premium puzzle—the estimates of the coefficient of relative risk aversion obtained using the C-CAPM in a model with equity are far too large to be intuitively plausible;
- the risk-free rate puzzle—the estimates of the rate of time preference obtained using the C-CAPM are negative or too large;
- the stock market volatility puzzle—stock prices are more volatile than consumption volatility implies in conjunction with the C-CAPM.

It would not be surprising if the “Canadian risk premium puzzle” were traceable to the C-CAPM.

A third problematic feature of the paper is the high volatility of expected real rates—up to 50 per cent higher than the volatility of expected inflation, in contrast with the results for the United States and the United Kingdom. It might not be surprising to find ex post real rates to be highly volatile, but the ex ante volatility seems to have unintuitive implications about the high-frequency components of monetary policy reactions. Since this result is obtained by using the VECM and not the C-CAPM, this particular puzzle cannot be blamed on the latter model.

In conclusion, this paper provides useful analysis of an interesting policy-related question. The results are intuitively plausible and deserve careful consideration. However, when examined in detail, questions about the paper’s modelling strategy cast some doubt on the accuracy of the quantitative results, particularly in the context of the case studies. The qualitative patterns are most likely right, but the quantitative conclusions

should be read with caution until further evidence appears that would support the less-plausible empirical estimates.