Discussion

Angelo Melino

Introduction

It is a matter of arithmetic to show that, in the absence of arbitrage, an asset's price can be written as the discounted expected value of its payoff, where the expectation is constructed using what is called the *risk-neutral* distribution. It is also a matter of arithmetic to express the (not necessarily unique) risk-neutral probabilities of an event as a mix of the market's assessment of the event's probability and an adjustment for bearing "risk." It is tempting to use the risk-neutral probabilities and a model for pricing risk to uncover the market's expectations of certain events. But it is not straightforward.

Variations on an Identity

In order to understand the link between expected inflation and forward rates, let us start with the following identity [see equations (5) and (6) in the authors' paper]:

$$\begin{aligned} E_t\pi_{t,k,q} &= f_{t,k,q} - E_t r_{t+k,\,q} - [f_{t,k,q} - E_t R_{t+k,\,q}] \\ &\quad - [E_t R_{t+k,\,q} - E_t r_{t+k,\,q} - E_t \pi_{t+k,\,q}] \\ &= f_{t,k,q} - E_t r_{t+k,\,q} - \phi^f_{t,k,q} - \phi^\pi_{t,k,q}. \qquad (1) \end{aligned}$$

I want to emphasize that, because it is an identity, equation (1) can be derived without invoking conditional normality of consumption growth and inflation. In fact, it holds regardless of how we choose to measure (or define) consumption $\{c_{t+k}\}$, the rate of inflation $\{\pi_{t+k,q}\}$, the nominal term structure $\{R_{t+k,q}\}$, the real term structure $\{r_{t+k,q}\}$, or the term structure of forward rates $\{f_{t,k,q}\}$. Finally, the identity (1) also holds for any interpretation of the expectations operator E_t. Given all the choices, which makes the most sense?

It depends. If our objective is to generate good forecasts of inflation to help the Bank of Canada achieve its inflation targets, then it makes sense to take the left-hand side of equation (1) as the rational expectation of

$$\ln(P_{t+k+q}/P_{t+k}),$$

where P_{t+k} is the level of the CPI at time $t+k$. But even for the Bank, other objects are also interesting.

We may want to uncover the *public's* expectation of future inflation. Although I find the assumption of rational expectations compelling as a first approximation, we may want to at least entertain the idea that the public is making systematic errors in their forecasts of inflation. Perhaps they need time to be convinced that the Bank's commitment to fighting inflation really does represent a new regime. At least in the transition, the public's expectations may systematically overpredict inflation while the Bank generates credibility.

We may be interested directly in the expectations formed using the *risk-neutral* distribution. The risk-neutral distribution contains exactly the same information as the asset prices that are used to construct it—no more and no less—but sometimes just rearranging information makes it easier to interpret (witness the numerous ways that are used to summarize the term structure of interest rates).

Finally, we may be interested in forecasting the *Bank of Canada's* expectation of inflation. The Bank goes through various cycles in assessing the prospects of the Canadian economy, but the shortest major cycle is quarterly. Between these quarterly cycles, the Bank needs to monitor and update its beliefs about the path of the economy. In this paper, the authors take the left-hand side of equation (1) to be the quarterly forecasts of inflation that come out of one of the Bank's models. They look to see whether they can use the information in forward rates to "predict their quarterly prediction" at quarterly frequencies. This is just a first step. Their objective, I infer, is to eventually use available asset price data to predict their quarterly prediction at weekly or even daily frequencies.

Variations on the Forward-Rate Rule

I will assume that our objective is to uncover the rational expectation of subsequent inflation. The authors focus on the following forward-rate rule:

$$\hat{\pi}^{e}_{t,k,q} = a_k + b_k f_{t,k,q} \tag{2}$$

where $\hat{\pi}^{e}_{t,k,q}$ is the estimate of

$$E_t \hat{\pi}^{e}_{t+k,\,q}.$$

The value of the coefficients a_k and b_k are easy to derive from the identity (1) (just add

$$\hat{\pi}^{e}_{t,k,q} - E_t \pi_{t+k,\,q}$$

to both sides of the equation) using the familiar Theil–Griliches specification-error analysis. The paper does so under the maintained hypotheses that the forward and inflation premiums

$$\phi^{f}_{t,k,q} \text{ and } \phi^{\pi}_{t,k,q} \text{ (respectively),}$$

along with

$$\hat{\pi}^{e}_{t,k,q} - E_t \pi_{t+k,\,q},$$

are identically 0. The values of these coefficients are of indirect interest. The main questions are the following:

(i) Is there useful information about future inflation in forward rates?

(ii) Can we get a better forecast of inflation than (2)?

The paper shows that forward rates are correlated with inflation forecasts. This is a weak requirement. Let us raise the bar a bit and ask a more demanding question: Are movements in forward rates useful for *revising* our forecasts of inflation? This suggests that we entertain something along the lines of the following specification[1]:

$$\hat{\pi}^{e}_{t,k,q} = a_k + b_k f_{t,k,q} + c_k \hat{\pi}^{e}_{t-1,k+1,q}. \tag{2$'$}$$

Investigating specifications along the lines of (2′) seems especially appropriate given one of the paper's main motivations: To provide high-frequency updates to policy-makers that can supplement the monthly and quarterly models they now rely on. Taking this motivation to its logical

1. I ignore issues involving unit roots.

conclusion, we would like to find *time paths* for the coefficients that tell us how to update the forecasts between, say, the quarterly cycles used to generate $\hat{\pi}$. My prior, based on casual observation and the fact that inflation contracts are not actively traded,[2] is that there is very little information that arrives on a daily, or even weekly basis, about the path of inflation so that the coefficient b_k in (2′) should be close to 0. Talk is cheap and a better discussant would have run the regressions. I will leave that to the authors.

Under the maintained hypotheses of the paper, the only way to improve on equation (2) is to obtain independent information about expected real interest rates. Two sources seem obvious, albeit problematic. For about a decade now, the Government of Canada has issued Real Return Bonds. The market is thin and the returns on these bonds reflect both expectations of future real rates and an inflation premium, so they are not perfect. But in this context, they do not have to be. The only question is whether or not there is information in the real return term structure that can be used to improve on the forecast that is implicit in (2), namely

$$r^e_{t,k,q} \equiv f_{t,k,q} - \hat{\pi}^e_{t,k,q}.$$

A second source of information about the expected path of Canadian real interest rates is the expected path of U.S. real interest rates. Again, the link is not perfect, but it does not have to be. I think it would be relatively straightforward and interesting to add U.S. forward rates to (2). Does including either (or both) of these signals of real interest rate expectations lead to a better forecast of inflation? This is purely an empirical question. I will leave it to the authors.

Finally, let us turn to the maintained hypotheses of the paper. I am not going to say much about the difference between the rational expectations and the forecast of subsequent inflation obtained using the quarterly vector error-correction model described in the authors' paper. But I have published papers attacking the expectations hypothesis, so I have to comment on this notion that forward (and inflation) premiums are small and economically not meaningful. To be fair, the authors do not maintain that premiums are 0 but devote a good deal of their paper to convince us that these premiums *turn out* to be very small. They invoke the consumption-based capital-asset pricing (C-CAPM) model with isoelastic preferences and an assumption that consumption and inflation have a joint lognormal conditional distribution to

---

2. I vaguely remember, but have not been able to confirm, that contracts on the CPI were traded briefly in Chicago. The contracts failed to generate enough interest and the market was abandoned. One of the reasons given for the failure of derivatives on the CPI to generate sufficient trading volume was the paucity of high-frequency information on the path of the price level.

solve for the equilibrium forward and inflation premiums. According to their calculations, the equilibrium premiums are probably one or two (and at most a dozen) basis points, so they can be ignored.

I am not convinced. This particular version of the C-CAPM is notoriously misleading. For instance, if we use this model to price equities, we would conclude that the expected real return for holding equities is only about 20 to 40 basis points, much less than the historical average of approximately 700 basis points. If we use this model to estimate the social cost of business cycles, we would conclude that the average citizen would be willing to pay no more than a few pennies to eliminate the cycle entirely! With such a track record, how can we take seriously this model's prediction that forward and inflation premiums should be very small? At the very least, the authors should experiment with a variant of the C-CAPM that is not so obviously in violation of the data. For example, Campbell and Cochrane (1995) have reverse-engineered a model of habit persistence that rationalizes a wide variety of empirical regularities about asset prices. I would be interested in seeing what sort of premiums the Campbell and Cochrane model generates.

However, a more direct and, in my mind, more convincing argument that forward premiums are not negligible is to invoke rational expectations and document their predictability. Many authors have done so and find that forward premiums are large and predictable. This finding is so well established and so robust that there is little point in repeating the results. I will simply give an example. For the case $k = 1$ and $q = 4$, I used the authors' sample period and found that 40 per cent of the variation in the realized value of $\phi^{f}_{t,k,q}$ was predictable. The ratio of the standard deviation of the predicted forward premium to that of quarterly changes in inflation was also about 40 per cent. These are not small numbers and should disabuse everyone of the notion that forward premiums are 0.

Do time-varying premiums affect the validity of the forward-rate rule given by (2)? No. The interpretation of the coefficients changes, as does the decomposition of the yield curve into information about expected inflation, future real interest rates, and premiums. But it is still a perfectly valid exercise to use forward rates to predict subsequent inflation. What does change, however, is that there is now a new possibility to do better. If we can find some independent measure of the premiums, we can add them to the right-hand side of (2) to get a cleaner measure of inflation expectations. Research by my colleague Walid Hejazi, confirmed by the paper presented earlier at this conference by Gravelle et al., suggests that forward premiums. can be proxied by a measure of the conditional variance of innovations to the yield curve. I know less about the behaviour of inflation premiums, and would like to see more research on this topic, but I also doubt that they are small and unimportant. In short, the world would be a simpler place if forward and inflation premiums were 0. But it ain't so. Ignoring these premiums means ignoring the possibility of improving upon the forward rate as a predictor of subsequent inflation.

References

Campbell, J. and J. Cochrane. 1995. *By Force of Habit: A Consumption-based Explanation of Aggregate Stock Market Behaviour.* National Bureau of Economic Research Working Paper 4995.