

Discussion

Glen Donaldson

The paper by Levin, Mc Manus, and Watt addresses a very interesting question: “How can the Bank of Canada ascertain what financial market participants believe about the probable future behaviour of the Canada–U.S. exchange rate, given that the Bank cannot directly ask everyone in the market what they think?”

This is a particularly important question given that the Bank (and everyone else) would like to know how markets view the Bank’s monetary policy actions and pronouncements.

An Analogy

A simple analogy might be to consider the problem someone might face if they attempted to determine my body weight. One way to ascertain how much I weigh is to directly ask me: “How much do you weigh?” This would be akin to asking everyone in the financial market what they think about the potential future behaviour of the Canadian dollar in light of the Bank of Canada’s recent policy announcement. Unfortunately, in reality it may be difficult to obtain such information in this direct manner, and thus an indirect method must be employed.

One possible indirect way to guess someone’s body weight might be to observe what the person eats for lunch and then deduce how much they weigh based on how much they ate. To do this we could begin by hypothesizing that heavier people eat more food than light people eat. Then if we were told that the only thing someone ate for lunch was a small side salad, we might guess the person weighs 90 pounds. Conversely, if we were

told the person ate three steak dinners in one sitting, we might guess they weigh 500 pounds.

In this example, we are using our hypothesized relationship between food consumption and body weight to deduce how much a person weighs from what we observe about their food consumption. There are, of course, several problems with this approach. First, our hypothesis could be wrong—maybe heavier people do not always eat more than light people eat. For example, a heavy person might eat a small meal because he or she is on a diet, while a light person might eat a large meal in an effort to increase their weight. In this case, our deduction technique would deliver a conclusion exactly opposite to reality. This problem is usually referred to as “model specification error.”

Second, we may have the correct model, but may not observe all the relevant data. For example, we may observe what the person eats for his or her appetizer, but not observe the main course. This would not be a problem if food consumption in the main course was perfectly correlated with appetizer consumption, but could be a problem otherwise. For example, a small person may eat only an appetizer and no main course, while a large person may eat a big main course in addition to an appetizer. Our lack of data concerning main course consumption would therefore inhibit the ability of our model to deliver accurate body weight estimates. Such a problem might be called “sampling error.”

A third problem, often referred to as “measurement error,” would arise if we could only visually observe what someone ate without being able to accurately measure calorie content. A T-bone steak may look big, for example, but have low calories since the item may be mostly bone, while a small sirloin may actually have more calories.

I could continue the analogy further to discuss a variety of other possible errors and biases, but I hope the point is obvious by now—all indirect deduction methods are subject to many possible sources of error because they are dependent on the model and assumptions employed. This does not mean that one should become paralyzed at the wide range of model and assumption choices and thus do nothing. However, it does mean that we need to proceed cognizant of the fact that the quality of the results from any analysis depends on the quality of our model and assumptions, and thus that our conclusions should be viewed with appropriate caution. Given the many potential sources of specification, data, and other errors in the procedure employed by Levin, Mc Manus, and Watt, such caution is particularly warranted in this case. That is why I thought the preceding analogy might be instructive.

The Procedure

Let us now consider the indirect deduction methodology employed by Levin, Mc Manus, and Watt, which I shall list in point form along with a note indicating possible sources of error at each step:

1. Hypothesize a mathematical model to describe the way exchange rates evolve through time. Source of possible error: Wrong model for exchange rate evolution.
2. Derive some mathematical formulas to tell us how an exchange rate futures option would be priced if exchange rates really did behave the way the authors have hypothesized. Source of possible error: Markets may not satisfy all necessary conditions for the options-pricing equations to be true (e.g., no market frictions.)
3. Collect data on the prices of exchange rate futures and futures options contracts. Source of possible error: Only exchange-traded options data are employed; no over-the-counter contracts data were used.
4. Taking the options-price data as given, use the mathematical models to reverse-engineer what the financial market must have believed about the probable behaviour of future exchange rates such that the options would have had the prices we observed in the market. In other words, find the set of probability beliefs such that, when the belief parameters are plugged into the mathematical pricing model, the model fits the data. Sources of possible error: Different criteria for fitting models to data; the theory requires a continuum of strike prices for the options, but we have only a few discrete price observations, so interpolation is required.
5. Assume that the only force moving exchange rates during the time period under study is the market's anticipation, and then observation, of the Bank's *Monetary Policy Report*. Source of possible error: There may be other factors moving exchange rates and options prices.
6. Indirectly deduce the effect of the *Monetary Policy Report* on the market's beliefs by observing changes in the estimated probability density functions (PDFs) derived from the models and data conditional on all of the previous assumptions being true. Source of possible error: Model and assumptions may be incorrect.

This procedure therefore gives us an estimated PDF, which supposedly reveals the market's beliefs concerning the potential behaviour of future exchange rates. We can study the way this estimated density function changes in response to Bank policy announcements to deduce what the financial market believes about the effects of the policy. This is kind of like guessing a person's body weight by observing how much they eat.

Interpreting the Results

The PDFs estimated by Levin, McManus, and Watt can be viewed in their entirety, as in the distribution plots in their paper. Alternatively, information about the distributions can be summarized by various statistics, such as the mean, variance, skew, and kurtosis of the distribution. I prefer the statistical summary approach, as it allows me to focus on certain features of the distributions that I can identify as particularly important.

One of the most impressive things this technique can do is estimate the market's "uncertainty" concerning future movements in the exchange rate. Uncertainty is often measured by the standard deviation—or volatility—of the distribution, since this provides an indication of how "spread out" probabilities are among potential alternatives. More volatility means more dispersion, which is often taken to imply more uncertainty.

Also important to gauging uncertainty is a measure of how likely market participants believe a big swing in rates might be. This is indicated by the amount of mass in the extreme tails of the estimated probability distribution. The more likely is an event that would move the exchange rate far from where it now sits, the more mass there would be in the distribution tails. One way to obtain an indication of this tail mass is by the kurtosis of the distribution. *Kurtosis* actually measures how pointed the peak of the distribution is. But if the peak is pointed, the tails are usually fat because the mass has to go somewhere. If the mass is not on the peak then it is often, though not always, in the tails.

Another way to see the effects of policy announcements is to look at trading volume. The argument goes that new information is often viewed differently by different traders, so that it leads to a divergence of opinion, which in turn leads to trading activity. High trading volume is therefore often taken as a sign that the market is uncertain what to make of new information, such as a Bank policy announcement.

To summarize, assuming that Levin, McManus, and Watt have deduced an accurate image of the market's beliefs, then increased uncertainty will be revealed by an increase in volatility, kurtosis, and trading volume. We can therefore use the paper's results to investigate the effects of Bank policy announcements. Levin, McManus, and Watt present their results by plotting the entire probability distribution. In my comments, I will instead plot key statistics over time to help focus more clearly on the uncertainty question.

Case Studies

Consider first Figure 1 in my comments, which covers the five days surrounding the 15 May 1997 announcement, which was supposed to indicate a change in the direction of Bank policy.

In Figure 1, the first line shown in the legend plots the absolute change in the futures price from day to day. The jump in this line on the day of the policy announcement suggests that the 15 May announcement changed the market's belief concerning the level of the expected future exchange rate.

The second line in Figure 1 plots volatility, which rose following the policy announcement. This increase suggests that the announcement raised uncertainty over future exchange rates.

The last line in Figure 1 plots the kurtosis of the estimated probability distribution. Note the peak on the day of the announcement. This kurtosis peak suggests that market traders at first interpreted the Bank's announcement as a sign that large swings in the exchange rate were suddenly more likely than had previously been the case. However, time passed and kurtosis fell, suggesting that traders considered extreme rate swings less likely. This up-down kurtosis behavior suggests that the market may have had difficulty interpreting the Bank's policy announcement. This view is supported by the trading volume curve in Figure 1. Note that volume shot up on the announcement day—and stayed up—as traders digested the new information. Volume fell only after two heavy days of post-announcement trading.

Figure 1 reveals the power of the PDF approach to uncover changes in market beliefs about Bank policy. According to Levin, Mc Manus, and Watt, the May 1997 report was supposed to signal a change in policy direction. The uncertainty this change created seems evident in the options market data. Of course, this conclusion is based on the key assumptions that the market was reacting to Bank policy, and not to other stimuli, and that the model employed accurately captures market sentiment.

Next consider Figure 2, which covers the November 1997 announcement that was supposed to indicate a stable monetary policy. Assuming that the market anticipated stability, the nature of the November announcement suggests that there should not have been any increase in uncertainty in the market. Figure 2 supports this view, showing little change in prices, volatility, or trading volume on the announcement day. Kurtosis actually fell throughout this period, suggesting that the market believed that wide swings in the dollar were becoming increasingly unlikely. Contrasted with Figure 1, Figure 2 depicts the effects of a calming—as opposed to uncertainty-increasing—policy announcement.

Figure 1
Monetary Policy Report, 15 May 1997


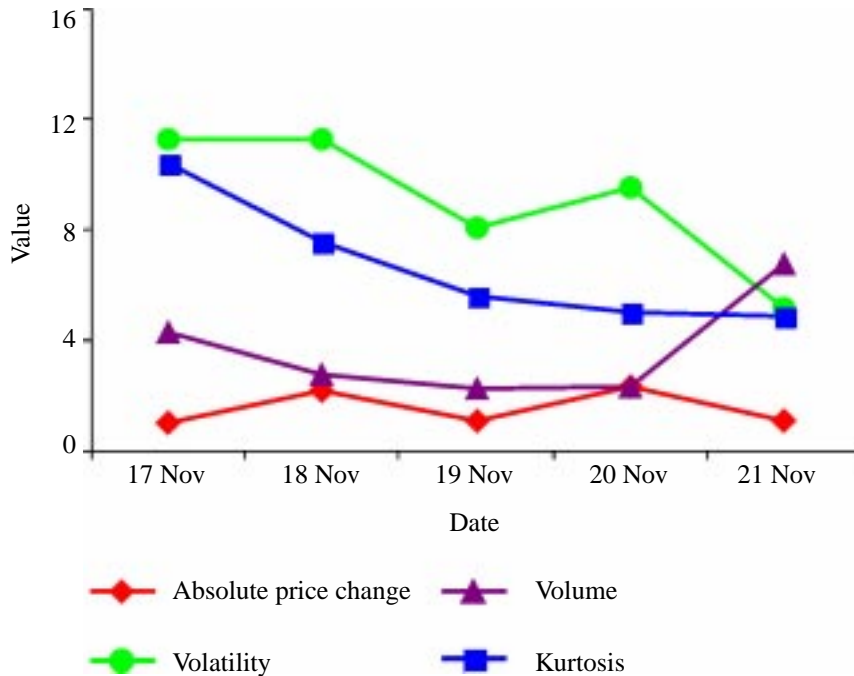


Figure 2
Monetary Policy Report, 19 November 1997



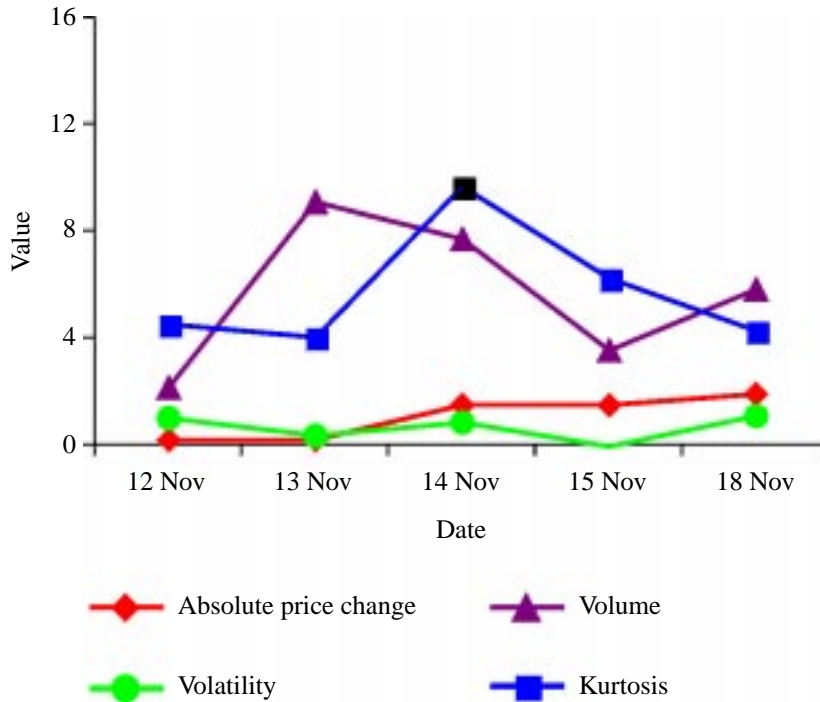
Finally, consider Figure 3, which plots our statistics from the November 1996 announcement. The November 1996 report was also supposed to have signalled a smooth continuation of previous policy and was therefore hypothesized to be a calming influence on the market. However, the statistics derived from the estimated PDFs, which are themselves derived from the options data, suggest that this may not have been the case.

The stable volatility curve in Figure 3 reveals that volatility did not increase following the announcement, consistent with a calming policy effect. However, kurtosis spiked on the announcement day, which suggests that the *Monetary Policy Report* released that day heightened the market's belief that wide swings in the exchange rate were suddenly more likely than before. Increased uncertainty about the *Report* is also suggested by the increase in trading volume, although oddly enough the increase occurred the day before the announcement, not the day of the announcement. Thus, the message we get from volatility and price changes runs counter to the signals from kurtosis and trading volume.

Perhaps the best way to describe the evidence from November 1996—as depicted in Figure 3—is “mixed.” There are no clear signals in either direction, and even some apparent inconsistencies between the various statistics. It almost seems as if the market was reacting to some force other than the Bank's *Report*, which is of course entirely possible. Indeed, Figure 3 shows why caution is required when trying to extract information about beliefs from financial market data. As in my analogy of guessing someone's body weight by observing how much they eat, one cannot always be certain that the data are giving you the signal you think you are getting. The model might be wrong or the analysis otherwise polluted such that an inaccurate image is rendered. In sum, Figures 1 and 2 seem easy to interpret and deliver the expected results, but Figure 3 is problematic.

The difficulty in interpreting Figure 3 brings me to my final two points. First, all of the statistics and distribution estimates produced in the Levin, McManus, and Watt paper really do need to have confidence intervals around them. These are estimates, after all, and like any estimate they are not certain. Before this technique can be usefully applied in practice, a procedure for calculating confidence intervals around the estimates must first be developed and employed.

My last point of caution is that the probability distributions Levin, Mc Manus, and Watt have calculated in their paper are actually risk-neutral probabilities, not the true probabilities investors use when making decisions. Thus, any results from the Levin, Mc Manus, and Watt procedure are based on the assumption that changes in estimated distributions are being driven by changes in pure event probabilities, not by changes in risk tolerance. This

Figure 3***Monetary Policy Report, 14 November 1996***

may be an acceptable assumption, but more work will be required to convince a skeptic.

I found the Levin, Mc Manus, and Watt paper a very interesting and thoughtful analysis of the issues at hand. It provides a potentially useful way to deduce the market's reaction to the Bank's policy announcements. There are definitely some bugs to be worked out, but I think the authors are moving in the right direction and I urge them to continue their research. This PDF-recovery technique is becoming widely used in industry—and to some extent by the Fed and other central banks—and is therefore a tool the Bank of Canada will need to possess if it wants to keep pace.