Discussion

Pierre Duguay

Some years ago, Chuck was presenting an address to the Canadian Macro Study Group. His intention had been to show how economic analysis had shaped policy developments at the Bank of Canada. But as he looked back, he realized that much of the analysis that had underpinned the move to inflation targeting, the conduct of monetary policy in the absence of reserve requirements, and the identification and mitigation of systemic risk in clearing and settlement systems, had been done in-house, without much contribution from the academic literature.

When it comes to inflation targeting, Lars has pretty much single-handedly redressed the balance. At the Jackson Hole Symposium last year, Chuck marvelled at his characterization of flexible inflation targeting as the solution to a dynamic optimal-control problem. I had an opportunity recently to compare the practical conduct of monetary policy in Canada with the theoretical solution derived by Lars. Rarely have theory and practice meshed so harmoniously.

So we were thrilled at the Bank when Lars not only accepted our invitation to contribute to this Festschrift, but chose to deal with the question put to him by Chuck in Jackson Hole: How do you deal with low-probability extreme events? Do you set policy on the basis of the mean of the distribution, at the cost of always being wrong? Or do you set policy on the basis of the mode of the distribution, so that you might at least be right most of the time, though at the cost of being a little more wrong if the extreme event does materialize?

Lars resisted the temptation to assume the problem away. He acknowledges that there would be little point trying to pre-empt an extreme event if there were still time to counter its effect after it happens. But he asks: What if

there weren't time? What if any policy action after the event could not undo its sizable impact on the target variable? Then we are left with a genuine trade-off. Lars leaves us no wiggle room.

But he does come up with a useful taxonomy to address the issue: The answer depends on your loss function. And, though he begs to be excused for not providing a definite answer, Lars certainly makes a very convincing case against the mode and in favour of the mean or the median as a measure of central tendency. He shows that it takes a rather bizarre loss function to give prominence to mode projections. And he reminds us that the mode may not be unique. I have no quarrel with any of these conclusions.

The distribution of forecasts for Canadian GDP growth and CPI inflation in 2004 (Figure 1) illustrates well the potential problems with the mode. And, of course, the odds are still better than even that the actual outcome will fall outside that range of projections.

While I appreciate Lars' tight set-up of the problem, I fear that it is somewhat contrived. First, it focuses on a single period; but the conduct of monetary policy is a repeat game. The object of policy is to keep the trend of inflation on target over time. If an event produces a lasting disturbance, monetary policy can always be adjusted after the event to undo its effect on future inflation. Thus, any unavoidable deviation from target that might arise from such an event would only be temporary. This would a fortiori be the case if the event did not have a lasting effect. Consequently, it seems to me that Chuck's question does beg re-examination in a multi-period framework. That may not alter the main conclusion, but it would raise some interesting questions such as: How does the central bank maintain credibility if it systematically misses the target by γa? How can it validate its view about the low-probability event? How accountable can it really be?

Second, there is something artificial in singling out one extreme event. There will always be some risk of large unexpected events (a quadrupling of oil prices, an LTCM, an Enron). There is an infinity of possibilities, and I find it difficult to justify singling out any one in particular. The nature of shocks is that they are unpredictable. Having the distinction to be flavour of the month does not make a possible event more worthy of attention than others, though it would surely expose the central bank to "I told you so" criticism if it materialized. Is that enough to warrant special attention? That, of course, depends on the loss function. But as Lars puts it: Whose loss function? Society's or the central bank's?

The focus on a single event can lead to a distorted perception of probabilities. How often have we seen analysts assigning probabilities of, say, 70 per cent to their base case and 30 per cent to an alternative scenario,

Figure 1

Distribution of forecasts from Consensus Economics Inc., April 2003

GDP growth
2004

Mean, median

3.35

CPI inflation
2004

Mean, median

2.05
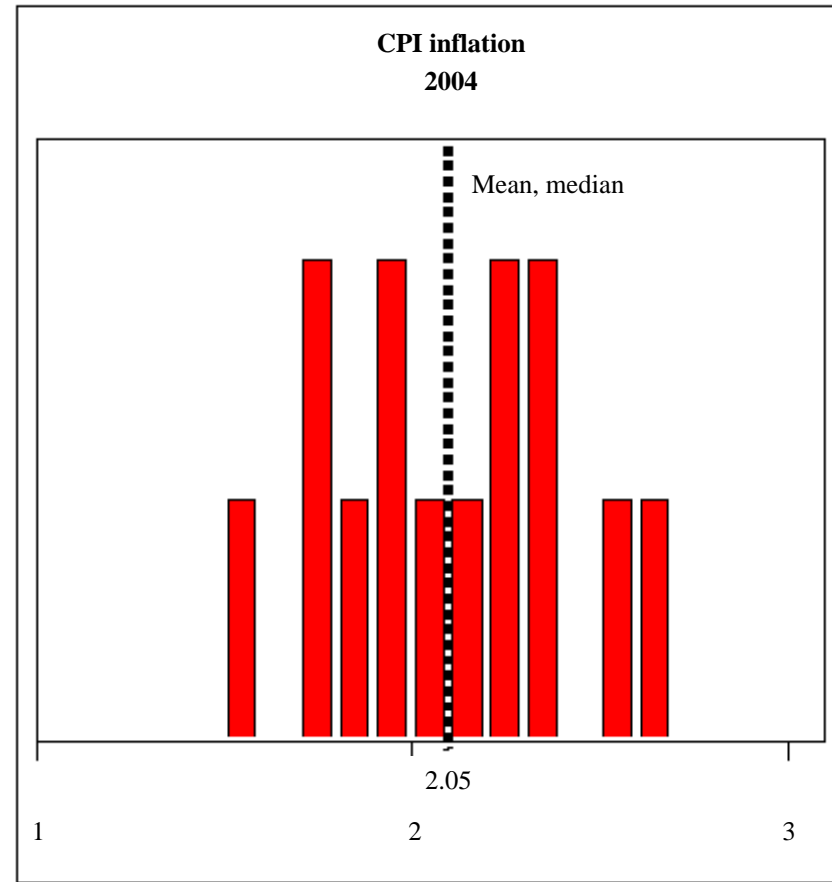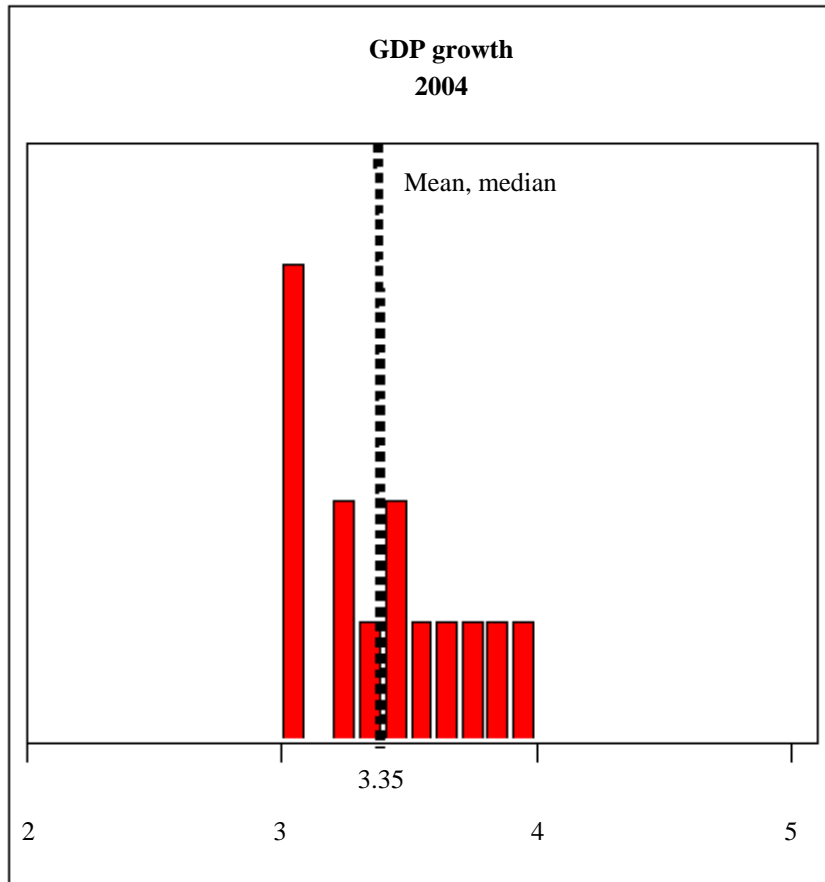
apparently unperturbed by the implied exclusion of any other possibilities? Now, if you feel generous, you will call it a semantic shortcut of little significance. And that may well be the case; but it still makes me uneasy.

Some of the discussions of possible extreme events have focused on the risk of a sharp asset-price movement. This strikes me as a very odd choice, given that asset prices, however hard to predict, are essentially endogenous. Typically, the event that is contemplated is the risk of a sudden correction. Correction presupposes that the price is judged to be out of line with fundamentals. And we all know that, sooner or later, asset prices will return to fundamentals. It therefore makes little sense to me to build a base-case projection on the assumption that asset prices would remain out of line indefinitely. Surely, the base-case projection should assume a return to fundamentals over some period. Uncertainty about timing would typically mean setting policy as if the return was spread over time, even if you suspected that once triggered by some random event the correction will take place suddenly. One could use Lars' framework to analyze which loss functions give a different outcome by comparing the case where the adjustment is assumed to be linear over n periods against one where there is a $1/n$ probability that the full adjustment takes place either in the current period or in any subsequent period until the adjustment occurs.

An intriguing variation on that theme is the case where the probability of a correction (and its severity) can be influenced by policy actions. Charles Goodhart has long argued, for example, that policy could reduce the severity of an asset-price correction by precipitating it early (pricking the bubble). In effect, this involves an intertemporal trade-off—missing the inflation target by a small amount now to avoid a larger deviation later, and goes well beyond the original question put to Lars. And it is not a question that can be examined in a one-period model.

Some of the events that are of interest to policy-makers arise from real-life non-linearities, such as the effect of bankruptcy risk on the cost of capital. These events may give rise to the possibility of multiple equilibriums. This possibility is generally ignored by the policy-maker's models, because they typically abstract from such non-linearities. Could these other equilibriums be tackled in the policy-maker's model as exogenous, low-probability events? Or should we try to model these non-linearities? As Lars reminds us, higher moments of the distribution do matter in non-linear models, and uncertainty equivalence does not hold.

One non-linearity that is receiving considerable attention these days is the nominal interest rate floor, which may limit the central bank's ability to respond to a large negative demand shock, thus heightening the risk of compounding the problem through deflation.[1] That would make it more critical to deal pre-emptively with low-probability extreme events.

To sum up, I found Lars' analysis insightful and persuasive. But I can't help thinking that our current preoccupation with low-probability extreme events is somewhat exaggerated and possibly counterproductive. The nature of shocks is that they are unanticipated. It seems to me that an alternative would be to investigate non-linearities and their consequences for the behaviour of the economy.

---

1. That risk is much reduced when expectations are well anchored to credible inflation targets. And the risk would disappear altogether under price-level targeting.