

Discussion

John Murray

Professor Goodhart has written a thoughtful and interesting paper on the conduct of monetary policy at the Bank of England. He uses a unique database, containing the inflation and output forecasts of the Monetary Policy Committee (MPC), to examine the Committee's behaviour over the past five years and to test a number of intriguing hypotheses. The descriptive nature of the paper provides a useful complement to the theoretical analysis presented by Lars Svensson, and offers a real-world example of how central bank actions might be conditioned by the prospect of low-probability extreme events (see Lars Svensson, this volume). Most important for our purposes here today, however, is that Professor Goodhart's paper reflects the sort of practical and insightful research that has characterized so much of Chuck Freedman's work over the past 30 years. Chuck has always been concerned not just with the theory of monetary policy but with its everyday application. In this, he and Professor Goodhart are kindred spirits.

Four Testable Hypotheses

Professor Goodhart's paper focuses on four related hypotheses. The first deals with the evolution of the MPC's inflation forecasts through time, and suggests that level of inflation averaged over all forecasts should converge on the Bank's inflation target as one moves n periods into the future (where n is the length of time that the MPC believes is required for monetary policy to have its full effect on the economy and is assumed to be about eight quarters). It also suggests that the standard deviation of the MPC's inflation forecasts should reach a minimum at n . Both predictions can be tested using Goodhart's database with a view to seeing (i) if the MPC in fact aimed for the mandated target of 2.5 per cent, and (ii) what the MPC's estimate of n was.

The second hypothesis starts with the observation that monetary policy actions influence inflation by first affecting the level of real economic activity. Assuming the economy is currently operating at or near potential, and inflation is not too far off course, one might expect (i) that the growth of output averaged over all forecasts would converge to the Bank's estimate of potential growth over a reasonably short time frame, and (ii) that the standard deviation of the output forecast would reach a minimum at time q , where q represents the MPC's estimate of how long it takes for monetary policy to have its full effect on the real economy and q is presumably less than n .

Goodhart's third testable hypothesis involves the relative size of the standard deviations for output growth and inflation. More specifically, he suggests that the standard deviation of forecast output growth should be significantly larger than that of inflation, provided the MPC is behaving as it should and adjusting interest rates to hit a predetermined inflation target of 2.5 per cent. Evidence that the standard deviation of output growth is smaller than that of inflation would indicate that the Committee has been more concerned with stabilizing output growth than inflation. It would also indicate that the MPC has been following an operating strategy much like Orphanides and Wieland (2000) and Orphanides and Wilcox (2002) have described for the United States. They have suggested that as long as inflation lies within an acceptable band, the Fed assigns a near zero weight to it and focuses instead on stabilizing growth in the real economy. When applied to the activities of the MPC, this strategy would predict that as long as inflation remained within the 1.5 to 3.5 per cent band established by the Treasury, the Bank of England would largely ignore the target and concentrate on real output growth.

The final hypothesis concerns possible asymmetries in the distribution of MPC forecasts, which could reveal interesting information about the MPC's utility function in the presence of asymmetric risk. Goodhart contends that the period under observation was characterized by persistent upside risk to inflation, owing to a suspected overvaluation of the pound, which MPC members worried could correct at any time. Assuming the MPC's utility function was quadratic, one would expect the MPC's mean forecast of inflation to equal 2.5 per cent, with the median and, especially, the mode of their forecast distribution lying below 2.5 per cent. Evidence that the median or the mode was closer than the mean to 2.5 per cent would suggest that the exponent on the utility function was less than 2.0 and that the MPC tended to place a low to zero weight on extreme outcomes in their policy deliberations.

Results

The four hypotheses described above were tested using data spanning the 23 quarters from May 1997 (when the MPC was first established) to November 2002 (when the most recent numbers were available). The main findings can be summarized as follows.

The evidence uncovered for the MPC's inflation forecasts was broadly consistent with the predictions of the first hypothesis. Although the forecast level of inflation typically tracked lower during the first few quarters, it then reversed direction and steadily converged to a level slightly above 2.5 per cent. The standard deviation of the inflation forecasts, meanwhile, declined continuously over the two-year horizon. The pattern observed for the output-growth forecasts was similar to the one for inflation and tells a coherent story, with growth first declining slightly to around 2.1 per cent—below most estimates of potential growth—and then rebounding to about 2.6 per cent—somewhat above consensus estimates of potential growth. Contrary to Goodhart's expectations, however, the standard deviation of output growth shows no sign of stabilizing over the two-year forecast horizon and falls monotonically. In other words, contrary to the second hypothesis, there is no indication that q is smaller than n .

The results related to hypothesis three—the relative size of the standard deviations for inflation and output—provided no support for the Orphanides-Wieland-Wilcox view and instead suggested, according to Goodhart, that the MPC had dutifully targeted a 2.5 per cent inflation rate. The standard deviation for inflation was approximately half as large as that for output growth over the first six quarters of the forecast horizon, and then fell abruptly to even lower levels during the last two quarters. “This result [Goodhart claims] is entirely inconsistent with the Orphanides-Wieland-Wilcox hypothesis, and entirely consistent with the hypothesis that the MPC's overriding objective has been to drive forecast inflation back into line with target at $t = 7$ and $7 = 8$, particularly the latter” (p. 164).

The results for the fourth hypothesis, involving the MPC's loss function, contained more surprises and suggested that the MPC attempted to put the modal forecast on target, effectively ignoring widespread views that the risks to inflation were biased on the upside because of a suspected imminent depreciation of the pound. Goodhart offers three alternative explanations for this. First, the exponent in the MPC's loss function may be zero, causing the Committee to dismiss the asymmetric risks. Second, the MPC may have an inflationary bias—preferring upside inflation risk to downside inflation risk. Third, the MPC might have been willing to gamble that the asymmetric risks would not be realized over the forecast period or that, if they were, sufficient

time would still remain for the Committee to adjust their instrument settings and deal with the problem.

Comments

The results that Goodhart reports are, for the most part, reasonable and non-controversial, as are his explanations for the more surprising outcomes. I would not be doing my duty as a discussant, however, if I did not identify some questions that occurred to me as I read the paper.

The first and most obvious point that one could make—indeed, Goodhart makes it himself—is that the sample is necessarily short and covers a period that was remarkably benign. Inflation never wandered far from its target level and output growth was uncharacteristically stable. The differences and asymmetries that Goodhart observes, therefore, are necessarily rather small and (probably) statistically insignificant. Consequently, it might be dangerous to read too much into the results.

My second reservation involves the important distinction that needs to be drawn between output growth and output levels when one analyzes inflation-output dynamics. Owing to the data limitations that Goodhart describes (see page 160 of this volume), he is forced to conduct his analysis under the (implicit) assumption that the inflation process is driven by changes in output growth rather than by the output gap.¹ Since the latter is most relevant for his purposes, however, we should not be surprised that the standard deviation of output growth bears little relation to that of inflation. Goodhart acknowledges this potential problem in his introduction but should probably have given it greater recognition in the concluding sections.

A related point can be made with regard to Goodhart's interpretation of the relationship between the standard deviations of inflation and output growth. There is no reason for the standard deviation of output, in other words, to exceed that of inflation. Nor is it obvious that the results taken at face value contradict the Orphanides-Wieland-Wilcox hypothesis. The fact that policy-makers might attach a zero weight to inflation as long as it remains within an acceptable band (say 1.5 to 3.5 per cent) does not mean that they would ignore the midpoint and adjust interest rates only when inflation threatened to go outside the band. Since the probability of exceeding the band increases as one approaches the upper and lower limits, policy-makers would presumably attach some importance to keeping inflation near the centre even

1. Note that the MPC publishes its forecast for the growth of output, but not for the gap; although the latter measure is even more judgmental, many commentators regard it as more critical for informing the interest rate decisions (see page 157).

in an Orphanides-Wieland-Wilcox world. It is difficult to draw any firm conclusions, since the behaviour of the MPC under the two regimes would be observationally equivalent. The results are not entirely inconsistent with an opportunistic loss function.

One final point that I would make relates to the discussion of the MPC's loss function and the three competing hypotheses that Goodhart puts forward. Here again, observational equivalence—and perhaps “conceptual equivalence”—may complicate our interpretation of the results. If I understand correctly the implications of Lars Svensson's elegant analysis, one of the reasons that policy-makers might adopt the sort of laissez-faire strategy inherent in Goodhart's third hypothesis is that the exponent in their loss function is zero. The first and third explanations that he provides, in other words, are essentially identical, although he assigns a zero probability to the first and a 75 per cent probability to the third. It is quite possible, however, that I am missing something here.

Conclusion

Let me conclude by noting how much I enjoyed reading Professor Goodhart's paper and how interesting I found the results. Although I have questions about parts of his analysis, he has identified a number of intriguing issues that warrant further work. I do not believe that similar data are publicly available for other central banks, but perhaps this will change as they strive to become more transparent, allowing outside academics to replicate and extend Goodhart's work. In any event, it should be possible for researchers at other central banks to accept Professor Goodhart's challenge and apply his tests to their own internal databases.

“It would be interesting to do a companion exercise for other countries, but I do not know of any countries where comparable data would make that feasible” (p. 167).

References

- Orphanides, A. and V. Wieland. 2000. “Inflation Zone Targeting.” *European Economic Review* 44 (7): 1351–87.
- Orphanides, A. and D.W. Wilcox. 2002. “The Opportunistic Approach to Disinflation.” *International Finance* 5 (1): 47–71.
- Svensson, L.E.O. 2004. “Optimal Policy with Low-Probability Extreme Events.” This volume.