

Discussion

1. David Wilcox¹

In his seminal 1987 monograph entitled *Models of business cycles*, Robert Lucas challenged the premise of Bob Gordon's paper for this conference – and, indeed, much of the premise of the conference itself. As you will recall, Lucas demonstrates that in the context of a standard model of household decision-making, the welfare consequences of economic fluctuations are astonishingly small. In fact, the representative household in Lucas's model would be willing to give up less than 0.1 per cent of its consumption each year in return for being rid of the magnitude of fluctuations that have been typical of the experience in the United States since the end of World War II. A small welfare consequence indeed.

If Lucas's claim were correct, the global phenomenon of the stabilisation of real activity over the past two decades or so would be of little import. But I suspect that few, if any of us, in this room believe that Lucas's result is correct. Why not?

One line of attack has focused on two key aspects of Lucas's original formulation: his assumption that the household sector can be adequately modeled by positing a representative agent who suffers only the economy-wide average amount of income variability, and his specification of the representative agent's utility function. Gadi Barlevy (2005) at the Federal Reserve Bank of Chicago reviews the papers that have relaxed these two assumptions and concludes that the benefits of eliminating business cycle fluctuations could be equivalent to as much as 2½ per cent of lifetime consumption – far greater than the amount derived by Lucas.

Moreover, another line of attack has questioned the assumption implicit in Lucas's original calculation that business cycle stabilisation affects neither the average level nor the average rate of growth of consumption. For example, Barlevy (2003) pursues the idea that business cycle fluctuations are bad for the average pace of real growth because fluctuations cause a dearth of investment in some periods and a surfeit of it in others. If investment is subject to decreasing returns, this piling-up of investment in some periods reduces the average pay-off to a given amount of investment over time. On the empirical front, Ramey and Ramey (1995) show an empirical link between volatility and growth in international data.²

To be sure, this debate is not settled, but our collective understanding of the issue has advanced greatly over the past two decades; undoubtedly, it will be possible to say the same two decades from now. At the centre of this debate are some of

-
1. I am grateful to Karen Dynan, Doug Elmendorf, Dan Sichel and Peter Tulip for their comments on these remarks. The views expressed here are my own and are not necessarily shared by either the Board of Governors of the Federal Reserve System or my colleagues on the staff.
 2. A methodological note of relevance to this conference arises from this. In the early part of their paper, Ramey and Ramey use an estimation technique that assumes that the variance of output is fixed over time, though it can vary across countries. Later in the paper, however, Ramey and Ramey adopt an alternative technique that allows the variance of output to vary both over time and across countries.

the keys for understanding why the topic of this conference, and the topic of Bob Gordon's paper, are so very important.

With that, let me turn to Bob's paper. This paper puts a wealth of useful material on the table and I am confident it will be widely cited in the literature on the stabilisation of real activity. While there are many interesting facts demonstrated in the paper, the real contribution is the emphasis it places on using an estimated structural model as the vehicle for attributing credit for the 'Great Moderation' to specific underlying causes. This strikes me as precisely the right approach and one that should be pursued extensively in future work on this topic.³

That said, the implementation in this paper leaves me sceptical, to say the least. For as Bob himself notes,

Perhaps the most surprising result in this paper is that, when monetary policy is assessed solely in terms of alternative Taylor Rule reaction functions and their effect, there was no difference between the 'Greenspan' monetary policy in effect in 1990–2004 and the 'Burns' reaction coefficients in effect in 1960–79 (page 104 of Bob's paper).

Surprising indeed.

In fact, I find this conclusion impossible to swallow, partly because it flies in the face of so much of the earlier literature claiming to demonstrate that Greenspan has respected the 'Taylor Principle', which counsels central banks to raise the nominal policy rate at least one for one with increases in inflation, whereas the monetary policy-makers of the 1970s did not. Bob criticises that literature as being based on empirical specifications that are 'plagued' with serial correlation, and claims to show that any appearance of obedience to the Taylor Principle on Greenspan's part disappears once the plague is cured.

But quite aside from the academic literature, there is the evidence of one's own eyes: the period 1960–79 ended with the US perceiving itself as on the brink of economic disaster, due in no small part to the fact that inflation seemed to be galloping out of control. By contrast, the period 1990–2004 witnessed a gradual decline of inflation, effectively towards zero after taking account of measurement bias. Indeed, in the last years of Bob's sample, the Federal Reserve made clear that it perceived a risk, for a time, of inflation moving too low, before that possibility receded, due partly to a dose of unusually stimulative monetary policy. Moreover, far from seeing itself on the brink of disaster, the country has been enjoying a productivity revival during the past decade, and a case can be made that monetary policy had a hand in fostering that revival.

Another way of making essentially the same point is this: the inflation objective that Bob posits (2 per cent) was never met between the late 1960s and the early 1990s. It is hard to believe, looking at actual inflation as shown in Gordon's Figure 7, that the inflation outcome after 1990 reflected substantially the same monetary policy as the inflation outcome prior to 1979.

3. Roberts (2004), discussed below, also specifies a structural model for the same purpose.

The source of Bob's finding is difficult to pinpoint. One possibility is that it reflects his assumption that the inflation objective was constant at 2 per cent throughout the sample period. If, during the period since 1990, the inflation objective actually drifted downward, a specification estimated under the counterfactual assumption of a constant objective would suffer from serial correlation. This explanation is far from satisfying, however, because it fails to explain why the post-1990 coefficients should be so much more sensitive to serial-correlation correction than the earlier coefficients. Much remains to be sorted out.

I will mention here two other aspects of Bob's paper, but, in the interest of concision, not pursue them. First, he notes that his estimate of the slope of the Phillips Curve has not changed much over a very long period of time. This puts him at odds with a good deal of the remainder of the empirical literature⁴ – a place, I should hasten to add, where I suspect Bob does not mind being. This literature argues that the Phillips Curve is substantially flatter now than, say, 20 years ago. While the evidence on this question is susceptible to alternative interpretations, a number of my colleagues at the Fed are sympathetic to the idea that the Phillips Curve is flatter now than before. But rather than seeing any flattening of the Phillips Curve as a defeat for that construct, they interpret it as a victory for monetary policy. In particular, they square the facts by positing that inflation expectations are now less responsive to fluctuations in resource utilisation than they used to be. As a result, in a fully articulated structural model, there is no instability in the equation that relates current inflation to expected inflation and other variables; instead, the instability occurs in the equation relating inflation expectations to resource gaps and other variables.

A second empirical note pertains to Bob's assumption that the equilibrium real interest rate (often denoted r^* for short) is constant throughout the sample period. The equilibrium real interest rate is embedded in the intercept of the estimated Taylor rule, along with the inflation objective and the coefficient describing the central bank's response to deviations of inflation from the objective. In the course of our normal analysis of the current economic situation, my colleagues and I expend an enormous amount of effort analysing the forces bearing on r^* . That effort leaves us convinced that r^* exhibits meaningful variation over time as, for example, equity values rise and fall, and productivity accelerates and decelerates. Future work might thus attempt to allow for time variation not only in the inflation objective, but also in r^* .

In the remainder of these comments, I would like to highlight some research conducted by colleagues of mine at the Federal Reserve that provides some interesting complements to Bob's paper.

In the first of these papers, Karen Dynan, Douglas Elmendorf and Daniel Sichel (2005) argue that financial innovation has contributed to the reduced volatility of real GDP growth over the past few decades. The forms of innovation they cite include the development of improved credit scoring and risk-based pricing of

4. See, for example, Atkeson and Ohanian (2001) and Staiger, Stock and Watson (2001).

credit, the securitisation of mortgage and other loans, and the emergence of the junk bond market. On the empirical front, they demonstrate – among many other facts – that the correlation between income and saving was markedly higher during the second half of their sample period, ‘just as we would expect if households can now borrow more freely in order to smooth consumption’ (p 24). They also show that a variable intended to capture the effects of Regulation Q, which imposed a ceiling on deposit rates, explains a substantial amount of the variation in residential investment before 1984, but essentially none after 1984. Both of these findings, and a range of other evidence they present, are consistent with the idea that financial innovation should be added to the list of candidate explanations (good luck, better inventory management, better monetary policy, and so forth) for why the volatility of real GDP growth in the US has come down.

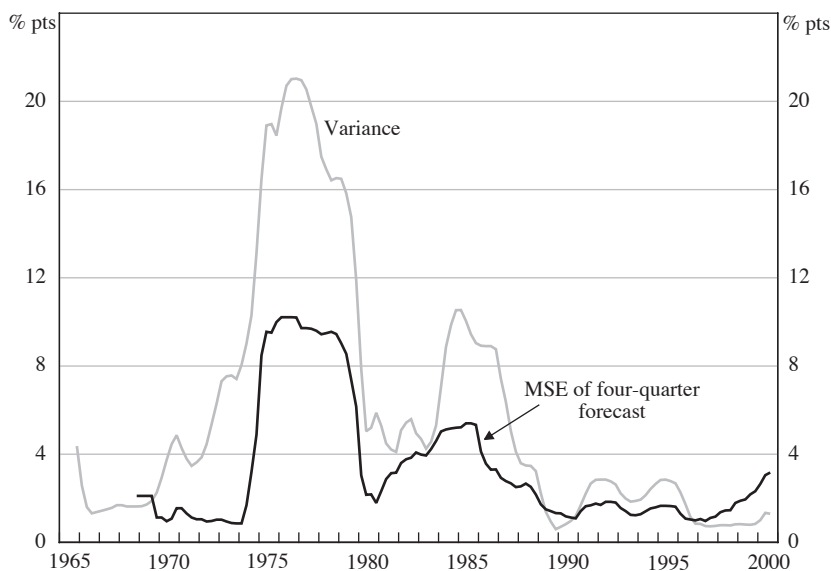
The second paper I wish to highlight is by John Roberts (2004), and is entitled ‘Monetary policy and inflation dynamics’. Like Bob, John explores the implications of changes in the conduct of monetary policy for changes in economic performance. John takes the alleged flattening of the Phillips Curve as the chief motivation for his paper – putting him in the camp that is at odds with Bob in this regard – but John addresses as well whether changes in the conduct of monetary policy can explain the reduced volatility of real GDP and inflation. John conducts his investigation using a three-equation model of the US economy similar to the one employed by Bob, though with enough differences to make a comparison of the two papers potentially very useful and good science. As well, John uses FRB/US, the Federal Reserve staff’s large-scale econometric model of the US economy.

In what I regard as a striking confirmation of Bob’s results, John finds that changes in the conduct of monetary policy can account for only a little of the reduction of GDP *growth* volatility, but ‘a large proportion of the reduction in the volatility of the GDP *gap*’, exactly consistent with Bob’s results. Oddly enough, however, and to prove that empirical economics is a tough, dangerous business, John finds only mixed results with respect to the ability of changes in monetary policy to explain the reduction in inflation volatility. Monetary policy does the whole job in John’s simple three-equation model, but accounts for very little of the reduction in inflation volatility in FRB/US.

The third paper of relevance to this conference is by Peter Tulip (2005). In a paper just completed, Peter argues that GDP growth variability *per se* probably does not reduce economic welfare very much (otherwise, why would seasonal variation generally be ignored in the literature on the Great Moderation?), so he shifts the focus to unpredictability. Putting the forecast team at the Federal Reserve Board under the microscope, Peter asks whether the variability of *errors* in predicting the growth of GDP has declined in the same way that the variability of GDP growth itself has. His answer is this: ‘less than you might have thought’.

Peter’s findings are summarised in four key charts – numbers 4–7 in his paper. I will focus on Chart 5 (reproduced below as Figure 1), which pertains to four-quarter-ahead forecasts of real GDP growth. On the basis of this figure and its companions, Peter draws the following conclusions:

Figure 1: Variance and Unpredictability of Four-quarter-ended US GDP Growth



Note: Dates refer to the end point of a 5-year rolling window

Source: Tulip (2005)

- First, as shown by the grey line, the variance of four-quarter-ended real GDP growth has come down in real-time data, just as Gordon and others have documented using fully revised versions of the data.
- Second, as shown by the black line, unpredictability has declined as well, though by distinctly less. Indeed, the predictable component seems to have all but disappeared, and since the early 1990s, the mean-squared error of our forecast has exceeded the variance of GDP growth itself, the variable we are trying to predict.⁵
- Third, the decline in unpredictability is much more evident at short horizons (four quarters or less) than at long ones.
- Fourth and finally, inflation variability has come down, and – in a contrast with the GDP results – so has inflation unpredictability.

One interpretation of Peter's results is that monetary policy has succeeded so well in its pursuit of macroeconomic stabilisation that it has squeezed *all* the predictable variation out of GDP growth. Lack of forecastability is exactly what one would have expected, on the basis of a simple control-theory perspective, as the end result of a fully successful stabilisation policy.

5. The disappearance of the predictable component of real GDP growth seems to contradict the finding of Blanchard and Simon (2001) and Cecchetti, Flores-Lagunes and Krause (this volume) that data from the US show no significant evidence of a change in the dynamics of real GDP growth.

In the fourth and final paper I want to mention, Sean Campbell (2005) explores the related question of stock market volatility. One might have expected a substantial reduction in the variability of real activity to cause a similarly substantial reduction in the variability of asset returns. Indeed, as Sean shows, this intuition is validated in the context of a standard consumption-CAPM (capital asset-pricing model): when fundamental underlying uncertainty goes down, the variability of asset returns declines as well. But the real world has behaved differently. As Sean documents, the variability of returns in the US has declined only slightly over the last half century or so. How can these two facts be reconciled? Faced with a collision between theory and facts, Sean makes the wise decision to confront the facts and throw out the theory. In place of the workhorse consumption-CAPM, Sean substitutes the model developed by John Campbell and John Cochrane (1999), that gives a prominent role to habit formation. The aim of Campbell and Cochrane in developing this model is to explain the size of the equity premium – a puzzle of long standing in the finance literature. The beauty of Sean’s research strategy is that he takes a model invented for one purpose and poses an altogether different question to it – namely, what should be the consequences for equity returns of a decline in the uncertainty about fundamentals? And the answer he derives is ‘not much’. That is, the model predicts that much of the variability in asset returns derives from habits, which are intrinsic to the utility of consumers and not affected by the character of the external environment. When the external environment changes, the behaviour of asset returns changes as well, but only a little. In other words, Sean shows that the Campbell-Cochrane model delivers a very realistic answer to a question that is altogether different from the one it was invented to explain. In short, there is no puzzle in the seeming disconnect between reduced volatility of real GDP and little to no reduction in the variability of asset returns.

A concluding question is this: is the ‘Great Moderation’ more likely to prove permanent or transitory? Forever is an awfully long time, but certainly some of the evidence seems encouraging. The innovations in inventory management that have been highlighted by many authors will not go away, nor will the financial innovations discussed by Dynan *et al* (2005). Neither does the extreme volatility of government purchases noted by Bob Gordon’s paper seem likely to return. As far as monetary policy is concerned, there is every reason to believe that the lessons that have been learned over the past twenty or thirty years will not be forgotten, and thus that whatever gains have accrued as a result will not be lost. The real wild card is the volatility of the supply shocks that seemed to have buffeted the real economy so dramatically in earlier times. If supply shocks are as important for the volatility of real activity as Gordon suggests, then a return to high volatility is well within the realm of possibility. On the other hand, it would be hard to argue that the last few years have been free of supply shocks – the most obvious example, though not the only one, being the fluctuations in the price of oil. And through it all, real activity around the world seems to have persevered amazingly well. Perhaps the most hopeful possibility is that the structure of the global economy has become more flexible, allowing even substantial supply shocks to be absorbed with much less distress than would have been the case in years gone by. If so, the Great Moderation may prove a long-lived phenomenon.

References

- Atkeson A and L Ohanian (2001), 'Are Phillips curves useful for forecasting inflation?', Federal Reserve Bank of Minneapolis *Quarterly Review*, 25(1), pp 2–11.
- Barlevy G (2003), 'The cost of business cycles under endogenous growth', NBER Working Paper No 9970.
- Barlevy G (2005), 'The cost of business cycles and the benefits of stabilization', Federal Reserve Bank of Chicago *Economic Perspectives*, 29(1), pp 32–49.
- Blanchard OJ and J Simon (2001), 'The long and large decline in U.S. output volatility', *Brookings Papers on Economic Activity*, 1, pp 135–164.
- Campbell SD (2005), 'Stock market volatility and the Great Moderation: fundamental news, return news, and the role of habit formation in asset pricing', Board of Governors of the Federal Reserve System, mimeo.
- Campbell JY and JH Cochrane (1999), 'By force of habit: a consumption-based explanation of aggregate stock market behaviour', *Journal of Political Economy*, 107(2), pp 205–251.
- Dynan KE, DW Elmendorf and DE Sichel (2005), 'Can financial innovation explain the reduced volatility of economic activity?', paper presented at the Carnegie-Rochester Conference Series on Public Policy (Financial Innovation, Risk and Fragility), University of Rochester, April 15–16.
- Lucas RE (1987), *Models of business cycles*, Basil Blackwell, Oxford.
- Ramey G and VA Ramey (1995), 'Cross-country evidence on the link between volatility and growth', *American Economic Review*, 85(5), pp 1138–1151.
- Roberts JM (2004), 'Monetary policy and inflation dynamics', Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series No 2004-62.
- Staiger D, JH Stock and MW Watson (2001), 'Prices, wages and the U.S. NAIRU in the 1990s', in AB Krueger and RM Solow (eds), *The roaring nineties: can full employment be sustained?*, The Russell Sage Foundation and The Century Foundation Press, New York, pp 3–60.
- Tulip P (2005), 'Has output become more predictable? Changes in Greenbook forecast accuracy', Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series No 2005-31.

2. General Discussion

A key focus of the discussion following Robert Gordon's paper was the assessment of the relative performance of various Federal Reserve Governors. Several participants felt that using the coefficients from the estimated Taylor rule to compare the performance of Burns, Volcker and Greenspan provides an unfair comparison, given that each reaction function is conditioned on the environment each Governor inherited and experienced. Moreover, some participants questioned whether the assessment could accurately capture the effect Volcker had in establishing credibility for the Federal Reserve. For this reason, there was support for the suggestion that

a better comparison would be to ask what Greenspan would have done if he had inherited the Governorship in 1979. In response to all this, Bob Gordon argued that the question he is asking is not which Governor was best, but rather, how the economy would have looked under different policy regimes. He also argued that the role of credibility is partially captured by the specific way he models supply shocks. In particular, he argued that better-anchored inflation expectations are manifest in a reduction in the magnitude of food-energy shocks, given that this variable is specified as the difference between headline and underlying inflation.

There was also some discussion about the specific modelling approach used in the paper. In support of David Wilcox, a number of participants questioned the appropriateness of a constant rate for target inflation across all periods. Some thought that the failure to accommodate for a change in average inflation would help to explain the apparent autocorrelation in the Taylor rule equation. There was also some concern expressed about the technique of examining changes in the residuals over time, particularly for the output equation, since these may reflect the performance of the model, rather than changes in actual economic outcomes. To this end, it was suggested that the focus on the *unpredictability* of output growth in Peter Tulip's (2005) research, mentioned in the comments by David Wilcox, is a useful alternative.

In a similar vein, one participant challenged the omission of fiscal policy from Bob Gordon's model, suggesting that the appreciation of the US dollar in the early 1980s – which was instrumental in holding down inflation – was due to the fiscal expansion of the time. They argued that the exclusion of this variable may influence the output equation, increasing the size of output errors. Similarly, there was considerable support for the argument that the decline in import prices in the early 1980s should be attributed to monetary policy, rather than beneficial supply shocks, given that the exchange rate is a channel of monetary policy.

Finally, there was some discussion about whether the slope of the Phillips Curve has flattened over the past decade, picking up on the comments made by David Wilcox. One participant questioned whether Bob Gordon had tested the stability of the coefficients on the supply shocks over time, while another argued that there is clear evidence of a reduced autoregressive coefficient on inflation after 1990. However, Bob Gordon maintained his view that the slope of the Phillips Curve has not changed significantly throughout the post-war period, referring participants to his earlier work on this issue.