Discussion

Nicholas Rowe

Dupasquier and Ricketts are looking for evidence of non-linearities in the short-run Phillips curve; they ask whether the slope of the short-run Phillips curve is the same everywhere (linearity), or whether it varies according to the output gap, or the level of inflation, or other such factors.

They conclude from their study that there is tentative evidence in favour of a convex short-run Phillips curve, so that successive increases in the output gap cause successively larger increases in inflation, but that there is less evidence for other non-linearities. I concentrate my remarks on this conclusion, because it is both their strongest, and also their most interesting result.

Before turning to my main theme, I might comment that their econometric method, which in general strikes me as being not only sophisticated but careful, has two minor problems. The first is that they estimate non-linearities of the short-run Phillips curve, but impose linearity, and verticality, on the long-run Phillips curve. If the long-run Phillips curve is not in fact linear and vertical, some of the non-linearity or non-verticality of the true long-run Phillips curve *may* be attributed by the estimation technique to non-linearity of the short-run Phillips curve. Second, they model the short-run Phillips curve with inflation as the dependent variable (so that residuals are assumed to be uncorrelated with output), rather than with output as the dependent variable (so that residuals would be assumed uncorrelated with inflation). I am unconvinced about which is the right way to do it. I can imagine both independent shocks to inflation (the goods and services tax, for example) *and* independent shocks to output (floods). If, for instance, measurement errors for output (as the independent variable) had

constant variance, they would bias the estimated slope of the short-run Phillips curve downwards, but would not affect linearity. But if the variance of those measurement errors were correlated with output, the bias would vary with output, and could make a linear curve appear non-linear, or vice versa. Testing to see if the results were robust to reversing the regression equation would satisfy these concerns.

But my main argument here is that Dupasquier and Ricketts' tentative evidence for a convex short-run Phillips curve in Canada provides much stronger support of the policy of price stability than is apparent at first sight.

Dupasquier and Ricketts interpret the Phillips curve as a relationship between *output* and inflation. But the Phillips curve can also be interpreted as a relationship between *employment* and inflation (or, as originally conceived by Phillips, between *unemployment* and inflation). Since output and employment are positively correlated over the business cycle, the two interpretations are equivalent for many questions. But the two interpretations are not equivalent if the question is the linearity of the Phillips curve, unless output and employment are themselves linearly related. If there is diminishing marginal product of labour, for instance, a linear relationship between output and inflation implies a concave relationship between employment and inflation (increases in employment cause diminishing increases in inflation); and a linear relationship between employment and inflation implies a convex relationship between output and inflation (successive increases in output cause successively larger increases in inflation). To put this another way, only if Okun's law is linear does linearity of the *output*-inflation Phillips curve provide a necessary and sufficient condition for linearity of the *employment*-inflation Phillips curve.

The above observation provides a way to test the validity of any empirical procedure that estimates non-linearities in the output-inflation relationship. First, repeat the procedure in exactly the same way and estimate non-linearities in the *employment*-inflation relationship. Second, estimate the non-linearities in the output-employment relationship. Third, check to see if the three estimates of non-linearities are mutually consistent. If they are not consistent, the empirical procedure is clearly not robust. If they are consistent, the results of the initial procedure should inspire more confidence.

My observation, that linearity of the output-inflation relationship is not necessarily the same as linearity of the employment-inflation relationship, also leads us to ask which of the two relationships it is that matters for monetary policy.

Suppose the Bank of Canada wants to maximize the average level of output. Suppose the Bank can choose the average level of inflation, and can

also choose the variance of inflation (or can at least reduce the variance of inflation, if it chooses, by investing in research and collecting information on the state of the economy). With a fully linear output-inflation relationship, with a vertical long-run Phillips curve and short-run Phillips curves all having the same constant slope, the Bank's problem is trivial. Any average level of inflation, and any variance of inflation, will lead to exactly the same average level of output (equal to the natural rate). Any investment the Bank makes in discovering the appropriate policy to achieve price stability, however defined, will be wasted. A non-linear Phillips curve, in some sense, is needed to justify the attention we pay to monetary policy and our unwillingness to replace the Bank of Canada by a roulette wheel.

But not all non-linearities would justify a monetary policy aiming at price stability. A concave short-run Phillips curve would warrant a policy of *increasing* the variance of inflation to increase average output. And a short-run convex Phillips curve that got steeper at high rates of expected inflation would warrant *high* average inflation to reduce the variance of output (unless high average inflation caused the variance of inflation to increase by enough to overcome this effect).

If the Bank's goal is to maximize the average level of *output*, the linearity of the *output*-inflation Phillips curve is what we must estimate if we want to evaluate the Bank's pursuit of price stability. If the Bank's goal is instead to maximize the average level of *employment*, we should instead estimate the linearity of the *employment*-inflation Phillips curve. But a more compelling assumption than either of these is that the Bank's goal is (at least ultimately) to maximize the average level of utility. If so, we must estimate the linearity of the *utility*-inflation Phillips Curve.

Suppose initially that the utility (U) of the representative individual is linear in output (Y) (because of constant marginal utility of consumption) but non-linear in employment (N), with increasing marginal disutility of employment, then

$$U = Y - V(N), \qquad V' > 0 \quad V'' > 0 .$$

And suppose that the production function has diminishing returns to labour,

$$Y = F(N) \qquad F' > 0 \quad F'' < 0 .$$

Inverting the production function yields

$$N = G(Y) \qquad G' > 0 \quad G'' > 0 .$$

Substituting the inverse production function into the utility function yields

$$U(Y) = Y - V(G(Y)) \qquad U'' < 0 .$$

Utility is now a concave function of output, since $V(\cdot)$ and $G(\cdot)$ are both convex. Because extra output requires increasing quantities of extra labour, and extra labour gives increasing quantities of extra disutility, utility first rises with output, at a decreasing rate, reaches (presumably) a maximum, then falls. In a competitive economy with no distortions, the natural rate coincides with the top of the utility function, but my argument does not depend on this. If, because of monopoly power, distorting taxes, or minimum wages, for instance, the natural rate of output is below the optimum (which certainly accords with the general presumption that booms are good times and recessions are bad times), an increase in output during a recession would still raise utility more than would an equivalent increase in output during a boom.

Substituting the output-inflation Phillips curve into the utility function yields a convex *utility*-inflation Phillips curve, unless the *output*-inflation Phillips curve is very concave. And even more, if we assumed diminishing marginal utility of consumption (and hence output), rather than the constant marginal utility of consumption assumed above, the utility-inflation Phillips curve would be even more convex.

In short, standard microeconomic assumptions imply that the utility-inflation relationship is more convex than the output-inflation relationship, which in turn is more convex than the employment-inflation relationship. If the Bank's goal is to maximize the average level of *utility*, rather than the average level of output, then a convex utility-inflation short-run Phillips curve is sufficient for the policy conclusion that the Bank should try to minimize the variance of output, which in turn *normally* means that it should minimize the variance of inflation. (I have put in the weasel-word "normally" here, because if there are supply shocks, like the GST for example, which cause the short-run Phillips curve to fluctuate vertically, minimizing the variance of output requires the Bank to accommodate the resulting inflation). This means that Dupasquier and Ricketts' finding—that the output-inflation short-run Phillips curve is probably convex—provides a much stronger support for the Bank's goal of price stability than would at first appear. With a convex utility-inflation short-run Phillips curve, a monetary policy that minimized the variance of unanticipated inflation (which is presumably *part* of what is meant by price stability) would maximize expected utility.

My demonstration above, that utility is a concave function of output over the business cycle, rests on a simple representative-agent model; it assumes in effect that all unemployment is equally shared, which is clearly false. Suppose instead that each individual is either fully employed or fully unemployed, and that the business cycle has no effect on the utility of an individual in either category, but merely varies the proportions of the

population who find themselves in each category. (The truth lies somewhere between these two extremes, since some individuals find themselves partly unemployed during a recession, or working more overtime in a boom.) Here, the average utility of the population could be linear in output (provided that unemployment is linear in output), since it is just a weighted average of the utilities of individuals in each category. But if recessions increase the *duration* of spells of unemployment (which they do), rather than just the probability of spells, and if utility is concave in duration of unemployment (a 1 per cent probability of two months' unemployment is worse than a 2 per cent probability of one month's unemployment), which seems reasonable, then expected utility is still a concave function of the unemployment rate, the utility-inflation Phillips curve is still convex, and price stability is warranted.