How Much Should Monetary Policy Take Output into Account, and What is the Best Way to Do So? *

Comments on Gaspar and Smets (2002) and Wolfers (2003)

Lars E.O. Svensson
Princeton University
www.princeton.edu/~svensson

I would like to thank the organizers for the privilege of discussing two of the fine papers presented here today. My assignment is to discuss Gaspar and Smets (2002) and Wolfers (2003), as well as to give my own view on how much monetary policy should take output into account and the best way of doing so.

Gaspar and Smets

Let me start with the Gaspar and Smets paper. I find it very interesting and well worth studying. Let me start from the standard objective for monetary policy that involves stabilizing inflation around an inflation target and stabilizing output around an output target equal to potential output. This can be formulated as an intertemporal loss function to be minimized,

$$(1) \qquad \mathrm{E}_t \sum_{\tau=0}^{\infty} (1-\delta)\delta^{\tau} L_{t+\tau},$$

with the quadratic period loss function

$$(2) \qquad L_t = (\pi_t - \bar{\pi})^2 + \lambda (y_t - \bar{y}_t)^2.$$

Here, E_t denotes expectations conditional on information available in period t, δ is the discount factor and fulfills $0 < \delta < 1$, π_t is inflation, $\bar{\pi}$ is the inflation target, y_t is output, $\bar{y}_t$ is potential output, and $y_t - \bar{y}_t$

* Presented at the Council on Foreign Relations/International Finance conference "Stabilizing the Economy," New York, July 11, 2002, and revised in February 2004. I thank Kathleen Hurley for editorial and secretarial assistance. Remaining errors and expressed views are my own.

is the output gap.[1] Finally, λ is a given positive weight on output-gap stabilization relative to inflation stabilization.

The main message in this paper is that, from a social point of view, a better macroeconomic outcome is achieved, if the central bank is assigned a loss function with a lower λ than the one corresponding to the social loss function. That is, we should have "weight- conservative" central banks, with a lower weight on output-gap stabilization than society's weight.

I remain unconvinced by the arguments presented in the paper. Instead, I believe that there is a good case for why the loss function of the central bank should have the true social weight on output-gap stabilization.

In order to demonstrate why I am not convinced by the paper's main point, let me start with section 2 of the paper. As a background, we can recall the old analysis by Kydland and Prescott (1977) and Barro and Gordon (1983). There is a Lucas-type Phillips curve, where the output gap is increasing in the excess of actual inflation over previously expected inflation. Furthermore, there is a loss function like the one above, except that the output target is not potential output but y^*, where by assumption the output target exceeds potential output, $y^* > \bar{y}_t$. Finally, the central bank is assumed to minimize the loss function under discretion, that is, without being able to commit to a particular systematic policy.

A standard result of Kydland-Prescott and Barro-Gordon is that, under these circumstances, an "average inflation bias" arises. That is, average inflation will exceed the inflation target, and the excess will be proportional to λ times the average excess of the output target over average potential output. Rogoff (1985) then suggested that it is better to assign a lower λ to the central bank, to have a weight-conservative central bank. This lowers the average inflation bias towards the inflation target. But unfortunately, this lower λ also creates "stabilization bias." That is, the relative stabilization of inflation and output is no longer optimal. In this particular setup, there will be too much output variability and too little inflation variability. Therefore, it is not optimal to reduce λ all the way to zero. In that case, the average inflation bias would be eliminated, but the output variability would be too large.

But, a much better way to handle the problem of average inflation bias, a way that is actually much closer to what has apparently happened in the real world, is to have "output-conservative" central banks. That is, the output target should be reduced to equal potential output. This is also consistent with the data.

[1] It is practical to write the weights as $(1 - \delta)\delta^{\tau}$, since then the weights sum to zero, and the intertemporal loss function can be seen as a weighted average of expected future period losses.

Rogoff's idea of a weight-conservative central bank as a reason for lower inflation implies that countries with lower inflation should, everything else equal, have higher output variability. But, in the data, this is not the case. Low-inflation countries do not have higher output variability. Thus, the idea of output-conservative central banks is a better empirical explanation of low inflation. As a recommendation for monetary-policy objectives, it also has the attractive property that it does not induce any stabilization bias.

Let us move away from the Lucas-type Phillips curve used in the old Kydland-Prescott and Barro-Gordon analysis to the newkeynesian Phillips curve used in this paper, where current inflation depends on expected future inflation, the current output gap, and a current serially correlated cost-push shock. Assume that the output target is potential output, so there is no problem with average inflation bias. In this setup, Clarida, Gali and Gertler (1999) have shown that, under discretion, a stabilization bias results. Furthermore, by reducing λ, this stabilization bias can be reduced. This appears to be a second argument for a low λ. Unfortunately, there are two problems with this argument.

First, the direction of the stabilization bias is actually not robust but model dependent. In Svensson (1997), with a Lucas-type Phillips curve with output persistence, the direction of the stabilization bias is the same as in the newkeynesian model with serially correlated cost-push shocks. However, in Clark, Goodhart and Huang (1999), with a Lucas-type Phillips curve with inflation rather than output persistence, the stabilization bias goes the other way. In that case, the central bank should have a higher λ than that of society. Thus, there is no reason to believe that the direction of the stabilization bias is robust.

Second, there is a better solution to the policy problem, namely a commitment to the optimal targeting rule. The optimal targeting rule is, for this loss function and model, of the form (Svensson and Woodford, 2004)

$$\pi_t - \bar{\pi} = -\frac{\lambda}{\kappa}[(y_t - \bar{y}_t) - (y_{t-1} - \bar{y}_{t-1})]. \tag{3}$$

That is, the "inflation gap" between inflation and the inflation target should be negatively proportional to the change in the output gap, with the proportionality coefficient λ/κ, where κ is the slope of the short-run Phillips curve. This simple and verifiable targeting rule will lead to the first-best optimal policy, with neither average inflation bias nor stabilization bias. Thus, there is no need to modify the λ.

Section 3 of the paper, on potential-output uncertainty and output-gap stabilization, relies on a paper by Smets and Wouters (2002). There are two kinds of shocks. One kind is a serially correlated supply shock that does affect potential output, and therefore there is no trade-off between stabilization of inflation and stabilization of the output gap. The central bank might as well stabilize inflation without any weight on output-gap stabilization, as if it had a zero λ. The other kind is a serially correlated demand shock. It does not affect potential output and does create a trade-off between stabilizing the output gap and inflation. If λ is positive, the presence of these demand shocks implies that the central bank must compromise between stabilizing inflation and stabilizing the output gap.

Suppose that the central bank acts under discretion. Suppose that the central bank cannot identify the separate shocks, and instead considers two extreme assumptions: (i) All shocks are supply shocks. This means behaving as if there is no tradeoff between stabilizing inflation and stabilizing the output gap, that is, as if λ equals zero. (ii) All shocks are demand shocks. This means taking the tradeoff and the positive λ into account. Finally, suppose that the central bank is extremely risk averse, so that it wants to minimize the maximum loss that arises when its assumption (i) or (ii) is wrong. That is, when the central bank assumes (i) but (ii) is right, or assumes (ii) but (i) is right. Smets and Wouters then show that the maximum loss is minimized if the central bank assumes (i) and hence behaves as if λ equals zero.

Gaspar and Smets interpret this result as another piece of support for a weight-conservative central bank. Again, I am not convinced. First, I believe that the result depends on the direction of the stabilization bias, and as mentioned, that direction is not robust but model dependent. Second, and more fundamentally, the linear-quadratic setup chosen here is a special case of the linear-quadratic model with imperfect information about the state of the economy that is examined in Svensson and Woodford (2003). For this case, certainty-equivalence holds. That is, the central bank should make its best estimate of the underlying state of the economy, the conditional mean of the unobservable state variables. It should then behave as if those conditional means were certain. This implies stabilizing inflation and the output gap taking the true λ into account. Furthermore, the central bank should estimate the state of the economy and the two kinds of unobservable shocks with an optimal Kalman filter. Thus, even if the central bank cannot identify the shocks perfectly, it should still construct the best estimate of the shocks and potential output and use those. It should not just assume that the shocks are either all supply shocks or all demand shocks. In particular, the optimal targeting rule will work fine, if inflation and the output gap are replaced by the best estimate of inflation and the output gap. This will best minimize the true expected loss.

Section 4 of the paper deals with private-sector learning and output-gap stabilization. It is inspired by a paper by Orphanides and Williams (2004), in which the private sector forms inflation expectations by

running regressions of current inflation on lagged inflation, in which case there will be quite a bit of persistence in expectations formation. The main result in section 4 is that inflation and output-gap stabilization is more successful if the central bank's λ is lower than society's. This is because it reduces the persistency of private-sector inflation expectations.

But is this mechanical learning, where the private sector mechanically runs regressions, really reasonable? The analysis may indeed apply reasonably well to the U.S., given the lower level of transparency of monetary policy and the uncertainty about any implicit inflation target there. But things are different in other countries. In countries with explicit inflation targeting and more transparent monetary policy, central banks do affect inflation expectations more directly via the inflation target, the inflation reports, published forecasts, and much more explicit communication with the private sector. Inflation-targeting central banks have been quite successful in making their inflation targets credible and in anchoring private-sector inflation expectations on those inflation targets. This has reduced the persistence of inflation in those countries. Orphanides and Williams (2004) indeed also consider the case where the inflation target is announced and known to the private sector. This simplifies the private sector's learning problem and inflation expectations become much more stable and less persistent.

Furthermore, the analysis in section 4 is under the assumption that the central bank does not understand that the private sector forms expectations by running regressions. But Orphanides and Williams (2004) show that, if the central bank understands how the private sector forms expectations and takes this into account, the optimal policy is quite insensitive to λ. Thus, in that case, it does not follow that the central bank's λ should be lower than society's.

Wolfers

Let me go on to my other assignment, Wolfers (2003). I find it a fascinating and challenging paper. The main message is that the general public's loss function is different from (2), approximately linear in inflation and quadratic in unemployment, with a high weight on unemployment variability. This suggests that there are large gains from stabilizing unemployment but little or no gain from stabilizing inflation. Thus, the main message is quite different from that of Gaspar and Smets (2002), and there is indeed some tension between these two papers. Unfortunately, I received the Wolfers paper too late before the conference to be able to provide a thorough discussion of it, so I only have a few superficial points.

The most interesting part of the paper is the one where the author actually tries to estimate people's preferences over inflation and output or unemployment. Some of the regressions are on levels and

standard deviations. I am not convinced that this is best. Most of the preferences used in theoretical work involve levels and squares or variances. Therefore, I find the regressions on levels and squares easier to interpret.

Wolfers regresses levels of happiness against levels of inflation and squared unemployment and simplifies by setting nonsignificant coefficients equal to zero. We can think of this as resulting in a loss function (where loss is negative happiness) of the form

$$L_t = \pi_t + \lambda_u u_t^2 . \tag{4}$$

That is, the loss function is linear in inflation rather than quadratic in the inflation gap and quadratic in unemployment. I denote the weight on unemployment stabilization by λ_u. Because of Okun's Law, the relative weight on unemployment stabilization is different from the relative weight on output-gap stabilization, λ. A simple version of Okun's Law is that a change of the unemployment rate of one percentage point corresponds to a change of the output gap of some 2–2.5 percentage points. Under this assumption, λ equals λ_u divided by 4–6.25. The value of λ_u that Wolfers finds is about eight (cf. Table 5, the ratio between the coefficient on unemployment squared and the coefficient on inflation). Then the corresponding value of λ is about 1.3–2, which may perhaps seem a bit on the high side. On the other hand, because (4) is linear in inflation, this is not the weight relative to that on inflation-gap *stabilization* but relative to that on the inflation *level* , which makes the comparison between (2) and (4) more complex.

The form of (4) implies that there are substantial gains from stabilizing unemployment but that the variability of inflation does not matter. This is rather different from the message of Gaspar and Smets. But there are several problems with Wolfers's results. First, the confidence interval around the λs is likely to be quite wide, because the confidence intervals on the coefficients that form the λ is high. It would be interesting to see explicit confidence intervals on λ. They would not be difficult to construct.

Second, there are some serious aggregation problems. The results are based on data from several countries over several years for many different individuals in many different situations. Each individual has his or her own individual preferences. Then all these preferences are aggregated. This procedure obviously results in some serious aggregation problems and difficulties in the interpretation of the results.

Third, the answers people give, how they evaluate their individual situation and their views on inflation and unemployment, must be interpreted as local estimates for individuals in very different situations.

Some face low unemployment and low inflation, others face high unemployment and high inflation. They could also be in different sectors of the same economy with different levels of unemployment.

Since all these results are local, they should not be used globally, that is, for very different levels of inflation and unemployment. However, suppose we do interpret the estimated loss function above as a good global loss function. Furthermore, suppose we believe in a vertical long-run Phillips curve. Then, in the long run, monetary policy cannot do anything about the average unemployment level, but it can do something about inflation. More precisely, the loss function (4) above, being linear and increasing in inflation, suggests that monetary policy should let inflation go to minus infinity. Do we really believe that result?

My own view on monetary policy and the output gap

My own view on monetary policy in general and monetary policy and the output gap in particular is, very simply, that central banks should be explicit, transparent and truthful. Some inflation targeting central banks are doing a fine job in this regard already. But they could do a few more things. For example, they could actually be explicitly about their loss function. They could try out different loss functions internally, evaluate the use of such loss functions, and if the evaluation is positive decide on which loss function to apply, before going public.

Let us again consider the standard intertemporal loss function, (1), with the period loss function, (2). There are a few parameters that need to be specified. First, the discount factor, δ, is for all practical purposes very close to one. It is therefore not very controversial. Second, the inflation target is already explicit in an increasing number of countries. Third, in several inflation-targeting countries, actual inflation is close to the target. There does not seem to be any average inflation bias in several countries. Therefore, it is clear in all but the rarest cases, the output target is potential output, not something higher. Potential output is an unobservable variable that needs to be estimated. This is difficult, both conceptually and empirically. Still, an increasing number of central banks are becoming more explicit about their measures of potential output.

The main sticking point is λ, the relative weight in the loss function on output-gap variability. One problem is whether people other than academics and sophisticated central bankers can understand the meaning of λ. I believe they can. Indeed, for a δ close to one, the intertemporal loss function is approximately equal to

(6) $$(\mathrm{E}[\pi_t] - \bar{\pi})^2 + \mathrm{Var}[\pi_t] + \lambda \mathrm{Var}[y_t - \bar{y}_t],$$

the squared deviation of the long-run (that is, unconditional) mean of inflation from the inflation target, plus the long-run variance of inflation, plus the weight λ times the long-run variance of the output gap. Thus, λ directly expresses the weight the central bank puts on a natural measure of output-gap variability, the long-run variance, relative to the weight on inflation variability and long-run deviations of inflation from the inflation target. A λ equal to one means that the central bank is equally concerned about output-gap variability as it is about inflation variability. A λ equal to a half means that the central bank is only half as concerned about output-gap variability as it is about inflation variability. This is not impossible to understand, I believe.

Thus, I believe that central banks should think about their λ, decide on the appropriate size of it, and go public about it. If there is a monetary-policy committee (MPC) as is the case in most central banks, each committee member may have their own individual λ. As suggested in Svensson (2003a), then they can simply vote about the committee's λ. A simple voting procedure to decide on λ would result in the median λ as the committee's λ, by the median-voter theorem. The median is good, because it implies that any extremist will not have any impact on the committee's λ. Thus, I believe there is no need to manipulate or modify the λ; instead central banks should be explicit about their true λ and apply that.

An important aspect of monetary policy is that there are lags between policy actions and the impact on inflation and output. Central banks cannot do anything about current inflation and output; they are predetermined by previous private- and public-sector decisions. The only thing is to look ahead and try to affect future inflation and output. This means that forecasts of inflation and the output gap have to be integrated in the decision process. Unfortunately, the papers presented at this conference ignore that and instead assume that central banks can affect inflation and output instantaneously. I personally prefer to work with models when the lags are explicit and taken seriously.

A good central bank must be able to construct conditional mean forecasts and make decisions from these. Theory tells us that mean forecasts are more relevant than mode or median forecast. Furthermore, a good central bank has to make forecasts of future potential output, so as to be able to make forecasts of the output gap. Many central banks have internal forecast of the output gap, some are reluctant to publish them, but some do indeed publish them.

The bank's forecasts are of course conditional on the bank's view of the transmission mechanism and on the bank's estimate of the current state of the economy. They must necessarily take into account a lot of judgment. There is no way you can allow models to exclusively determine policy: you always need plenty of judgment. In some of my work, I try to be more explicit about the way judgment should be

used (Svensson, 2003b). Furthermore, the bank must be able to do these forecasts conditional on alternative instrument rate paths. By looking at alternative forecasts for alternative instrument-rate paths, the MPC will get an idea of the tradeoffs involved between stabilizing inflation and stabilizing the output gap.

Once the MPC has an idea of the tradeoffs involved, it can find the optimal forecasts of inflation and the output gap, that is, the best compromise between the inflation forecast and the output gap forecast. This can be done in several different ways. If you have an explicit loss function, the staff can always assign a numerical loss to any forecast and explain which loss function has been used. The staff can also do experiments with targeting rules, that is, more explicit conditions that the inflation and output-gap forecasts must fulfill. Once the optimal forecasts of inflation and the output gap have been found, the corresponding instrument-rate path is the optimal instrument-rate plan. The current instrument setting is then the first element in the optimal instrument-rate plan.

Finally, the current instrument-rate setting is announced, together with the corresponding inflation and output-gap forecasts and the instrument-rate plan. Note that there is no room here for forecasts based on the assumption of a constant interest rate, an assumption that unfortunately is still used by several central banks. These constant-interest-rate forecasts should be abandoned and thrown on the garbage heap of history.

This is my view or vision of what it means to be explicit, transparent, and truthful. Although we are not there yet, we are actually not that far in some countries. The central banks of New Zealand, the U.K., and Sweden are not that far from this vision.

References

Barro, Robert J., and David R. Gordon (1983), "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy* 91, 589-610.

Clarida, Richard, Jordi Gali, and Mark Gertler (1999), "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37, 1661-1707.

Clark Peter B., Charles A.E. Goodhart, and Haizhou Huang (1999), "Optimal Monetary Policy Rules in a Rational Expectations Model of the Phillips Curve," *Journal of Monetary Economics* 43, 497-520.

Gaspar, Vitor, and Frank Smets (2002), "Monetary Policy, Price Stability and Output Gap Stabilization," *International Finance* 5:2, 193-211.

Kydland, Finn E., and Edward C. Prescott (1977), "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85, 473-491.

Orphanides, Athanasios, and John C. Williams (2004), "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in Bernanke, Ben S., and Michael Woodford, eds., *Inflation Targeting,* University of Chicago Press, Chicago, forthcoming.

Rogoff, Kenneth (1985), "The Optimal Degree of Commitment to a Monetary Target," *Quarterly Journal of Economics* 100, 1169-1190.

Smets, Frank, and Rafael Wouters (2002), "Strict Inflation Targeting: A Robust Monetary Policy Strategy under Potential Output Uncertainty," working paper, European Central Bank.

Svensson, Lars E.O. (1997), "Optimal Inflation Targets, 'Conservative' Central banks, and Linear Inflation Contracts," *American Economic Review* 87, 98-114.

Svensson, Lars E.O. (2003a), "The Inflation Forecast and the Loss Function," in Paul Mizen, ed. (2003), *Central Banking, Monetary Theory and Practice: Essays in Honour of Charles Goodhart, Volume I*, Edward Elgar, 135-152.

Svensson, Lars E.O. (2003b), "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules," *Journal of Economic Literature* 41, 426-477.

Svensson, Lars E.O., and Michael Woodford (2003), "Indicator Variables for Optimal Policy," *Journal of Monetary Economics* 50, 691-720.

Svensson, Lars E.O., and Michael Woodford (2004), "Implementing Optimal Policy through Inflation-Forecast Targeting," in Bernanke, Ben S., and Michael Woodford, eds., *Inflation Targeting,* University of Chicago Press, Chicago, forthcoming.

Wolfers, Justin (2003), "Is Business Cycle Volatility Costly? Evidence from Surveys of Subjective Well-Being," *International Finance* 6:1, 1-26.