General Discussion: Implications of a Changing Economic Structure for the Strategy of Monetary Policy

Chair: Malcolm D. Knight

Mr. Kohn: I find the discussion of difference versus level rules and approaches more than a little artificial. My point takes up perhaps a little bit on the latter part of Chairman Greenspan's remarks yesterday. It seems to me, particularly in an uncertain world, that the central bank needs to use all the information it has and cannot restrict itself even from using estimates of things like output gaps. Some notions of an output gap and a natural rate of interest are particularly important for forward-looking policy. You cannot forecast without some sense of the levels of these variables.

Vince Reinhart's aggressive hurry-up-and-wait policies, it seems to me, are even more dependent upon at least some notion of the output gap and the interest rate gap. How else do you know where to wait after you have hurried up? Changes also are important because they give a lot of feedback on the gaps. So, if you see changes in inflation and unemployment rates, if you see the economy responding more or less strongly than you expected to the stance of policy, you need to pay a lot of attention to that because that may tell you that your estimates of the interest rate gaps and output gaps are wrong. Policy, particularly in an uncertain world, must be totally information-encompassing and use both changes in variables and estimates of gaps, while recognizing that any gap or

potential gap is an unobservable variable and will be subject to considerable revision.

Mr. Berg: I have two short comments—one is for researchers and one is for policymakers. The framework of this paper could be used in order to discuss and analyze a little bit deeper the issues we discussed yesterday regarding the business cycle. In particular, it would be interesting to know if the preferences of central banks have changed over time. Therefore, it would be interesting to estimate the parameters of the objective function jointly with the parameters of model of the economy, as has been done in some recent papers by Richard Dennis and Giorgio Primiceri and others. Such an exercise would probably support what Chris Sims said yesterday. Systematic monetary policy has become more aggressive against inflation, but the less aggressive policy pursued in the 1970s in the United States, for example, was not loose enough to cause the high inflation then. The inflation peak in the 1970s is probably, to a large extent, due to exogenous shocks. Such an exercise would also be important in order to find out if the output gap is a significant variable in the Fed's objective function or enters estimated rules only because it affects future inflation.

The second comment: Reading this very good paper got me to thinking about in what ways central banks can contribute to reducing uncertainty. It is quite clear that we cannot do away with parameter uncertainty or uncertainty regarding exogenous disturbances. But what about imperfect information? By being a little more specific about targets and forecasts, a central bank could do something. In particular, there might be room for being more specific about the judgments, which are so important in the forecasting process.

The discussion yesterday on the Borio and White paper, for example, showed there is no consensus on the role of asset prices for monetary policy. Some of us, including myself, think that monetary policy should respond to asset prices to the extent they are deemed to affect aggregate demand and inflation. Others claim that central banks should react to asset price misalignments over and above the reaction

to inflation forecasts. Glenn Stevens underlined that financial crises are low-frequency events. This means it may take many years, in my view, to reach a consensus. In order for a central bank to reduce uncertainty regarding the rule of asset prices, we could already today be a little more specific about our judgment regarding asset prices. We could more clearly explain how the development of asset prices is assumed to affect aggregate demand, output, and inflation. Or, if it is not assumed to do this, we should explain it. We could be a little bit clearer about our own judgments.

Mr. Meyer: I want to expand on the comments of David Longworth and Don Kohn on level versus change rules. I am going to have a somewhat sharper set of conclusions here. Carl, I think you have overstated the potential benefits of a change rule. Indeed, a change rule may be extremely dangerous, particularly in the circumstance that you are worried about—structural change.

You talk about attenuation. You talk about change rules. Longworth and Kohn were talking about a third approach, which I will call "continuous updating." Let me give you a very simple example. Start with the level of the funds rate, depending upon the unemployment gap (I'll do it with the unemployment gap). Take a first difference of that. The change in the funds rate now depends on the change in NAIRU and the change in unemployment. What you are effectively telling us to do is to throw away the change in the NAIRU. It is really nice to get rid of the NAIRU. It is not even in the rule and we are uncertain about it. This seems very attractive unless the structural change is a change in the NAIRU. What was going on in the second half of the 1990s was very, very much about that. What you are telling is, "Forget about the structural change, forget about the decline in the NAIRU, and pretend it didn't happen." I want to ask you to evaluate for us what the change rule would have instructed policymakers to do in the second half of the 1990s compared with what actually happened.

One last point: You tell us that policy rules that focus on the actual objectives are more robust. Your rules do not function on the objec-

tives. You picked an artificialist and fed it. A change in the gap, when it is really the level of the unemployment gap in particular that is important, you cannot throw that away. So, you have to use all available information. You have to continuously update your estimates of the gap and NAIRU, particularly in periods when changes in those are at the heart of the challenge to policymakers.

Mr. Mussa: I appreciate this type of work. It is useful and illuminating, but I don't think its importance should be exaggerated. Bob Lucas, in his AEA presidential address last year, asserted forcefully the proposition that the maximum that improvements in stabilization policy in the United States over the last 50 years, stabilization had achieved the optimum of smoothing out all of the around-trend variability of consumption in the United States. The maximum contribution to welfare would have amounted to no more than two-tenths of 1 percent of GDP. I don't exactly agree with that for a variety of reasons I don't want to go into here, but if I focused not on the last 50 years but on the last 20 years or so, the U.S. economy has had two comparatively mild and brief recessions. The inflation rate has come down from around 4 percent when the Fed knocked it down in the early 1980s to now 1 or 2 percent.

It is difficult for me to imagine macroeconomic stability performance in the longer term that is ever going to be much better than that. There are real shocks to the economy that are going to induce macroeconomic fluctuations, and even the best attuned monetary policy is not going to smooth all of those things out. We are talking at the margin here. Keep doing what you are doing now, and if you can improve it a little bit at the margin that is fine.

The main lesson from experience is: Avoid the big mistakes. The Fed made a big mistake in the 1930s when monetary policy failed to respond adequately to the contraction of the money supply. That was a very costly mistake. Let's not do that one again! Also in the succession of business-cycle expansions we had until the 1980s—where successively in each business-cycle expansion inflation started out a little

higher than it started out in the previous one and ended up higher than it ended up in the previous one—we gradually ratcheted the inflation rate up until it reached 13 percent at an annual rate. In the last expansion of late 1970s and 1980-81, that was a serious mistake of monetary policy. That lesson has been learned. The cost that was paid to restore Federal Reserve credibility in the deep and prolonged recession of the late 1970s and early 1980s will be remembered and not repeated. What we are talking about now, at least for the United States, is second order.

That is not true, however, for all countries around the world. Frenkel, Ortíz, and others have mentioned that. The situation in many emerging-market countries is very different. When the credibility of monetary policy is not firmly established, then the stakes are very much bigger.

Mr. Sims: I also think that the exercises in this paper are interesting, informative, and important, but the paper makes claims about the contrast between Bayesian and robust control methods much too strong. This is unfortunate for two reasons. One is the actual language of policy discussion is the language of Bayesian decision theory. You could see that in Alan Greenspan's speech yesterday. His description of the policy process, the need to consider both the probability of deviations from model forecasts and the losses under various contingencies, and the need to recognize that no objective analysis of the data is ever going to be completely sufficient for the determining judgments about probability—that is all the Bayesian language. That is one reason to think it is unfortunate to force a sharp contrast between robustness and Bayesian methods.

On top of that, there is the fact that the contrast is based on a misconception. It is just wrong. At the beginning of the paper there is a claim that there can be circumstances in which it is impossible to put probability on a contingency. It is never impossible to put a probability on a contingency that we can imagine, and robust control cannot help us imagine things that we could not otherwise imagine. In fact, when we take decisions, we are always implicitly putting probabilities

on contingencies. Nonetheless, robust control exercises are quite useful. Why is that? It is because we have to recognize that the Bayesian paradigm, indeed any mathematical optimal decision-making paradigm that we have available today, does not properly recognize the inevitable limits on our analytic capacity. So, Bayesian exercises are always done with imperfect assessments of prior probabilities and, indeed, with incomplete lists of contingencies. Naïve Bayesian approaches that start out by staying, "Let's write down a prior," always end up with some crude approximation to a real prior. And they can easily end up putting insufficient weight on very important contingencies. The advantage of robust control from a Bayesian perspective is that it allows us to pick out worst-case scenarios, which may, indeed, be things that we will have ignored or forgotten to assess well if we try directly just to assess probabilities.

On the other hand, a robust control exercise should never stop with picking out the worst-case scenario. The worst-case scenario depends always on the range of alternatives that has been considered, and it can easily turn out to be something that is just so unlikely that it is unreasonable to base policy on it. As David Longworth's comment pointed out, "We look like we are doing robust control in times like 9-11." When times are normal and stable, the probabilities of the worst-case scenarios are so low, it does not make sense to condition policy on them. But in times like 9-11, we start giving them more probability.

Mr. Walsh: Let me go in reverse direction and start with Chris Sims. I would not fundamentally disagree with his point. In some respects, my purpose here was to contrast the two approaches, and I am sure I might have overstated the differences. But the purpose was also not to necessarily propose robust control as optimal or the way one would necessarily want central bankers to behave, but just to try to discuss some of the differences in the perspectives and how that might play out in specific cases.

One of the interesting things about thinking about alternative attitudes toward risk, which is one of the things that I started to think

about as a result of the robust control literature, is that it does tie back to Mike Mussa's point about Lucas' estimate of the costs of the business cycle, which he claimed to be quite trivial. That is basically a function of his use of standard expected utility theory, which makes the costs look pretty trivial. There is some interesting work that looks at different attitudes toward risk and that does a better job of explaining both macro fluctuations and things like the equity premium, which standard models have great difficulty explaining. These alternative preferences can suggest that the potential cost of business cycle fluctuations is much larger than the very trivial levels that Lucas came up with. That also ties into the issue of how we think about risk and how we evaluate risk and uncertainty.

In response to Larry Meyer, I don't focus specifically on unemployment-related rules. The Orphanides and Williams result was that as you added more uncertainty about the NAIRU, their optimal policy rule converged to this first-difference specification. In just the simplest case (say, a one off change in the NAIRU), let's say an increase that raises the actual rate of unemployment, you respond to that change. But then the unemployment rate is no longer changing, so you are not continuing to revise the interest rate. The level rules will leave you thinking there is a gap between the unemployment and the NAIRU until you have updated your forecast. I don't want to totally dismiss the necessity to keep a reference level in there as well. While I may have understated the contrast, you may have overstated it.

Now, let me come back to a point that Longworth made, and it relates to Don Kohn's point as well. I was using the term "guideline" in the David Longworth sense and not in the Lars Svensson sense. One of the things that central banks typically do is run model forecasts to help guide the choice of policy. Implicitly in that exercise one needs to include some specification as to how the future interest rate will evolve, unless one is adopting something like the Bank of England approach of looking at constant interest rate forecasts.

What assumptions do you make about the evolution of policy choices in the future? Simple rules are one way of capturing the systematic behavior that policymakers might follow. They can be useful for constructing model forecasts. In part, one of the lessons I draw from recent work is that it can be useful not to be restricted to the Taylor-type formulation, as has been the case for the last few years, but to look at alternative specifications of those rules, particularly ones that may be relevant when there are serious issues of measurement error.

Finally, looking at worst-case scenarios is also a useful benchmark for evaluating policy to try to see what the downside risk would be of particular instrument paths, even though that doesn't necessarily mean you are adopting a formal robust control approach or designing your policy specifically to address that worst-case scenario.

In the Jensen paper on nominal income growth targeting and my earlier work on speed-limit targeting, the notion of targeting rule was to think about those variables—nominal income growth or the change in the output gap—as being among the objectives of the central bank. The central bank behaved optimally to minimize fluctuations in inflation and changes in the output gap. Here, it is a slightly different use. I am looking at simple rules in which the objectives are defined in terms of inflation, output gap volatility, and I don't change those, but instead look at alternative simple instrument rules that might include levels versus differences—just reflecting the difference between a targeting rule and a simple rule.

Mr. Cotis: I liked the paper and I found it very interesting and very instructive. I am a bit worried, however, about reaching premature policy conclusions. And, as you said, the results can be highly model dependent. One of my worries, for instance, is your focus on first-difference rules. I understand the rationale for dropping the level of the output gap in your optimal monetary rules, but it may have drawbacks. Dropping the level of the output gap would make communication to the public barely audible. It is a bit unnatural and it's really not the way policymakers think. They think in terms of levels. Dropping output

gap levels may also lead to misguided policies. To give just an example: What would happen if the impact of the output gap on inflation were either small or slow to come? Here, I would take the case of Continental Europe today. The current situation is one of high nominal inertia with the euro area inflation stuck around the 2 percent medium term target, combined with a substantial and slow-moving negative output gap, in a context of subdued recovery. What would the difference rule entail in this situation? Basically, it would lead to monetary inaction. And this lack of responsiveness from monetary policy could lead in turn to a situation where the economy is locked into a high unemployment equilibrium trap in the long term. In all, keeping in output gap levels seems unavoidable in that they provide monetary rules with an essential error correction mechanism.

Mr. Freedman: I want to continue with the same theme of the level of gap versus change of gap. I found the paper very interesting, but as I was reading it I was thinking to myself that I have spent the last eight or nine years trying to convince participants—both in the market and in the economy—that we look at the level of the gap. For example, in the current case in the United States, even if you got 3½ or 4 percent growth one year, it does not mean that interest rates have to rise if you started out with quite a lot of slack.

On a more technical level, the case that is often cited is the case that Orphanides talked about, in which there was a major policy error in the 1970s because of a lack of appreciation of the change in the potential rate of growth. We are sensitive to that. We have learned a lot from that episode. We are now more sensitive to the possibility of making those types of errors. That is why Dave Longworth commented that we look at a lot of information. In your own estimates—the little model-building exercises—you use data covering that period as well. If you took out that period and just looked at, say, a later period—I don't know if you have enough data—maybe just the earlier and the later period and leave out the period of the 1970s where there was clearly an error resulting from a misestimate of the size of the level of the gap, would you get the same results favoring change of gap over level?

Mr. Rogoff: One recent development in monetary theory, certainly over the last five to 10 years, draws attention to the rethinking of the Brainard result when you have uncertainty. Do you have a smaller response or a larger response? Brainard basically said you do a little less when you are uncertain. And Hansen and Sargent said, "No, the opposite, if there is some chance of a bad outcome, you have to overkill and react very aggressively."

The thing the paper draws attention to, once we go beyond very simple rules and once we go beyond very simple frameworks, is how difficult it is for the models to come to grips with what the right answer is. It is a very perplexing myriad of results that one gets. It echoes the point the Chairman made about we cannot rely on any one single model.

If I might draw an analogy of this question of overreaction, that might be interesting. I have friends who are a couple who taught at a Midwestern university, let's just say it was in the Kansas City Federal Reserve District, and they read Truman Capote's *In Cold Blood,* and became terrified that something would happen at night. When the husband traveled, he decided that his wife should sleep with a shotgun at her bed. Even though she was a fairly good marksman, she didn't have her glasses when she was sleeping. But with the shotgun, she could just point it at the wall and it would blow out the wall to destroy any intruder. I think you see the analogy here, which is, if it were an intruder, it's fine. If it were my friend, who happened to come home early from his trip and was fumbling for the keys, it might not have had such a happy outcome. It is very hard to know this issue from simple models, whether to respond over aggressively or not.

Mr. Walsh: Just to follow up on that, robust control would tell her not to fire the shotgun because the worst-case scenario is that it is her husband who comes home.

Mr. Longworth: On the output gap, it is certainly true that policymakers are frustrated by the fact that it is so uncertain, but we don't

want to throw the baby out with the bath water. Therefore, I agree strongly with Don Kohn's comments. What we have to do is the best we can with all the information that we have on the output gap. Yes, it may mean that relative to a world where the gap was certain we should down-weight it, but we still have to concentrate on the level of the estimated gap.

Mr. Walsh: I'd like to make a quick note on the issues of communication that were raised by a couple of the participants. One of the reasons I originally started looking at speed-limit policies was from reading the press releases from the FOMC in 1999, at a time when interest rates were being increased. The language was very explicit that the FOMC was concerned about the growth rate of demand relative to the growth rate of supply, which is essentially the change in the output gap. Certainly that was a period in which the Fed was perceived as communicating well in terms of its policy intentions. On that respect anyway, you can keep your focus on the objectives of stabilizing fluctuations in real economic activity, while controlling inflation and communicate that in a variety of ways, some of which may involve discussion of where the economy is going relative to potential instead of just in terms of the level of the output gap.