————. 2003b. "Expectations and the Stability Problem for Optimal Monetary Policies." *Review of Economic Studies* 70:807–24.

————. 2006. "Monetary Policy, Expectations and Commitment." *Scandinavian Journal of Economics* 108:15–38.

————. 2009a. "Expectations, Learning and Monetary Policy: An Overview of Recent Research." In *Monetary Policy under Uncertainty and Learning,* K. Schmidt-Hebbel and C. E. Walsh, 27–76. Santiago: Central Bank of Chile.

————. 2009b. "Learning and Macroeconomics." *Annual Review of Economics* 1:421–51.

————. 2009c. "Robust Learning Stability with Operational Monetary Policy Rules." In *Monetary Policy under Uncertainty and Learning,* K. Schmidt-Hebbel and C. E. Walsh, 145–70. Santiago: Central Bank of Chile.

————. 2013. "Learning as a Rational Foundation for Macroeconomics and Finance." In *Rethinking Expectations: The Way Forward for Macroeconomics,* edited by R. Frydman and E. S. Phelps, chapter 2. Princeton, NJ: Princeton University Press.

Gaspar, V., F. Smets, and D. Vestin. 2006. "Adaptive Learning, Persistence and Optimal Monetary Policy." *Journal of the European Economic Association* 4:376–85.

Molnar, K., and S. Santoro. 2010. "Optimal Monetary Policy When Agents Are Learning." Norges Bank Working Paper 2010/8.

Orphanides, A. 2002. "Monetary Policy Rules and the Great Inflation." *American Economic Review, Papers and Proceedings* 92:115–20.

————. 2003a. "Monetary Policy Evaluation with Noisy Information." *Journal of Monetary Economics* 50:605–31.

————. 2003b. "The Quest for Prosperity Without Inflation." *Journal of Monetary Economics* 50:633–63.

Orphanides, A., and J. C. Williams. 2002. "Robust Monetary Policy with Unknown Natural Rates." *Brookings Papers on Economic Activity* 2:63–118.

————. 2005a. "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning and Expectations." *Journal of Economic Dynamics and Control* 29:1927–50.

————. 2005b. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." In *The Inflation-Targeting Debate,* edited by B. Bernanke and M. Woodford, 201–34. Chicago: University of Chicago Press.

Preston, B. 2008. "Adaptive Learning, Forecast-Based Instrument Rules and Monetary Policy." *Journal of Economic Dynamics and Control* 32:3661–81.

Primiceri, G. E. 2006. "Why Inflation Rose and Fell: Policy-Makers' Beliefs and US Postwar Stabilization Policy." *Quarterly Journal of Economics* 121:867–901.

Sargent, T. J. 1999. *The Conquest of American Inflation.* Princeton, NJ: Princeton University Press.

————. 2008. "Evolution and Intelligent Design." *American Economic Review* 98:5–37.

Sargent, T. J., N. Williams, and T. Zha. 2006. "Shocks and Government Beliefs: The Rise and Fall of American Inflation." *American Economic Review* 96:1193–224.

Discussion

Christopher Sims had technical comments. Constant gain learning is easier to analyze in a theoretical model than in one with constant parameter change, but there is no excuse in this empirical exercise to use it instead of a Kalman filter with an explicit model for parameter change. The Kalman

filter distinguishes between periods when there is a lot to be learned about parameters and when there is not, and there could be a different historical pattern derived from it. It does not discount mechanically, but according to how the explanatory variables move. Sims then moved to Honkapohja's discussion on the research of Thomas Sargent, clearing up the issue that Sargent's early work is at complete odds with his later work. Early on, Sargent felt that this entire issue dealt with escape dynamics, whereas later on he contradicts that. In the end, it might be econometrics is to blame for bad policy. In order to get bad policy, you need to have people stuck on theories that are bad. Using this flawed theory, people were estimating rich econometric models, but they should have known that the theory was not working. Everyone was stuck on the two equation models, and that was the problem.

Lars Svensson liked the chapter, but felt it had a big inconsistency. The authors seem to perform optimal control without optimal filtering. If the state of the economy is unobservable, optimal control also means optimal filtering, in that you estimate the underlying state of the economy. In this chapter, the natural rate is unobservable. A rational policymaker would conduct optimal policy and also try to estimate the natural rate. In this type of linear model, to estimate the natural rate is to use a Kalman filter and update priors of the natural rate, which might lead to a huge misperception. If inflation takes off, then the modeler should realize they have the wrong idea about the natural rate and the Kalman filter provides the weights one can put on their indicators. Svensson thought this was a problem with much of Orphanides, work, in that he does not take into account whether misperceptions are the best unbiased estimates of the relevant states or not. It might be interesting to see this experiment when the policymaker has the chance to update his or her estimate of the natural rate given the realization of inflation and output, and it might lead to strikingly different results. Svensson also stressed that the techniques to do all of this were known in the 1970s. There was an ambitious optimal control problem being done at the Federal Reserve Board that Chairman Burns did away with. That was being done with backward-looking variables, and Svensson and Michael Woodford have done work on the same problem using forward-looking variables. The algorithm is easily obtainable. What is being done here is an approximation to that problem.

Andrew Levin felt this chapter did not identify what went wrong during the Great Inflation. With regards to 1965 to 1969, Chairman Martin made it clear that his highest concerns were eliminating inflation pressures and keeping orderly financial markets. Levin had a difficult time finding any quotes where Chairman Martin refers to natural unemployment rates or an output gap. The FOMC tightened in 1969 knowing that a recession might be the cost of such tightening. Therefore, for the authors to place a weight of 16 on stabilizing the output gap in their loss function is not plausible from 1965 to 1969. But that was just stage one, and Chairman Martin ended that stage by admitting he felt he failed to control inflation. Stage two was 1976

to 1980. The figure that the authors provide in their chapter shows inflation expectations coming down in 1976, and by 1980 they are around 2.5 to 3 percent. Levin cited his work with John Taylor as evidence that this did not happen, thus it cannot be a natural rate misperception that led to the spike in inflation expectations that was actually realized.

Otmar Issing admitted he has a learned a lot from the work of Orphanides. To him, this chapter left the lasting message that policymakers do not know as much as they used to think they did. Could this happen in 2008, or 2020? Every policymaker should have the warning sign about pretense of knowledge. The more policymakers think they know, the more dangerous their actions become. Take Walter Heller, for example. He was the chairman of a group of economists that came to Germany in the 1950s, a group that made strong recommendations for extremely expansionary policy. Issing admired the German people for respecting the opinions of Heller, yet rejecting them. Heller came back later and admitted he had been wrong and praised German economic policy. It is not that the policymakers were ignoring the recommendations—they were aware of the risks of being wrong.

Vitor Gaspar suggested the Mark Twain quote: "It ain't what you don't know that gets you into trouble. It's what you know for sure that just ain't so." Gaspar felt the authors provided a marvelous illustration of the fundamental insight of Milton Friedman, in that when you design optimal/feasible policy, you have to take account of what you do and cannot possibly know. The simple rule here seems to be a good illustration of just that.

Bennett McCallum felt Levin's comments had merit. Yet even if this study does not really speak to Martin's behavior and is not completely historically correct, it provides a nice cautionary tale. It emphasized the kind of pressures that were there historically and continue to be there. The only reservation McCallum had about the chapter was the estimated model. The trouble is that without the constraints on lagged and expected future variables with weights of 0.5, where do you go? Half of the parameters are not estimated. If you leave them all free, you will not get nearly as nice of results. Can one really take the results seriously then? Would you tell an undergraduate student that this is a full-blown estimated macroeconomic model?

Seppo Honkapohja agreed with Sims's comments on Sargent's literature, and adds that he thinks there is still a misperception element at play. One could also go and say there are multiple models being learned about, and that policymakers chose the wrong model.

Benjamin Friedman paralleled some of Sims's comments, and thought that much of the conversation was remarkably like sermon-speak—that is, there is a certain scripture or certain religious figures that indicate anything that you think is good you attribute as an interpretation of something that religious figure believed in. The notion that writing down an interest rate reaction function with a term in the deviation of inflation from something and a term in the change of the unemployment rate, and deriving the values

of that, does not strike Friedman as something that Milton Friedman advocated. If there are figures in the profession whom we all admire, it does not mean we can attribute any policy we see as good policy to that person.

Williams began the response with comments directed toward Sims and Svensson. In a previous paper with Orphanides, the authors did do policy evaluation using unconditional moments and the standard monetary policy using the optimal Kalman filter. This was optimal policy using rational expectations. When you add in the learning, you show that with learning and the combination of mistakes, policymakers want to put more weight on output in their reaction function. Williams himself admits he does not know the deep parameters of the Kalman filter. In fact, the chapter presented here has a section on the optimal Kalman filter, and the authors admit they want to work on that next.

Orphanides continued the response by emphasizing the criticisms people make about the real-time natural estimates used here. Orphanides strongly believed that these are the estimates that policymakers would have used had they had sophisticated models at the time, referencing previous comments made by Rudd on the methodology being used at the time and how similar it is to what is used now. From the mid-1970s onward, the real-time estimates being used are those the Congressional Budget Office was using. If one tries to do this in a Kalman filter, what are the signal-to-noise ratios you assume? What are the parameters of the model? Orphanides felt the model here is very close to what happened in real time and reflects the best efforts of the profession. It is not a new concept, as was pointed out by Svensson. Optimal control was something the Federal Reserve had been working on, and William Poole probably remembers it. With regards to Friedman's remarks, Orphanides acknowledged he is a great admirer of the work of Milton Friedman, but to interpret Milton Friedman's policy teaching as one that necessarily uses money growth numbers seems like an extremely narrow viewpoint. A much deeper lesson from Orphanides's perspective that we can take from Milton Friedman is that one should always respect the limits of knowledge, the ability of using economic models for stabilization purposes, and the idea of measuring something that is "natural." If you actually estimate these beasts, why do optimal control? To continue, what kind of policy would approximate the lesson of Milton Friedman? Orphanides was convinced it would be some sort of difference rule. This is not different from the kind that Knut Wicksell proposed after he defined what the "natural rate" was. Afterwards comes the question about implementation of such a rule. To respond to McCallum, Orphanides referenced a previous paper in the *Journal of Money, Credit, and Banking* where he and Williams estimated a model with free parameters, and that works quite well in terms of estimation. What the authors wanted to do here was restrict parameters to bring the model as close as possible to modern New Keynesian models that respects the limits that the forward-backward hybrid combination can have

on the parameters. If we do not hold the weights at 0.5, then the estimated parameters are far out of bounds. This worried McCallum, but Orphanides thought that the New Keynesian model was a good tool. Orphanides rejected Levin's views on the attempt to reconstruct history, which is not what the authors were trying to do. The authors do not believe chairmen Martin and Burns followed the recommendations of the Federal Reserve Board staff at the time. What the authors attempted to do is ask the following question: Suppose they did follow the recommendation of an optimal control policy using modern techniques—would they have avoided the big mistakes? The answer this chapter provides is no, because of the misperceptions. The basic lesson is that because of learning and misperceptions that result in over-expansionary policy, one would not be able to control inflation. It is key that in order to understand the Great Inflation, one needs to realize the economic profession's obsession with the ability to stabilize both the real side of the economy and inflation. Even with the very best models, we have failed. In other papers, Orphanides and Williams have tried to match history, but misperceptions always create mistakes. In reference to Sims and constant gain learning, the authors used constant gain learning as a parable for how actual people in the economy, as opposed to the hypothesized rational expectations agents, might be using past data to form expectations. The authors tried to capture the evolution of the expectations formation process that results from suspicion that things might be changing. Orphanides admitted that they could have time variation in parameters, and then they would have a Kalman filter giving an optimal way to shift the parameters. In the end, the authors wanted to approximate James Stock and Mark Watson. They are good at forecasting, and how do they change their vector autoregressions?