# Real-Time Model Uncertainty in the United States: "Robust" Policies Put to the Test*

Robert J. Tetlow
Division of Monetary Affairs
Federal Reserve Board

I study forty-six vintages of FRB/US, the principal macro model used by the Federal Reserve, as measures of real-time model uncertainty and examine the robustness of commonly applied, simple monetary policy rules. Model uncertainty turns out to be a substantial problem: key model properties differ in important ways across model vintages, as do the optimized parameterizations of candidate rules. Among the simple monetary policy rules considered are rules that eschew feedback on the output gap, rules that target nominal income growth, and rules that allow for time variation in the equilibrium real interest rate. Many rules that previous research has shown to be robust in artificial economies would have failed to provide adequate stabilization in the real-time, real-world environment seen by the Federal Reserve staff. I identify certain policy rules that would have performed relatively well, and characterize their key features to draw more general lessons about the design of monetary policy under model uncertainty.

JEL Codes: E37, E5, C5, C6.

---

[*] The working paper version of this paper and others can be found at http://www.roberttetlow.com. I thank Flint Brayton and Dave Reifschneider for helping me interpret historical model changes and for valuable discussions. I thank Bill Branch, Ben Friedman, and an anonymous referee for helpful comments, and Sean Taylor and Trevor Davis for dedicated research assistance. All remaining errors are mine. The views expressed in this paper are those of the author alone and do not necessarily represent those of the Federal Reserve Board or other members of its staff. Author contact: Robert Tetlow, Federal Reserve Board, Washington, D.C. 20551. E-mail: rtetlow@frb.gov.

## 1. Introduction

> "We have involved ourselves in a colossal muddle, having blundered in the control of a delicate machine, the working of which we do not understand."
>
> — John Maynard Keynes, *The Great Slump of 1930* (December 1930)

Over the twenty years since the Taylor (1993) rule was introduced, there has been an explosion of work studying the characteristics of monetary policy rules in general and simple, interest rate feedback rules in particular. While considerable insight has come out of this literature, so has a fundamental critique, namely that results formulated in this way may not be robust to misspecification of the underlying model. Keynes's metaphor of the economy as a "delicate machine ... which we do not understand" seems as apt today as it was in 1930.

It follows from that lack of understanding that a principal concern for policymakers is uncertainty, and how to deal with it. The fast-growing literature on model uncertainty seeks answers to this question; see, inter alia, Levin, Wieland, and Williams (1999), Tetlow and von zur Muehlen (2001), Onatski and Williams (2003), Levin et al. (2006), Brock, Durlauf, and West (2007), and Taylor and Wieland (2012).[1] This strand often employs the *rival models* method of analysis wherein the researcher posits two or more alternative models of the economy and employs statistical or decision-theoretic techniques to find a policy rule that performs "well" in each of the posited models; see, e.g., McCallum (1988). While this approach has produced interesting and useful results, it is hampered by the artificiality of the environment in which it has typically been employed. In nearly all cases, the models under consideration are highly abstract and do not fit the data well, useful perhaps for making narrow

---

[1] Two other aspects of uncertainty, relevant to monetary policymaking, are *parameter uncertainty* (see, e.g., Brainard 1967, Söderström 2002, Walsh 2004, and Kimura and Kurozumi 2007) and *data uncertainty* (Aoki 2003, Jääskelä and Yates 2005). These subject areas should be regarded as complementary to the study of model uncertainty.

points, but not to be taken seriously as tools for monetary policy design.[2]

Virtually absent from the above characterization of the literature is the real-time analysis of model uncertainty. At one level, this is not surprising; after all, while it is easy to conceptualize changing views about what the true model might be, it is more difficult to imagine the laboratory in which such an analysis could be conducted. That is, however, exactly what this paper provides. My laboratory is the Federal Reserve Board staff and the FRB/US model. I examine time variation in model properties, and hence model uncertainty, as it was seen *in real time* by the Federal Reserve Board staff. I do this using forty-six vintages of the Board staff's FRB/US model, or four per year, that were actually used for forecasting and policy analysis during the period from July 1996 to October 2007, examining how the model specification, coefficients, databases, and stochastic shock sets changed from vintage to vintage as new and revised data came in. The advantage provided is that I can focus on those aspects of model uncertainty that are germane to policy decisions, using a model that is used to formulate advice for those decisions.

The relevance of the model is unquestionable: since its introduction in July 1996, the FRB/US model has been used continuously for communicating ideas to the Board of Governors and the Federal Open Market Committee (FOMC). All of the staff's alternative scenarios focusing on domestic economic issues during this period were conducted using the model; forecast confidence intervals are computed using FRB/US, as are optimal policy exercises that are presented to the FOMC. In his 1998 monograph on his time as Vice Chairman of the Federal Reserve Board, Alan Blinder notes (p. 12) the important role that FRB/US simulations played in guiding his thinking; and Blinder and Yellen (2001) made extensive use of simulations of the FRB/US models to explain the boom in the U.S. economy in the 1990s.

---

[2] An illuminating exception to this rule is the paper of Levin et al. (2006), which uses an estimated DSGE model. Comments on this paper by Walsh (2006) express doubts that the current generation of DSGE models is sufficiently advanced to be taken seriously for this purpose.

Armed with these forty-six vintages of the model, I ask whether the policy rules that have been promoted as robust in one environment or another are in fact robust in this real-world context.

The first policy rule I study, and the one that serves as my benchmark for comparison, is the familiar Taylor (1993) rule, although I use parameterizations that are optimal for the model vintages that are interesting. I also consider three rules that take up the argument of Orphanides (2001), among others, that the inherent difficulty in conditioning policy on unobservable constructed variables like output gaps means that policy should eschew feedback on latent variables altogether. Two candidate rules follow Bennett McCallum (1988) by keying off of nominal output growth. A nominal output growth rule establishes a nominal anchor but unlike, say, an inflation-targeting rule, makes no explicit call on whether shocks are real or nominal; because of this, it is arguably less susceptible to misspecification of the supply side and of the incidence of supply shocks. Two rules build off of the finding of Levin et al. (2006) to the effect that policy should respond to nominal wage inflation instead of price inflation. In this way, the policymaker pays particular attention to the labor market, arguably the part of the economy that, from a neoclassical perspective, is the most distorted.

This paper goes a number of steps beyond previous contributions to the literature. As already noted, it goes beyond the extant rival models literature through its novel and efficacious focus on models that are actually used in a policy environment. It also goes beyond the literature on parameter uncertainty. That literature assumes that parameters are random but the model is fixed over time: misspecification is simply a matter of sampling error. Model uncertainty is a thornier problem, in large part because it often does not readily lend itself to statistical methods of analysis. I explicitly allow the models to change over time in response not just to the data but to the economic issues of the day.[3] Finally, as already noted, this paper

---

[3] There have been a number of valuable contributions to the real-time analysis of monetary policy issues. Most are associated with data and forecasting. See, in particular, the work of Croushore and Stark (2001) and a whole conference on the subject, details of which can be found at http://www.phil.frb.org/econ/conf/rtdconfpapers.html. An additional, deeper layer of real-time analysis considers revisions to unobservable state variables, such as potential output; Athanasios Orphanides, alone or with co-authors, has

does all this within a class of possible models that is undeniably realistic.

The analysis presented herein is, of course, based on the U.S. economy and the FRB/US model. However, uncertainty, in its various forms, is of concern for monetary authorities the world over as it is for other decision makers. Real-time data issues and data uncertainty more generally have garnered a great deal of attention in the United Kingdom; see, e.g., Garratt and Vahey (2006) and Garratt et al. (2007) and references therein. In the euro area, Giannone, Reichlin, and Sala (2005) study real-time uncertainty for its implications for monetary policy.[4]

The rest of this paper proceeds as follows. The second section begins with a very brief discussion of the FRB/US model in generic terms, and the model's historical archives. The third section compares model properties by vintage. To do this, I document changes in real-time "model multipliers" and compare them with their ex post counterparts. The fourth section computes optimized Taylor-type rules and compares these to commonly accepted alternative policies in a stochastic environment. The fifth section examines the stochastic performance of candidate rules for two selected vintages, the December 1998 and October 2007 models. A sixth and final section sums up and concludes.

## 2. Forty-Six Vintages of the FRB/US Model

The FRB/US model came into production in July 1996 as a replacement for the venerable MIT-Penn-SSRC (MPS) model. The main objectives guiding the development of the model were that it be useful for both forecasting and policy analysis; that expectations

---

been at the vanguard of this issue; see, e.g., Orphanides et al. (2000). See also Giannone, Reichlin, and Sala (2005) for a sophisticated, real-time analysis of the history of FOMC behavior.

[4] With regard to fiscal policy, Cimadomo (2008) and Giuliodori and Beetsma (2008) uncover important implications of real-time data uncertainty and, indirectly, model uncertainty for the measurement of fiscal stance and the conduct of fiscal policy for the OECD countries and the euro area, respectively. Whole conferences have been organized just on the need for real-time data for the euro area; e.g., the Center for Economic Policy Research conference "Needed: A Real Time Database for the Euro-Area," June 13–14, 2005, in Brussels (http://www.cepr.org/meets/wkcn/1/1632/papers/).

be explicit; that important equations represent the decision rules of optimizing agents; that the model be estimated and have satisfactory statistical properties; and that the full-model simulation properties match the "established rules of thumb regarding economic relationships under appropriate circumstances," as Brayton and Tinsley (1996, p. 2) put it. To address these challenges, the staff included within the model's structure an expectations block, and with it, a fundamental distinction between intrinsic model dynamics and expectational dynamics. Two versions of expectations formation were envisioned: VAR-based expectations and perfect foresight, although mixtures of those two classes of expectations formation can and have been implemented.

The key features influencing the monetary policy transmission mechanism in the FRB/US model are the effects of changes in the federal funds rate on asset prices and from there to expenditures. From this, long-term real interest rates are determined, which in turn affect stock prices (and hence private wealth) and exchange rates. The model's wage-price block has always shared the same basic features: sticky wages and prices, expected future excess demand in the goods and labor markets that influences price and wage setting, and a channel through which productivity affects real and nominal wages. That said, as we shall see, there have been substantial changes over time in both (what we may call) the interest elasticity of aggregate demand and the effect of excess demand on inflation.

In this paper, I will be working exclusively with the VAR-based expectations version of the model. Typically it is the multipliers of this version of the model that are reported to Board members when they ask "what-if" questions. This is the version that is used for forecasting and most of the policy analysis by the Federal Reserve staff, including, as Svensson and Tetlow (2005) demonstrate, policy-optimization experiments. Thus, the pertinence of using this version of the model for the question at hand is unquestionable. What could be questioned, on standard Lucas critique grounds, is the validity of the simple-rule optimizations, given that expectations are not fully rational. However, the period under study is one entirely under the leadership of a single Chairman, and I am aware of no evidence to suggest that there was a change in regime during this period. So, as Sims and Zha (2006) have argued, it seems likely that the perturbations to policies encompassed by the range of policies studied below are not large enough to induce a change in expectations formation.[5] Moreover, in an environment such as the one under study, where changes in the non-monetary part of the economy are likely to dwarf the monetary policy perturbations, it seems safe to assume that private agents were no more rational with regard to their anticipations of policy than the Federal Reserve staff was about private-sector decision making.[6] In their study of the evolution of the Federal Reserve beliefs over a longer period of time, Romer and Romer (2002) ascribe no role to the idea of rational expectations. Moreover, Rudebusch (2002) shows that issues of model uncertainty are often of second-order importance in linear rational expectations models. Thus the VAR-based expectations case is arguably the more quantitatively interesting one.

There is not the space here for a complete description of the model. The working paper version of this article contains a more generous description of model features and a more detailed characterization of economic history over the period under study with reference to how that history affected the construction of various model vintages; see also Brayton and Tinsley (1996), Brayton et al. (1997), Reifschneider, Tetlow, and Williams (1999), and Tetlow and Ironside (2007).

Since its inception in July 1996, the FRB/US model code, the equation coefficients, the baseline forecast database, and the list of stochastic shocks with which the model would be stochastically simulated have all been stored in model archives, one for each of the eight forecasts the Board staff conducts every year. Because it is releases of National Income and Product Accounts (NIPA) data that typically induce reassessments of the model, I use four archives per year, or forty-six in total, the ones immediately following NIPA preliminary releases.[7]

---

[5] The model's VAR-based expectations code does allow for shifting in the public's perceptions of the long-run real federal funds rate and the target rate of inflation, but not in the elasticities conditional on those beliefs. See Brayton and Tinsley (1996) for details.

[6] A complete set of rational expectations vintages of the FRB/US model does not exist and, in any case, working with those models is computationally infeasible.

[7] The archives are listed by the precise date of the FOMC meeting in which the forecasts were discussed. For present purposes, such precision is not necessary, so

In what follows, I experiment with each vintage of the model, comparing their properties in selected experiments. Consistent with the real-time philosophy of this endeavor, the experiments I choose are typical of those used to assess models by policy institutions in general and the Federal Reserve Board in particular. They fall into two broad classes. One set of experiments, *model multipliers*, attempts to isolate the behavior of particular parts of the model. A multiplier is the response of a key endogenous variable to an exogenous shock after a fixed period of time. An example is the response of the level of output after eight quarters to a persistent increase in the federal funds rate. The other set of experiments judge the stochastic performance of the model and are designed to capture the full-model properties under fairly general conditions. So, for example, I will compute by stochastic simulation the optimal coefficients and economic performance of simple rules, conditional on a model vintage, a baseline database, and a set of stochastic shocks.[8]

Model multipliers have been routinely reported to, and used by, members of the FOMC. Indeed, the model's sacrifice ratio—about which I will have more to say below—was used in the very first FOMC meeting following the model's introduction. Similarly, model simulations of alternative policies have been carried out and reported to the FOMC in a number of memos and official FOMC documents.

The archives document model changes and provide a unique record of model uncertainty. As we shall see, the answers to questions a policymaker might ask differ depending on the vintage of the model. The seemingly generic issue of the output cost of bringing down inflation, for example, can be subdivided into several more precise questions, including (i) what would the model say is the output cost of bringing down inflation today? (ii) what would the model

---

I shall describe them by month and year. Thus, the forty-six vintages I use are, in 1996, July and November; then, typically thereafter the months would be January (but often February), May, August (but occasionally July), and November (but twice October and once December). Nothing of importance is lost from the analysis by excluding every second vintage from consideration.

[8] Each vintage has a set of variables that are shocked for stochastic simulations using bootstrap methods. The list of shocks is a subset of the model's complete set of residuals because other residuals are treated not as shocks but as measurement error. The set of shocks varies with the model vintage according to vintage-specific variable data construction and the period over which the shocks are drawn.

of today say the output cost of bringing down inflation would have been in December 1998? and (iii) what would the model have said in December 1998 was the output cost of disinflation at that time? These questions introduce a time dependency to the issue that rarely appears in other contexts.

## 3. Model Multipliers in Real Time and Ex Post

In this section, I consider the variation in real time of selected model multipliers. In the interest of brevity, I devote space to just three multipliers. The first is the sacrifice ratio; that is, the cumulative annualized cost measured in terms of increased unemployment over five years of permanently reducing the inflation rate by 1 percentage point. The second is the funds rate multiplier, defined here as the percentage change in the level of real output after eight quarters that is induced by a persistent 100-basis-point increase in the nominal federal funds rate.[9] In the parlance of an undergraduate textbook closed-economy model, these two multipliers represent the slope of the Phillips curve and the slope of the aggregate demand curve, respectively. To add an international element, I add an exchange rate multiplier—specifically, the percentage change in real GDP associated with a 10 percent appreciation of the trade-weighted exchange value of the U.S. dollar. The sacrifice ratio is the outcome of a five-year simulation experiment; the other multipliers are measured in terms of their effects after eight quarters.

It is easiest to show the results graphically. But before turning to specifics, it is useful to outline how these figures are constructed and how they should be interpreted. In all cases, I show two lines. The solid line is the real-time multiplier by vintage. Each point on the line represents the outcome of the same experiment, conducted on the model vintage of that date, using the baseline database at that point in history. Thus at each point shown by the solid line, the model, its coefficients, and the baseline all differ. The dashed line shows what I call the ex post multiplier. The ex post multiplier is computed using the most recent model vintage for each date; the

---

[9] These multipliers could have been defined differently; however, the qualitative conclusions would be no different for any reasonable alternative.

**Figure 1. Real-Time and Ex Post Sacrifice Ratios, by Model Vintage**

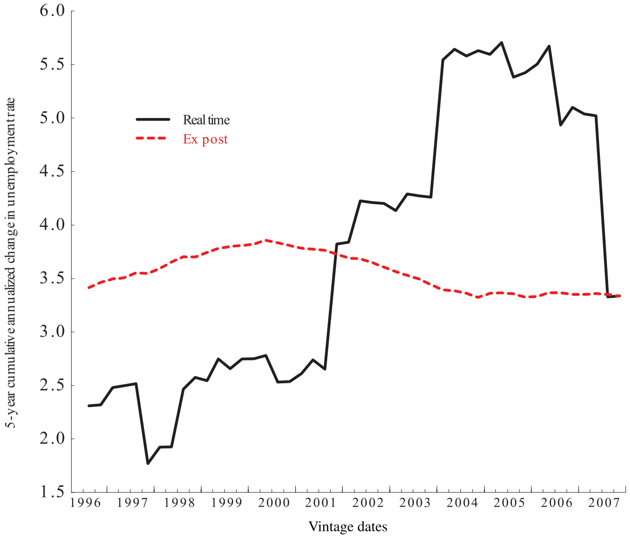

only thing that changes for each point on the dashed line is the initial conditions under which the experiment is conducted. Because the multipliers for linear models are independent of initial conditions, examining the dashed line gives an indication of how much of the time variation shown in the solid line arises from the initial conditions and associated model non-linearities.

Now let us look at figure 1, which shows the sacrifice ratio.[10] Let us focus on the dashed line first. It shows that for the October 2007 model, the sacrifice ratio is essentially constant over time. So if the staff were asked to assess the sacrifice ratio, or what the sacrifice ratio would have been in, say, December 1998, the answer based on the October 2007 model would be the same: about 3-1/4, meaning it would take that many percentage-point years of unemployment

---

[10] The experiment is conducted by simulation, setting the target rate of inflation in a Taylor rule to 1 percentage point below its baseline level. The sacrifice ratio is the cumulative annualized change in the unemployment rate, undiscounted, relative to baseline, divided by the change in personal consumption expenditure (PCE) inflation after five years. Other rules may produce somewhat different sacrifice ratios but the same profile over time.

to bring down inflation by 1 percentage point. Now, however, look at the solid line. Since each point on the line represents a different model, and the last point on the far right of the line is the October 2007 model, the dashed line and the solid line must meet at the right-hand side in this and all other figures in this section. But notice how much the real-time sacrifice ratio has changed over the twelve-year period of study. Had the model builders been asked in December 1998 what the sacrifice ratio was, the answer based on the February 1997 model would have been about 2-1/2. Prior to a revision in mid-2007 that was undertaken expressly, in larger part, to reduce it, the sacrifice ratio for vintages from 2004 to 2006 was of the order of 5-1/2, or more than double what it was in the 1990s.

The sacrifice ratio is a crucial statistic for any central bank model. On the one hand, it describes the cost of bringing down inflation, given that one inherits a higher inflation rate than is desired because of, say, a supply shock. From this perspective, a high sacrifice ratio is a bad thing. On the other hand, a high sacrifice ratio reflects a flat Phillips curve; that is, shocks to aggregate demand of a given magnitude will elicit smaller changes in inflation than otherwise. From this perspective, a high sacrifice ratio is a good thing. Which effect dominates depends on the incidence of supply and demand shocks.

The primacy of the model's sacrifice ratio to policy debates is clear from FOMC transcripts. It was, for example, a topic of discussion at the first FOMC meeting following the introduction of the model.[11] Similarly, the February 1, 2000 meeting of the FOMC produced this exchange between Federal Reserve Bank of Minneapolis President Gary Stern and then FOMC Secretary (subsequently Federal Reserve Board Vice Chairman) Donald Kohn:[12]

> Mr. Stern: Let me ask about the Bluebook [FRB/US model simulation] sacrifice ratio. I don't know what your credibility assumption is, but it seems really high.

---

[11] See http://www.federalreserve.gov/monetarypolicy/files/FOMC19960703meeting.pdf for a transcript of the July 2–3, 1996 meeting of the FOMC, pp. 42–47.

[12] Transcript, FOMC meeting, February 1, 2000, pp. 41–2. Available at http://www.federalreserve.gov/monetarypolicy/files/FOMC20000202meeting.pdf.

> Mr. Kohn: It is a little higher than we've had in the past, but not much. It is consistent with the model looking out over the longer run. It is a fairly high sacrifice ratio, I think, compared to some other models, but it is not out of the bounds.

Kohn was clearly aware that the model's sacrifice ratio had undergone some change and was rightfully cognizant of how it compares against alternative models. As it happens, the increases already incurred in the sacrifice ratio were only the beginning.

The climb in the model sacrifice ratio is striking, particularly as it was incurred over such a short period of time among model vintages with substantial overlap in their estimation periods. Of particular note is the sizable jump in the sacrifice ratio in late 2001 which arose when the staff began estimating the wage and price equations simultaneously, together with other equations to represent the rest of the economy, including a Taylor rule for policy. One might be forgiven for thinking that this phenomenon is idiosyncratic to the model under study. But other work shows that this result is not a fluke.[13] At the same time, as I have already noted, the model builders *did* incorporate shifts in the NAIRU (and in potential output), but found that leaning exclusively on this one story for macroeconomic dynamics in the late 1990s was insufficient. Thus, the revealed view of the model builders contrasts with the idea advanced by Staiger, Stock, and Watson (2001), among others, that changes in the Phillips curve are best accounted for entirely by shifts in the natural rate of unemployment. Toward the end of the decade, a reduction in the sacrifice ratio became an important objective of the specification and estimation of the model's wage-price block; success on this front was achieved through respecification of how long-term inflation expectations evolve over time.

---

[13] In particular, the same phenomenon occurs to varying degrees in simple single-equation Phillips curves of various specifications using both real-time and ex post data. One paper along these lines is Atkeson and Ohanian (2001). Primiceri (2005) estimates a time-varying VAR model to uncover substantial shifts in the time-series behavior of inflation. Cogley and Sargent (2005) estimate three Phillips-curve models simultaneously and apply Bayesian decision theory to explain why the Federal Reserve did not choose an inflation-stabilizing policy before the Volcker disinflation; that paper also finds substantial time variation in the output cost of disinflation.

Figure 2. Funds Rate Multipliers, by Model Vintage

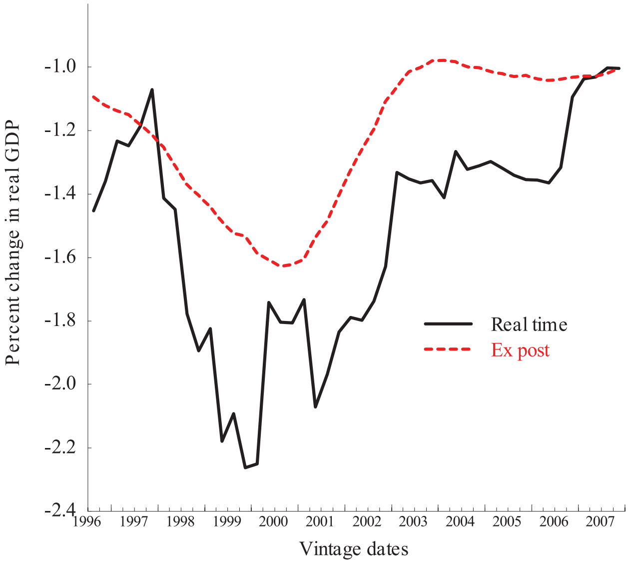

Figure 2 shows the funds rate multiplier; that is, the percentage decrease in the level of real GDP after eight quarters in response to a persistent 100-basis-point increase in the funds rate. In this instance, the dashed line shows important time variation: the ex post funds rate multiplier varies with initial conditions; it is largest, in absolute terms, at about 1.6 percentage point in late 2000, and lowest at the beginning and at the end of the period, at about 1 percent. The non-linearity stems entirely from the specification of the model's stock market equation, which is written in levels, rather than in logs, a feature that makes the interest elasticity of aggregate demand an increasing function of the ratio of stock market wealth to total wealth. The mechanism is that an increase in the funds rate raises long-term bond rates, which in turn bring about a drop in stock market valuation operating through the arbitrage relationship between expected risk-adjusted bond and equity returns. The larger the stock market, the stronger the effect.[14]

---

[14] The levels relationship of the stock market equation means that the wealth effect of the stock market on consumption can be measured in the familiar "cents

The real-time multiplier, shown by the solid line, is harder to characterize. Two observations stand out. The first is the sheer volatility of the real-time multiplier. In a large-scale model such as the FRB/US model, where the transmission of monetary policy operates through a number of channels, time variation in the interest elasticity of aggregate demand depends on a variety of parameters. Second, the real-time multiplier is almost always smaller than the ex post multiplier. The gap between the two is particularly marked in 2000, when the business cycle reached a peak, as did stock prices. At the time, concerns about possible stock market bubbles were rampant. One aspect of the debate between proponents and detractors of the active approach to stock market bubbles concerns the feasibility of policy prescriptions in a world of model uncertainty.[15] The considerable difference between the real-time and ex post multipliers during this period demonstrates the difficulty in carrying out historical analyses of the role of monetary policy; today's assessment of the strength of those monetary policy actions can differ substantially from what the FRB/US model implied in real time.

The final multiplier covered here is the effect of a sustained 10 percent appreciation of the real exchange value of the U.S. dollar on real output in the United States. The striking change in 1998 in figure 3 corresponds with a shift from a G10 aggregate of trade weights for foreign indexes to a G29 aggregate. The subsequent reversal began with a shift to chain weighting of domestic price indexes in 1999:Q3. In any case, without belaboring the details, the salient fact to take from this figure and others like it that are not shown here is first and foremost the variability of the elasticities.[16]

To summarize this section, real-time multipliers show substantial variation over time, and differ considerably from what one would

---

per dollar" form (of incremental stock market wealth). Also playing a role is the log-linearity (that is, constant elasticity) of the relationship between wealth and consumption.

[15] The "active approach" to the presence of stock market bubbles argues that monetary policy should specifically respond to bubbles. See, e.g., Cecchetti et al. (2000). The passive approach argues that bubbles should affect monetary policy only insofar as they affect the forecast for inflation and possibly output. They should not be a special object of policy. See Bernanke and Gertler (1999).

[16] The working paper version of this article shows and discusses another open-economy multiplier, namely the pass-through non-oil import prices into PCE price inflation. See also Tetlow and Ironside (2007).

**Figure 3. Real Exchange Rate Multipliers, by Model Vintage**

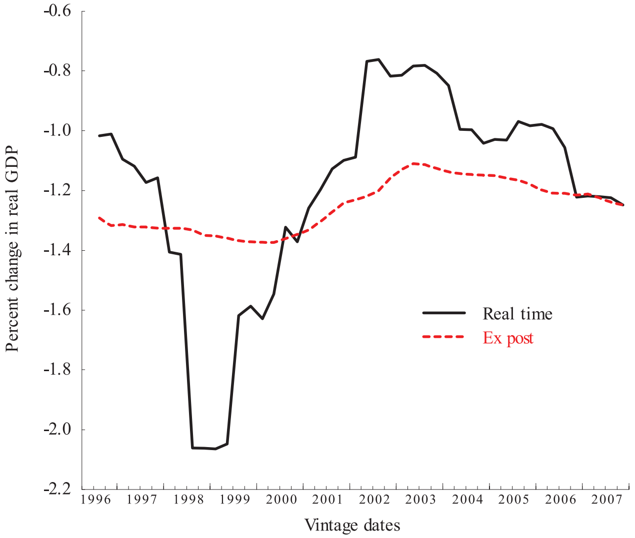

say ex post the multipliers would be. Moreover, the discrepancies between the two multiplier concepts have often been large at critical junctures over the period under study from 1996 to 2007. It follows that real-time model uncertainty is an important problem for policymakers. The next section quantifies this point by characterizing optimized policies, and their time variation, conditional on these model vintages.

## 4. Monetary Policy in Real Time

### *4.1 The Rules*

I seek rules that are simple, robust, and effective. In the current context, a monetary policy rule can be described as *robust* if (i) the optimized policy coefficients do not differ in important ways across models *or* (ii) the performance of the economy does not depend in an economically important way on rule parameterization. A robust policy rule can also be described as *effective* if (iii) it performs “well,” relative to some benchmark policy rule.

A popular simple monetary policy rule is the canonical Taylor (1993) rule. One reason the Taylor rule is advocated for monetary policy is because of its simplicity in that it calls for feedback on only those variables that are central in a broad class of models; see, e.g., Williams (2003) for an argument along these lines. And indeed many central banks use Taylor rules, or rules much like them, in the assessment of monetary policy and for formulating policy advice, including in the Board staff's presentations of policy options for the FOMC. In the U.S. case, Giannone, Reichlin, and Sala (2005) show that the good fit of simple two-parameter Taylor-type rules can be attributed to the small number of fundamental factors driving the U.S. economy; that is, the two arguments that appear in Taylor rules encompass all that one needs to know to summarize monetary policy in history. In what follows, I will use the Taylor rule, appropriately parameterized, as my benchmark for comparison.

Taylor rules have their detractors. Much of the earlier work on robust policy rules has focused on the importance of estimation and misperception of potential output and the associated mismeasurement of the output gap.[17] Accordingly, some of the rules I consider are those that have been suggested as prophylactics for this problem. In other instances, it is a broader class of latent variables that have been the object of concern. For example, the productivity boom in the United States in the second half of the 1990s brought about misperceptions not just of the *level* of the output gap but also of potential output *growth* going ahead; these concepts in turn have a bearing on the equilibrium real interest rate since in all but the smallest of open economies, the equilibrium real interest rate is determined, in part, by the steady-state growth rate of the economy. The two problems are related but different. Mismeasurement of the *level* of potential output, by itself, induces only a stationary error process. Missing a shift in the *growth rate* of potential represents a more persistent error that affects a wider range of variables, including the equilibrium real interest rate. Accordingly, some of the rules I consider below are intended to address the latter, more complicated, problem.

---

[17] See, e.g., Orphanides et. al. (2000), Orpanides (2001), and Ehrmann and Smets (2003).

The simple rules employed for this paper are summarized in table 1. In this section, I provide brief descriptions of their features. Most of the analysis is restricted to a class of optimized two-parameter policy rules. This keeps the rules on equal footing in that it is to be expected that adding extra optimized parameters should improve performance, at least for a given model. It also keeps computational costs to a manageable level. Nevertheless, as a check against possible idiosyncrasies in results, I do consider a limited number of three-parameter specifications.

#### *4.1.1 Two-Parameter Policy Rules*

The first rule is the familiar Taylor rule which, for short, I will often refer to as "TR." As indicated in the first row of table 1, and in the accompanying notes to the table, the rendition used here differs slightly from Taylor (1993) in that I allow the coefficients to differ depending on the simple-rule optimization discussed below, and I permit the "equilibrium" real rate, $rr^*$, to vary over time rather than fix it to the constant value of 2 used in the original. Also, in keeping with FOMC practice, the inflation rate used in the rule is the (four-quarter) PCE inflation rate rather the GDP deflator.

In my first bow to the output-gap mismeasurement problem, I also study an inflation-targeting rule (IT); that is, a rule that eschews feedback on the output gap altogether in order to avoid problems from the sort of data and conceptual revisions described in section 2 above, as suggested by Orphanides (2001). For this rule and several others, I allow for instrument smoothing, with the parameter $\alpha_r$, and allow the term $(1 - \alpha_r)(.)$ to pick up the steady-state level of the real interest rate.[18] In addition to IT, I investigate a price-level targeting counterpart of the same specification, PLT, where it should be understood that $p_t^*$ need not be a fixed number; it could instead be (and, for us, is) a predetermined trending path for the

---

[18] In nearly all works on optimized rules, the steady-state terms are omitted for two reasons: first, the models used are linear, so the steady state can be taken as zero; and second, no allowance is made for shifting steady states. (An exception is Orphanides and Williams 2002, who specifically consider $rr^*$ that shift over time.) Because I am using real models with real databases, and I am considering persistent deviations from steady state—indeed arguably this is a large part of the problem of interest—I need to retain these steady-state terms.

## Table 1. Summary of Simple Rules

| Name | Rule | Ref. |
|---|---|---|
| | *Two-Parameter Rules* | |
| Taylor Rule | $r_t = rr_t^* + \widetilde{\pi}_t + \alpha_Y(y_t - y_t^*) + \alpha_\pi(\widetilde{\pi}_t - \pi_t^*)$ | TR |
| Inflation | $r_t = \alpha_r r_{t-1} + (1 - \alpha_r)(rr_t^* + \widetilde{\pi}_t) + \alpha_\pi(\widetilde{\pi}_t - \pi_t^*)$ | IT |
| Price Level | $r_t = \alpha_r r_{t-1} + (1 - \alpha_r)(rr_t^* + \widetilde{\pi}_t) + \alpha_\pi(p_t - p_t^*)$ | PLT |
| Change in U-Rate | $\Delta r_t = \alpha_\pi(\widetilde{\pi}_t - \pi_t^*) + \alpha_{\Delta u}\Delta u_t$ | DUR |
| Nominal Output 1 | $r_t = \alpha_r r_{t-1} + (1 - \alpha_r)(rr_t^* + \widetilde{\pi}_t) + \alpha_{\Delta yn}(\Delta\widetilde{yn}_t - \Delta yn_t^*)$ | YN1 |
| Nominal Output 2 | $r_t = (rr_t^* + \widetilde{\pi}_t) + \alpha_Y(y_t - y_t^*) + \alpha_{\Delta yn}(\Delta\widetilde{yn}_t - \Delta yn_t^*)$ | YN2 |
| Wage Growth 1 | $r_t = \alpha_r r_{t-1} + (1 - \alpha_r)(rr_t^* + \widetilde{\pi}_t) + \alpha_{\Delta w}(\Delta\widetilde{w}_t - \Delta w_t^*)$ | WN1 |
| Wage Growth 2 | $r_t = rr_t^* + \widetilde{\pi}_t + \alpha_Y(y_t - y_t^*) + \alpha_{\Delta w}(\Delta\widetilde{w}_t - \Delta w_t^*)$ | WN2 |
| | *Three-Parameter Rules* | |
| Inertial Taylor | $r_t = \alpha_r r_{t-1} + (1 - \alpha_r)(rr_t^* + \widetilde{\pi}_t) + \alpha_Y(y_t - y_t^*) + \alpha_\pi(\widetilde{\pi}_t - \pi_t^*)$ | TRI |
| Potential Growth | $r_t = rr_t^* + \widetilde{\pi}_t + \alpha_{\Delta Y*}\Delta y^* + \alpha_Y(y_t - y_t^*) + \alpha_\pi(\widetilde{\pi}_t - \pi_t^*)$ | DY* |
| Inertial YN | $r_t = \alpha_r r_{t-1} + (1 - \alpha_r)(rr_t^* + \widetilde{\pi}_t) + \alpha_Y(y_t - y_t^*) + \alpha_{\Delta yn}(\Delta\widetilde{yn}_t - \Delta yn_t^*)$. | YNI |

**Notes:** $r$ is the nominal federal funds rate; $rr^*$ is the "equilibrium" real rate; $\pi$ is PCE inflation; $p$ is the price level; $y$ is real output; $w$ is the nominal wage rate; $u$ is the unemployment rate; $yn$ is nominal output; a $\sim$ overstrike indicates a four-quarter average; an * superscript indicates a target or equilibrium level of a variable. All variables not expressed as rates are in logs.

(log of the) price level such that successful targeting renders a positive average rate of inflation. The important distinction between a price-level target and an inflation target is that in the event of an inflation surprise, a price-level targeting regime is obliged not just to bring inflation back down to the target level but to bring inflation below target for a time in order to return the price level to its target path.

I will also analyze a Taylor-type rule that substitutes the change in the unemployment rate for the traditional output gap in order to allow a real variable to enter the rule while still minimizing the effects of misperceptions of potential output; see, e.g., Orphanides and Williams (2002). Notice that this rule, designated "DUR," is written in the first difference of the funds rate, a configuration that eliminates the need to condition on the equilibrium real interest rate. As such, the DUR takes a step towards insulation against persistent shocks to productivity and associated mismeasurements of $rr^*$.

Another much-touted rule is the nominal output growth rule, along the lines suggested by Bennett McCallum (1988) and revisited by Dennis (2001) and Rudebusch (2002). Its purported advantage is that it parsimoniously includes both prices and real output *growth* but without taking a stand on the split between the two; for this reason it is said to be able to withstand productivity shocks. Detractors note that because output typically leads inflation, responding to the sum of the two is not as obviously beneficial as presumed and that fluctuations in productivity growth will imply that the inflation rate is not pinned down, even in the long run. I experiment with two versions. The first is designated with the rubric "YN1," follows the formulation of McCallum and Nelson (1999), and nests the versions studied by Rudebusch (2002). However, because YN1 embodies output growth within its specification, albeit with its coefficient restricted to equal that of GDP price inflation, but not a term for the *level* of resources utilization, I augment the analysis by including a second version, YN2, which has the virtue of being identical to TR except that nominal output growth substitutes for inflation.

I also pick up on the finding of Levin et al. (2006) that a policy that responds to nominal wage inflation (WN1) instead of nominal price inflation performs well across a range of microfounded models.

In this way, the policymaker pays particular attention to that part of the economy that, from a neoclassical perspective, is arguably the most distorted. Like the nominal output growth targeting rule, because wage setting is supposed to reflect both price inflation and labor productivity, the nominal wage growth rule also has the merit of implicitly incorporating changes in trend productivity. Finally, to close the summary of two-parameter rules, in parallel fashion to my nominal output rules, I consider a second version of the nominal wage growth rule, WN2, that replaces the lagged instrument with the output gap.

#### *4.1.2 Three-Parameter Policy Rules*

As noted, the benchmark against which all my rules are to be compared is the optimized version of the Taylor rule. It is possible, however, that the two-parameter Taylor rule is too parsimonious to adequately respond to the myriad economic disturbances to which the economy is subjected. In order to address this possibility, I also explore an inertial Taylor rule—let us call it "TRI." This rule is the most commonly studied extension on the static Taylor rule; Williams (2003) argues that the inclusion of the lagged instrument can provide significant benefits in terms of economic outcomes in linearized New Keynesian models; see also English, López-Salido, and Tetlow (2013).

Lastly, acknowledging that persistent shocks to productivity may shift the equilibrium real interest rate, it seems prudent to consider conditioning policy specifically on potential output growth. At the same time, to be realistic, one should use not ex post measures of potential growth but rather the estimates that modelers were working with in real time. This can be done using the rule shown in the penultimate line of the table, called the potential growth rule (DY*), where $\Delta y^*$ is the *vintage-consistent estimate* of potential output growth. The terms $rr^*$ and $\alpha_{Y*}\Delta y^*$ together can be taken as a reworked estimate of the equilibrium real rate, one that corrects for changes in potential output growth.

Together, these rules encompass a broad range of the rules that have been proposed as robust to model misspecification, and do so in a generic way in that their arguments do not depend on idiosyncrasies of the FRB/US model.

### *4.2 The Policy Problem*

Formally, a policy rule is optimized by choosing the parameters, $\Phi$, within the policy rule, $r = \Phi(x)$, where $\Phi = \{\alpha_i, \alpha_j, \alpha_k\}$ $i, j, k \in \{\pi, y, r_{t-1}, \Delta y^*, \Delta yn, \Delta u, \Delta w\}, i \neq j \neq k$, to minimize a loss function, subject to a given model vintage, $x = f(\cdot)$, and a given set of stochastic shocks, $\Sigma$. In the present case, this is written in generic terms as

$$\underset{\langle\Phi\rangle}{MIN} \sum_{i=0}^{T} \beta^i \left[ \left(\pi_{t+i} - \pi^*_{t+i}\right)^2 + \lambda_y \left(u_{t+i} - u^*_{t+i}\right)^2 + \lambda_{\Delta r} (\Delta r_{t+i})^2 \right] \tag{1}$$

subject to

$$x_t = f(x_t, \ldots x_{t-m}, z_t, \ldots z_{t-n}, r_t, \ldots r_{t-p}) + v_t \quad m, n, p \geqslant 0 \tag{2}$$

and

$$\Sigma_v = v'v, \tag{3}$$

where $u$ is the unemployment rate, $u^*$ is the vintage-consistent estimate of the natural rate of unemployment, $r$ is the federal funds rate, $x$ is a vector of endogenous variables, and $z$ is a vector of exogenous variables, both in logs, except for those variables measured in rates. Note that $\pi, y, y^*, u, r, rr^*, w, yn \in x$ while $\pi^*, u^* \in z$.[19] In principle, the loss function, (1), could have been derived as the quadratic approximation to the true social welfare function for the FRB/US model. However, it is technically infeasible for a model the size of FRB/US. That said, with the possible exception of the term penalizing the change in the federal funds rate, the arguments to (1) are

---

[19] The intercept used in the policy rules, where applicable, designated $rr^*$, is a medium-term proxy for the equilibrium real interest rate. It is an endogenous variable in the model. In particular, $rr^*_t = (1-\gamma)rr^*_{t-1} + \gamma(rn_t - \pi_t)$, where $r$ is the federal funds rate, and $\gamma = 0.05$. As a robustness check, I experimented with adding a constant in the optimized rules in addition to $rr^*$ and found that this term was virtually zero for every model vintage. Note that relative to the classic version of the Taylor rule where $rr^*$ is fixed, this alteration biases results in favor of good performance by this class of rules.

standard.[20] The penalty on the change in the funds rate may be thought of as representing either a hedge against model uncertainty in order to reduce the likelihood of the federal funds rate entering ranges beyond those for which the model was estimated, or as a pure preference of the Committee. Whatever the reason for its presence, the literature confirms that some penalty is needed to explain the historical persistence of monetary policy; see, e.g., Sack and Wieland (2000).

The optimal coefficients of a given rule are a function of the model's stochastic shocks, as equation (3) indicates.[21] The optimized coefficient on the output gap, for example, represents not only the fact that unemployment rate stabilization—and hence, indirectly, output-gap stabilization—is an objective of monetary policy, but also that in economies where demand shocks play a significant role, the output gap will statistically lead changes in inflation in the data; so the output gap will appear because of its role in forecasting future inflation. However, if the shocks for which the rule is optimized turn out not to be representative of those that the economy will ultimately bear, performance will suffer. As we shall see, this dependence will turn out to be significant for the results.

### *4.3 Computation*

Solving a problem like this is easily done for small, linear models; FRB/US, however, is a large, non-linear model. Given the size of the model, and the differences across vintages, I optimized the policy rule coefficients employing a sophisticated derivative-free optimization procedure with distributed processing. Specifically, each vintage of the model is subjected to bootstrapped shocks from its stochastic shock archive. Historical shocks from the estimation period of the key behavioral equations are drawn.[22] In all, 1,500 draws of eighty

---

[20] Qualitatively speaking, the results are the same if the output gap is substituted for the unemployment gap in (1), provided the proper normalization of the weight is taken to account for the relative size of unemployment gaps and output gaps over the business cycle.

[21] The rules will be optimal in the relevant class, conditional on the stochastic shock set, (3), under anticipated utility as defined by Kreps (1998).

[22] The number of shocks used for stochastic simulations has varied with the vintage, and generally has grown. For the first vintage, forty-three shocks were used, while for the November 2003 vintage, seventy-five were used.

periods each are used for each vintage to evaluate candidate parameterizations. The target rate of inflation is taken to be 2 percent as measured by the annualized rate of change of the personal consumption expenditure price index.[23] The algorithm is described in detail in Gray and Kolda (2004) and Kolda (2004); here I provide just a thumbnail sketch. In the first step, the rule is initialized with a starting guess; that guess and some neighboring points are evaluated. In the case of two-parameter rules, I need only investigate four neighboring points: higher and lower, by some step size, for each of the two parameters, with the initial guess in the middle. The loss function is evaluated for each of the five points, and the one with the lowest loss becomes the center of the next cluster of five points. As the five points become less and less distinguishable from one another, the step size is reduced until the convergence criterion is satisfied.

Optimization of a two-parameter policy rule using a single Intel Xeon 2.8 GHz machine can take over ten hours, depending on the rule; distributed processing speeds things up. Because this exercise is computationally intensive, I am limited in the range of preferences I can investigate. Accordingly, I discuss only one set of preferences: equal weights on output, inflation, and the change in the federal funds rate. This is the same set of preferences that have been used in optimal policy simulations carried out for the FOMC; see Svensson and Tetlow (2005).

## 5. Results

### *5.1 The Taylor Rule*

Let us begin with the Taylor rule (TR). In this instance, I provide a full set of results—that is, optimized parameters for each of the forty-six vintages; later I will narrow the focus. The results are best summarized graphically. In figure 4, the solid line is the optimized coefficient for the TR on inflation, $\alpha_\pi$, while the dashed line is the feedback coefficient on the output gap, $\alpha_Y$. Perhaps the most noteworthy observation from figure 4 is the distinct upward creep, on average, in both parameters. The inflation response coefficient never

---

[23] For these experiments, any reasonable target will suffice since the stochastic simulations effectively randomize over initial conditions.

**Figure 4. Optimized Coefficients of the Taylor Rule, by Vintage**

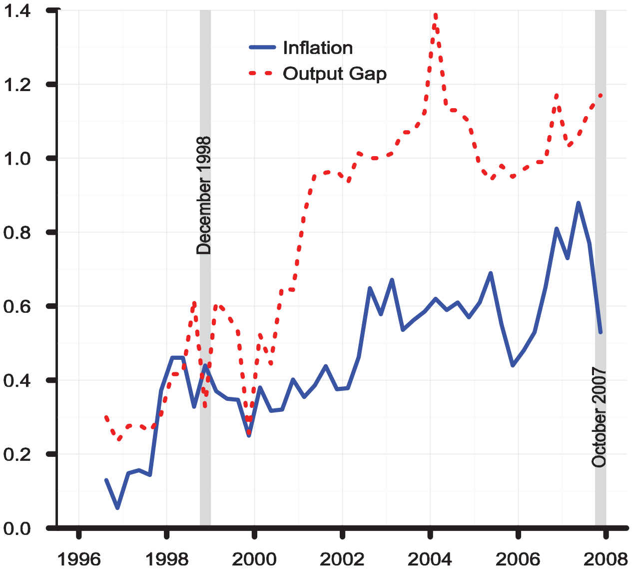

actually gets very large: it starts out quite low, and only in the new century does it climb above the 0.5 percent level of the canonical Taylor (1993) rule. The rise over time in the output-gap coefficient is more impressive. It too starts out low, with the first vintage in July 1996 at about 0.1, but then rises more or less steadily thereafter—the late-1999 dip aside—reaching values generally above 1 with the later vintages.[24]

The sharp increase in the gap coefficient in 2001 coincides with the inclusion of a new investment block which, in conjunction with changes to the supply block, tightened the relationship between supply-side disturbances and subsequent effects on aggregate demand, particularly over the longer term.[25] The new

---

[24] There is also a sharp jump in the gap coefficient over the first two quarters of 2001. One might be tempted to think that this is related to the jump in the sacrifice ratio, shown in figure 1. In fact, the increase in the optimized gap coefficient precedes the jump in the sacrifice ratio.

[25] In essence, the linkage between a disturbance to total factor productivity (TFP) and the desired capital stock in the future was clarified and strengthened

investment block, in turn, was driven by two factors: the earlier inclusion by the Bureau of Economic Analysis of software in the definition of equipment spending and the capital stock, and associated enhanced appreciation on the part of the staff of the importance of the ongoing productivity and investment boom. In any case, while the upward jump in the gap coefficient stands out, it bears recognizing that the rise in the gap coefficient was a continual process.

The point to be taken from figure 4 is that the time variation in model properties, described in section 3, carries over into substantial variation in the optimized TR policy parameters. At the same time, it is clear that time variation in the multipliers is not the sole reason why optimized TR coefficients change. In fact, changes in the stochastic structure of the economy are also in play. To the extent these differences in optimized parameters, conditional on the stochastic shocks, imply significant differences in economic performance, I can say that model uncertainty is a significant problem. I can examine this question by comparing the performance of the optimized TR against other plausible parameterizations. For this exercise and nearly all that follow, I narrow the focus to just two vintages: the December 1998 vintage and the October 2007 vintage. (The optimized Taylor-rule coefficients associated with these vintages are indicated in the figure by the gray bars.) These particular vintages were chosen because they were far apart in time, thereby reflecting as many different views of the world as this environment allows, and because their properties are among the most different of any in the set.

In the next section I examine the implications for economic performance of the TR and the other optimized simple rules for two selected model vintages.

## 5.2 Optimized Rules and Performance

### 5.2.1 Two-Parameter Rules

To this point, I have compared model properties and optimized policies but have had nothing directly to say about performance. This

---

so that an increase in TFP that may produce excess supply in the very short run can be expected to produce an investment-led period of excess demand later on.

section fills this void. I consider the performance, on average, of the model economies under stochastic simulation. I also expand the study to encompass the wider range of simple rules introduced in sections 4.1.1 and 4.1.2. At the same time, in order to make the computational costs feasible, I focus on results for the two selected vintages. Table 2 shows the performance for the complete set of two-parameter rules. The table is divided into two sections, one for each of the December 1998 and October 2007 vintages. In both sections, losses have been normalized on the performance of the optimized Taylor rule so that the efficacy of other rules can be interpreted as multiples of the TR loss.

Before delving into the numbers, it is useful to recall that the results in this table pertain to monetary authorities that understand the nature of the economy they control, including the shocks to which the economy is subject. That is, I am setting aside, for the moment, the issue of model uncertainty, which I take up in the next section. With this in mind, let us focus for the moment on the optimized parameters and normalized losses for the December 1998 vintage shown in table 2. The results show, first, why the TR has been a popular specification for policy design: it renders a very good performance with losses that are lower than nearly all of the alternatives. The one rule that outperforms TR is WN2, shown on line 8, which is identical to the Taylor rule but replaces price inflation with wage inflation. This rule is a version of the rule championed by Levin et al. (2006) on the grounds that in many models, it is wages that are the source of most nominal stickiness. It is not simply feedback on wages that is important to this result, however; the performance of WN1, on line 7, shows that a rule that replaces price inflation with wage inflation as the nominal anchor, but omits direct feedback on the output gap in favor of persistence in funds rate setting through the presence of a lagged funds rate, is the worst rule among those shown. There are other rules that are not far behind the TR in terms of performance, including the (change in) unemployment rate rule, DUR, line 4, with a loss only of 19 percent more than the Taylor rule, and the price-level targeting rule, line 3, which carries a loss only slightly above that of the Taylor rule. This latter result may seem familiar to results seen elsewhere that show strong performance of price-level targets. However, to the best of my knowledge, prior results have been exclusively for linear rational expectations

## Table 2. FRB/US Model Performance in Stochastic Simulation* (two-parameter rule optimizations, selected vintages)

| | **Parameters →** | **Anchor** | **Real** | **December 1998** | | | **October 2007** | | |
|---|---|---|---|---|---|---|---|---|---|
| **Line** | **Policy Rule ↓** | $i$ | $j$ | $\alpha_i$ | $\alpha_j$ | $Norm^\dagger$ | $\alpha_i$ | $\alpha_j$ | $Norm^\dagger$ |
| 1 | Taylor Rule | $\pi$ | $y$ | 0.44 | 0.33 | 1.00 | 0.53 | 1.17 | 1.00 |
| 2 | Inflation Target | $\pi$ | $r$ | 0.87 | −0.32 | 1.44 | −0.30 | −0.81 | 4.51 |
| 3 | Price Level | $p$ | $r$ | 8.14 | 0.46 | 1.03 | 14.70 | 1.14 | 0.96 |
| 4 | Change in U-Rate | $\pi$ | $\Delta u$ | 0.16 | −2.52 | 1.19 | 0.08 | −3.60 | 0.88 |
| 5 | Nominal Output 1 | $\Delta yn$ | $r$ | 0.20 | 0.93 | 1.35 | 0.37 | 0.90 | 1.42 |
| 6 | Nominal Output 2 | $\Delta yn$ | $y$ | 0.02 | 0.43 | 1.26 | 0.02 | 1.12 | 1.06 |
| 7 | Wage Growth 1 | $\Delta w$ | $r$ | 1.05 | −0.41 | 1.46 | 0.63 | −0.71 | 1.57 |
| 8 | Wage Growth 2 | $\Delta w$ | $y$ | 0.72 | 0.39 | 0.92 | 0.09 | 1.16 | 1.03 |

*Loss figures in the October 2007 columns cannot be compared with those in the December 1998 columns. †Average value for equation (1) from 1,500 stochastic simulations over twenty years, normalized so that losses are interpretable as multiples of the loss under the optimized Taylor rule.

models where the powerful role of expectations in strengthening the error-correcting properties of such rules is paramount. That good performance arises from a price-level target under the VAR-based expectations approach is remarkable. Of related interest is the fact that the price-level rule significantly outperforms the IT rule. Inflation targeting allows "bygones to be bygones" in the control of the price level, whereas price-level errors have to be reversed in price-level targeting regimes. Reversing price-level errors is a good thing when agents *know* that the central bank will do this, because anticipated reversals of the price level implies strongly anchored expectations for the inflation rate. When expectations are "boundedly rational," however, the conventional wisdom has been that bringing the price level back to some predetermined path will be all cost and no benefit. We see here that this is not so for the VAR-based expectations of the FRB/US model.

More generally, the performances of the other rules are not greatly different from the Taylor rule; as noted, the WN1 performs the worst, but its loss is only about 1-1/2 times that of the TR, not a good performance but not disastrous either. Evidently, controlling the economy of the December 1998 vintage is a relatively straightforward task.

Let us turn now to the far-right columns, where parallel results are shown for the October 2007 vintage. Here, once again, the TR does pretty well, on average, but in this instance there are two rules that do better, the price-level rule and the DUR. I have already noted that parameterizations of these rules did well in the December 1998 vintage. In addition, two other rules also performed almost as well as the TR: the YN2 and WN2. These rules share two important features. First, they employ feedback on a nominal variable that attempts to correct, albeit indirectly, for trend productivity growth and errors in its measurement. Second, they maintain feedback on the output gap. Thus, notwithstanding the mismeasurement issues associated with persistent changes in productivity *growth*, feedback on the output gap, which is subject to errors in productivity *levels*, is still beneficial, as can be seen by comparing line 6 with line 5, on the one hand, and line 8 with line 7, on the other. In other words, these two rules produce good results but not entirely for the reasons for which these rules were originally advocated.

The last word on this section of the table concerns, once again, the inflation-targeting rule, IT. Its performance controlling the October 2007 vintage could fairly be described as terrible, at 4-1/2 times the loss of the Taylor rule. Qualitatively, this is similar to the results for the December 1998 vintage, but quantitatively much worse. The reasons for this stem from the forementioned tightening of the linkages between the supply block of the model and subsequent aggregate demand fluctuations, together with the nature of the shocks that were incurred during the period over which the two rules are optimized. The rules for the December 1998 vintage are conditioned on shocks from 1981 to 1995, while the October 2007 vintage is conditioned on shocks from 1988 to 2002. The former period was dominated by garden-variety demand shocks, whereas the latter had large and persistent disturbances to aggregate supply—in particular, the productivity boom of the second half of the 1990s. Moreover, many of the key shocks borne during the more recent period were larger than was the case in the earlier period.[26] An implication of productivity booms is that they disrupt the "normal" time-series relationship between output (or employment) and inflation: when output fluctuations are dominated by demand shocks, and prices are sticky, output will statistically lead inflation, and the optimized parameters of rules like the Taylor rule will reflect that relationship. When demand shocks are the prevalent force behind output fluctuations, there is no dilemma for monetary policy: stabilizing output and stabilizing inflation are simultaneously achievable because they are more or less the same thing. It follows that policy can feed back on output (or its proxies) or inflation, and achieve good results either way. However when supply shocks drive cycles, inflation and output will tend to move in opposite directions, setting up a dilemma for the policymaker. Under these circumstances, responding to output and to inflation are no longer good substitutes for the purposes of

---

[26] This argument will clash with the intuition of readers familiar with the literature on the Great Moderation which suggests that shocks are smaller than they once were. The explanation is twofold: first, the period we are dealing with here is much shorter and has smaller residuals in both data sets. Just as important perhaps is a fallacy in the construction of the residuals in many studies that allege that shocks are smaller recently. The regressions from which these conclusions are drawn allow either a time trend or a free constant so that persistent supply-side shocks are mopped up in these terms.

minimizing losses, and responding strictly to inflation is insufficient for controlling output.

### *5.2.2 Three-Parameter Rules*

Table 3 tests the appropriateness of using the two-parameter Taylor rule, TR, as the benchmark by considering the simple extensions noted in section 4.1.2. In particular, the second row of the table shows that the performance of the Taylor rule extended to allow an optimized parameter on the lagged instrument—that is, the inertial Taylor rule, TRI—renders only slightly better performance than the TR itself, for either vintage. Moreover, the attempt through the use of a productivity growth term in the DY* fares worse, as shown in the third line.[27] The final two lines of the table exhibit the advantage of allowing feedback on the lagged instrument relative to the YN2.

This is the one case where adding the lagged instrument to a rule that already has a nominal anchor variable and an aggregate demand term pays off in a significant way. Still, none of these rules do markedly better than the Taylor rule despite the advantage of an added parameter. I conclude that using the Taylor rule as the benchmark is not erecting a straw man. Thus, I am satisfied that focusing attention, henceforth, on two-parameter policy rules is a suitable restriction.

My goal in this paper has been to uncover policies that are both effective and robust across models. To this point, I have identified rules that, when properly specified, perform well in contexts where they should perform well; that is, they are effective. The ones that do not—the inflation-targeting rule, and nominal income and wage growth rules that include the lagged instrument as their second argument—are not candidates as robust performers. Whether the effective rules are also robust is the subject of the next section.

---

[27] It should be the case that the addition of an added parameter cannot do worse than the best two-parameter rule. The contradictory result shown in the table is an artifact of occasional crashes in the optimization algorithm owing to the instability of the extended rule. Still, the instability of rule is, itself, a warning against such a rule.

### Table 3. FRB/US Model Performance in Stochastic Simulation (three-parameter rule optimizations, selected vintages)

| $\alpha_{ijk} \rightarrow$ | | **Anchor** | **Real** | **Added** | **December 1998** | | | | **October 2007** | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Line** | **Rule**↓ | $\boldsymbol{i}$ | $\boldsymbol{j}$ | $\boldsymbol{k}$ | $\boldsymbol{\alpha_i}$ | $\boldsymbol{\alpha_j}$ | $\boldsymbol{\alpha_k}$ | ***Loss***[†] | $\boldsymbol{\alpha_i}$ | $\boldsymbol{\alpha_j}$ | $\boldsymbol{\alpha_k}$ | ***Loss***[†] |
| 1 | TR | $\pi$ | $y$ | – | 0.44 | 0.33 | – | 1.00 | 0.53 | 1.17 | – | 1.00 |
| 2 | TRI | $\pi$ | $y$ | $r$ | 0.33 | 0.29 | 0.33 | 0.98 | 0.22 | 1.07 | 0.22 | 0.98 |
| 3 | DY* | $\pi$ | $y$ | $\Delta y^*$ | 0.38 | 0.36 | 0.10 | 1.04 | 0.41 | 1.23 | 0.29 | 1.02 |
| 4 | YN2 | $\Delta yn$ | $y$ | – | 0.02 | 0.43 | – | 1.26 | 0.02 | 1.12 | – | 1.06 |
| 5 | YN1 | $\Delta yn$ | $y$ | $r$ | 0.13 | 0.10 | 0.88 | 1.04 | 0.23 | 0.36 | 0.73 | 0.97 |

[†] Normalized losses. See the notes to table 2.

## 6. Robustness

I now turn to the principal issue, the robustness of optimized policies to model misspecification. The thought experiment is to imagine a policymaker who believes she is controlling the December 1998 economy model, but in half of the instances I discuss below, it turns out that it is the October 2007 vintage that is the true model. Those results are presented in table 4. Then, in table 5, I reverse the exercise by having our central banker assume she is controlling the October 2007 vintage, but it turns out that half of the time, it is the December 1998 vintage that is the correct model.

The same eight two-parameter rules as before are considered, for two vintages, comprising sixteen parameterizations. I subject both of these models to the same set of stochastic shocks as in the optimization exercise, for each candidate rule. As before, I am mostly interested in normalized losses where the normalization sets the loss under the appropriate optimized TR policy to unity (although I do show the absolute losses, for completeness). Before proceeding with the results, it is worth recalling, at the risk of oversimplification, that the December 1998 vintage is a model that sees the U.S. economy as being relatively stable and easy to control: rule parameterizations that are optimal for the December 1998 vintage are generally less aggressive than their October 2007 counterparts.

Beginning with the TR, where the normalized loss is unity by definition, we see that a policymaker who uses the October 2007 parameterization of that rule incurs losses about two-thirds higher than what she could have achieved had she known the true model; the Taylor rule is not particularly robust in this sense. The inflation-targeting rule, not a particularly good performer at the best of circumstances, is disastrous when misspecified, as shown on line 4. Among the top performers—at least when the true economy turns out to be the December 1998 vintage—are the price-level rule, lines 5 and 6; the change-in-unemployment rule, lines 11 and 12; and the wage growth rule that includes the gap, WN2, lines 15 and 16. Each of these rules performs at least as well as the Taylor rule when misspecified, and provides performance that is close to that of the TR

**Table 4. Normalized Model Performance for Optimized Two-Parameter Rules under Stochastic Simulation* (December 1998 model vintage)**

| | | | **Anchor Variables ($\alpha_i$)** | | | | **Real Variables ($\alpha_j$)** | | | **Dec. 1998 Loss** | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **Line** | **Rule** | **Vin.** | $\pi$ | $\Delta yn$ | $\Delta w$ | $p$ | $y$ | $r$ | $\Delta u$ | **Abs.** | **Norm.** |
| 1 | TR | D98 | 0.44 | | | | 0.33 | | | 17.6 | 1.00 |
| 2 | | O07 | 0.53 | | | | 1.17 | | | 29.2 | 1.66 |
| 3 | IT | D98 | 0.87 | | | | | −0.32 | | 25.4 | 1.44 |
| 4 | | O07 | −0.30 | | | | | −0.81 | | 406.0 | 23.00 |
| 5 | PLT | D98 | | | | 8.14 | 0.46 | | | 18.3 | 1.04 |
| 6 | | O07 | | | | 14.70 | 1.14 | | | 26.9 | 1.53 |
| 7 | YN1 | D98 | | 0.20 | | | | 0.93 | | 23.7 | 1.35 |
| 8 | | O07 | | 0.37 | | | | 0.90 | | 33.5 | 1.90 |
| 9 | YN2 | D98 | | 0.02 | | | 0.43 | | | 22.2 | 1.26 |
| 10 | | O07 | | 0.02 | | | 1.12 | | −2.52 | 31.0 | 1.76 |
| 11 | DUR | D98 | 0.16 | | | | | | −3.60 | 21.0 | 1.19 |
| 12 | | O07 | 0.08 | | | | | −0.41 | | 24.3 | 1.38 |
| 13 | WN1 | D98 | | | 1.05 | | | −0.71 | | 25.7 | 1.57 |
| 14 | | O07 | | | 0.63 | | | | | 27.8 | 1.58 |
| 15 | WN2 | D98 | | | 0.72 | | 0.39 | | | 16.2 | 0.92 |
| 16 | | O07 | | | 0.09 | | 1.16 | | | 29.4 | 1.67 |

*Selected rules and model vintages. Average losses from 1,500 draws of eighty periods each.

when properly specified.[28] The YN2 is not far off the mark set by the optimized Taylor rule.

Table 5 turns the exercise around by considering the case where the October 2007 vintage turns out to be the correct one. Misspecification of the Taylor rule is more costly here: the deterioration relative to the best policy parameterization is 80 percent. Once again, the TRI performs very poorly, while *most* of the rules that do include feedback on the output gap—the Taylor rule, the price-level rule, one of the nominal output rules, and the change-in-unemployment rule—all perform well. The one notable exception to the conclusion that feedback on the output gap is always a good thing is the WN2, where misspecification of the rule, as in line 16, results in large losses relative to the TR and most alternatives to it. Even here, though, it seems that it is feedback on wage growth that is the key to this result, as the rules in lines 13 and 14, which respond to wage growth and the lagged instrument, but not the output gap, perform even worse. What this tells us is that while a wage growth rule can turn in a very good performance, as it does when paired with the output gap on line 15, a good calibration is critical to its performance; the rule is not robust.

The PLT turns in an even stronger performance for the October 2007 vintage than it did for the December 1998 one. This result obtains notwithstanding that the parameterizations of the two rules differ significantly: the feedback parameters on the output gap are 1.14 and 0.46. As in rational expectations models, an important contribution to economic performance under this rule is that constraining the drift in the price level anchors inflation fluctuations. In both vintages of the FRB/US model, keeping inflation in check also limits cycling in long-term expected inflation. The stability of long-term inflation expectations reinforces the stabilizing force of policy, making output stabilization less critical than would otherwise be the case.

This case contrasts sharply with the change-in-unemployment rule, DUR. For this rule, feedback on inflation itself is slight at 0.08

---

[28] If I was to take a Bayesian perspective on this and assume that the two vintages are equally probable, the average values of the losses from the PLT, the DUR, and the YN2 are all less than that of the Taylor rule, for this model vintage.

### Table 5. Normalized Model Performance for Optimized Simple Rules under Stochastic Simulation* (October 2007 model vintage)

| | | | Anchor Variables ($\alpha_i$) | | | | Real Variables ($\alpha_j$) | | | Oct. 2007 Loss | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **Line** | **Rule** | **Vin.** | $\pi$ | $\Delta yn$ | $\Delta w$ | $p$ | $y$ | $r$ | $\Delta u$ | **Abs.** | **Norm.** |
| 1 | TR | O07 | 0.53 | | | | 1.17 | | | 17.7 | 1.00 |
| 2 | | D98 | 0.44 | | | | 0.33 | | | 31.9 | 1.80 |
| 3 | IT | O07 | −0.30 | | | | | −0.81 | | 79.9 | 4.51 |
| 4 | | D98 | 0.87 | | | | | −0.32 | | 134.0 | 7.57 |
| 5 | PLT | O07 | | | | 14.70 | 1.14 | | | 17.0 | 0.96 |
| 6 | | D98 | | | | 8.14 | 0.46 | | | 22.8 | 1.29 |
| 7 | YN1 | O07 | | 0.37 | | | | 0.90 | | 25.2 | 1.42 |
| 8 | | D98 | | 0.20 | | | | 0.93 | | 29.8 | 1.68 |
| 9 | YN2 | O07 | | 0.02 | | | 1.12 | | | 18.7 | 1.06 |
| 10 | | D98 | | 0.02 | | | 0.42 | | −3.60 | 24.6 | 1.39 |
| 11 | DUR | O07 | 0.08 | | | | | | −2.52 | 15.6 | 0.88 |
| 12 | | D98 | 0.16 | | | | | −0.71 | | 18.9 | 1.07 |
| 13 | WN1 | O07 | | | 0.63 | | | −0.41 | | 104.0 | 5.89 |
| 14 | | D98 | | | 1.05 | | | | | 41.9 | 6.38 |
| 15 | WN2 | O07 | | | 0.09 | | 1.16 | | | 18.3 | 1.03 |
| 16 | | D98 | | | 0.72 | | 0.39 | | | 57.1 | 3.22 |

*Selected rules and model vintages. 1,500 draws of eighty periods each.

**Figure 5. Optimized Coefficient of the Change-in-Unemployment Rule, by Model Vintage**

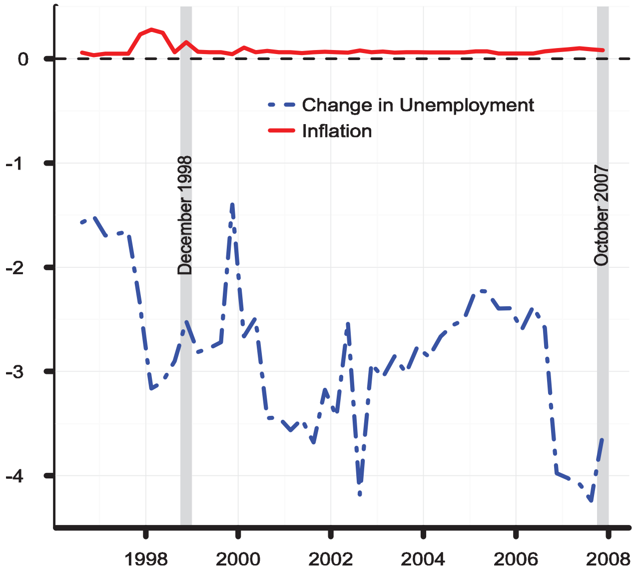

and 0.16. But feedback on the change in the unemployment rate is vigorous: –3.60 and –2.52. Thus, aggressively tempering fluctuations in unemployment is substituting for inflation (and price-level) control. The fact that the DUR is written in the first difference of the instrument, and therefore does not depend on estimates of the equilibrium real rate of interest, is also a factor; this means that the instrument can find the right level even when a productivity shock changes what that level should be. The DUR is the one rule of which I am aware that was tested, by Orphanides and Williams (2002), in an environment that allowed for persistent, unobserved shocks to the “natural rate of interest,” and was found to execute well.

The results for DUR in tables 4 and 5 suggest that it could be a robust rule. However, a closer look at the robustness of DUR is achievable by computing its optimized parameters for all vintages. The results of this exercise are shown in figure 5.

The figure shows that the coefficient on inflation, the solid line, is never much above zero, regardless of the vintage. By contrast, the coefficient on the change in the unemployment rate, the dashed line, jumps around somewhat with perhaps a slight tendency to increase, in absolute terms, over time. The range over the complete set of vintages for the coefficient on the change in the unemployment rate spans from a low of −1.4 for the November 1999 vintage to a high of −4.2 for the August 2007 vintage, considerably wider than the range encompassed by the December 1998 and October 2007 vintages, shown by the gray bars. The computations underlying figure 5 allow us to expand on the robustness analysis of tables 4 and 5 while focusing on the unemployment rate rule. This is done in table 6, where I consider the performance of the most extreme parameterizations of the rule in the two benchmark vintages.

The table shows that when either of the benchmark models is governed by the most extreme parameterization of the DUR rule, the small absolute coefficient on the (change in the) unemployment rate in the November 1999 vintage, the deterioration in control increases the loss relative to the best possible parameterization by a bit over 50 percent, as shown on line 2 of the table. The parameterization that rendered the largest coefficient on the change in the unemployment rate, the August 2007 vintage, gave coefficients that are not much different from the (chronologically close) October 2007 vintage. Thus lines 3 and 4 of the table are similar. Incremental losses, relative to the best possible DUR parameterization, of 50-some percent, are not particularly large in comparison with the results in tables 4 and 5.

## 7. Concluding Remarks

For central banks, the appropriate design of monetary policy under uncertainty is a critical issue. Many conferences are devoted to the subject, and the list of papers is lengthy and still growing. In nearly all instances, however, the articles, whether they originate from central banks themselves or from academia, have tended to be abstract applications. One posits an idealized model, or several models, of the economy and investigates, in some way, how misperceptions of, or perturbations to, the model affect outcomes. A good deal has been learned from these exercises, but results have tended to be specific to the environment of the chosen models. Moreover, the models themselves typically have not been representative of the models upon which central banks rely. It is difficult to know how serious a problem

**Table 6. Performance of Selected Parameterizations of Change-in-Unemployment Rate Rule**

| Line | Rule Parameterization | Coefficients | | Losses for Model | | | |
|---|---|---|---|---|---|---|---|
| | | | | December 1998 | | October 2007 | |
| | | $\alpha_{\Delta y}$ | $\alpha_{\pi}$ | Absolute | Normalized | Absolute | Normalized |
| 1 | December 1998 | −2.56 | 0.16 | 20.97 | 1.00 | 26.85 | 1.28 |
| 2 | November 1999 | −1.40 | 0.04 | 33.12 | 1.58 | 31.71 | 1.52 |
| 3 | August 2007 | −4.24 | 0.09 | 26.63 | 1.27 | 21.05 | 1.01 |
| 4 | October 2007 | −3.94 | 0.08 | 25.62 | 1.22 | 20.85 | 1.00 |

model uncertainty is if one cannot give a concrete and meaningful measure of uncertainty.

This paper has cast some light on model uncertainty and the design of policy in a much different context from the extant literature. I examined forty-six vintages of the model the Federal Reserve Board staff has used to carry out forecasts and policy analysis from 1996 to 2007. And I did so in a real-time context that focuses on the real problems that the Federal Reserve faced over this period. My examination looked at a number of simple policy rules that have been marketed as "robust." In the end, I uncovered a number of useful observations. First, model uncertainty is a substantial problem. Changes to the FRB/US model over the period of study were frequent and often important in their implications. The ensuing optimized policies also differed significantly in their parameterizations. Second, many simple rules that have been touted as robust turn out to be less appealing than one might have suspected. In particular, pure inflation-targeting rules turn out not to be robust. Third, adding an instrument smoothing term to a rule that already has a nominal anchor and a real variable contributes little to the robustness and efficiency of rules, at least in the environment studied here. Fourth, notwithstanding problems of mismeasurement of output gaps, it generally pays for policy to feed back on some measure of excess demand, regardless of the nominal anchor employed elsewhere in the rule. Fifth, a case can be made for designing simple rules that minimize the use of latent variables like potential output and the equilibrium real interest rate as arguments.

## References

Aoki, K. 2003. "On the Optimal Monetary Policy Response to Noisy Indicators." *Journal of Monetary Economics* 50 (3): 501–23.

Atkeson, A., and L. E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting?" *Quarterly Review* (Federal Reserve Bank of Minneapolis) 25 (1): 2–11.

Bernanke, B. S., and M. Gertler. 1999. "Monetary Policy and Asset Price Volatility." In *New Challenges for Monetary Policy,* 77–129. Kansas City: Federal Reserve Bank of Kansas City.

Blinder, A. S. 1998. *Central Banking in Theory and Practice.* Cambridge, MA: MIT Press.

Blinder, A. S., and J. L. Yellen. 2001. *The Fabulous Decade: Macroeconomic Lessons from the 1990s.* New York: Century Foundation Press.

Brainard, W. 1967. "Uncertainty and the Effectiveness of Monetary Policy." *American Economic Review* 57 (2): 411–25.

Brayton, F., A. Levin, R. Tryon, and J. C. Williams. 1997. "The Evolution of Macro Models at the Federal Reserve Board." *Carnegie-Rochester Conference Series on Public Policy* 47 (December): 43–81.

Brayton, F., and P. Tinsley, eds. 1996. "A Guide to FRB/US — A Macroeconomic Model of the United States." FEDS Paper No. 1996-42, Board of Governors of the Federal Reserve System.

Brock, W. A., S. N. Durlauf, and K. D. West. 2007. "Model Uncertainty and Policy Evaluation: Some Theory and Empirics." *Journal of Econometrics* 136 (2): 629–64.

Cecchetti, S., H. Genberg, J. Lipsky, and S. Wadhwani. 2000. *Asset Prices and Central Bank Policy.* The Geneva Report on the World Economy, Vol. 2. London: Center for Economic Policy Research.

Cimadomo, J. 2008. "Fiscal Policy in Real Time." ECB Working Paper No. 919.

Cogley, T., and T. J. Sargent. 2005. "The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty." *Review of Economic Dynamics* 8 (2): 528–63.

Croushore, D., and T. Stark. 2001. "A Real-time Data Set for Macroeconomists." *Journal of Econometrics* 105 (1): 111–30.

Dennis, R. 2001. "Inflation Expectations and the Stability Properties of Nominal GDP Targeting." *Economic Journal* 111 (468): 103–13.

Ehrmann, M., and F. Smets. 2003. "Uncertain Potential Output: Implications for Monetary Policy." *Journal of Economic Dynamics and Control* 27 (9): 1611–38.

English, W. B., D. López-Salido, and R. J. Tetlow. 2013. "The Federal Reserve's Framework for Monetary Policy—Recent Changes and New Questions." FEDS Paper No. 2013-76, Board of Governors of the Federal Reserve System. Available at http://www.federalreserve.gov/pubs/feds/2013/201376/201376pap.pdf. Forthcoming in *IMF Economic Review.*

Garratt, A., G. Koop, E. Mise, and S. P. Vahey. 2007. "Real-time Prediction with UK Monetary Aggregates in the Presence of Model Uncertainty." Birkbeck Working Papers in Economics and Finance No. 0714, Birkbeck School of Economics, Mathematics and Statistics.

Garratt, A., and S. P. Vahey. 2006. "UK Real-Time Macro Data Characteristics." *Economic Journal* 116 (509): F119–35.

Giannone, D., L. Reichlin, and L. Sala. 2005. "Monetary Policy in Real Time." In *NBER Macroeconomics Annual 2004*, ed. M. Gertler and K. Rogoff, 161–200. Cambridge, MA: NBER and MIT Press.

Giuliodori, M., and R. Beetsma. 2008. "On the Relationship between Fiscal Plans in the European Union: An Empirical Investigation Based on Real-Time Data." *Journal of Comparative Economics* 36 (2): 221–42.

Gray, G. A., and T. G. Kolda. 2004. "APPSPACK 4.0: Asynchronous Parallel Pattern Search for Derivative-Free Optimization." Unpublished Manuscript, Sandia National Laboratories.

Jääskelä, J., and T. Yates. 2005. "Monetary Policy and Data Uncertainty." *Quarterly Bulletin* (Bank of England) 45 (4, Winter): 477.

Keynes, J. M. 1930. *The Great Slump of 1930.* 2009 Kindle edition. Orlando: Signalman.

Kimura, T., and T. Kurozumi. 2007. "Optimal Monetary Policy in a Micro-Founded Model with Parameter Uncertainty." *Journal of Economic Dynamics and Control* 31 (2): 399–431.

Kolda, T. G. 2004. "Revisiting Asynchronous Parallel Pattern Search." Unpublished Manuscript, Sandia National Laboratories.

Kreps, D. 1998. "Anticipated Utility and Dynamic Choice." In *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997,* ed. D. P. Jacobs, E. Kalai, M. I. Kamien, and N. L. Schwartz. Cambridge: Cambridge University Press.

Levin, A. T., A. Onatski, J. C. Williams, and N. Williams. 2006. "Monetary Policy under Uncertainty in Microfounded Macroeconometric Models." In *NBER Macroeconomics Annual 2005*, ed. M. Gertler and K. Rogoff, 229–87. Cambridge, MA: NBER and MIT Press.

Levin, A., V. Wieland, and J. C. Williams. 1999. “Robustness of Simple Monetary Policy Rules under Model Uncertainty.” In *Monetary Policy Rules,* ed. J. B. Taylor, 263–99. Chicago: University of Chicago Press.

McCallum, B. 1988. “Robustness Properties of a Rule for Monetary Policy.” *Carnegie-Rochester Conference Series on Public Policy* 29: 173–204.

McCallum, B., and E. Nelson. 1999. “Nominal Income Targeting in an Open-Economy Optimizing Model.” *Journal of Monetary Economics* 43 (3): 553–78.

Onatski, A., and N. Williams. 2003. “Modeling Model Uncertainty.” *Journal of the European Economic Association* 1 (5): 1087–1122.

Orphanides, A. 2001. “Monetary Policy Rules Based on Real-time Data.” *American Economic Review* 91 (4): 964–85.

Orphanides, A., R. Porter, D. Reifschneider, R. Tetlow, and F. Finan. 2000. “Errors in the Measurement of the Output Gap and the Design of Monetary Policy.” *Journal of Economics and Business* 52 (1–2): 117–41.

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

Primiceri, G. 2005. “Time-varying Structural Vector Autoregressions and Monetary Policy.” *Review of Economic Studies* 72 (3): 821–52.

Reifschneider, D. L., R. J. Tetlow, and J. C. Williams. 1999. “Aggregate Disturbances, Monetary Policy and the Macroeconomy: The FRB/US Perspective.” *Federal Reserve Bulletin* (January): 1–19.

Romer, C., and D. Romer. 2002. “The Evolution of Economic Understanding and Postwar Stabilization Policy.” In *Rethinking Stabilization Policy,* 11–78. Kansas City: Federal Reserve Bank of Kansas City.

Rudebusch, G. 2002. “Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty.” *Economic Journal* 112 (479): 402–32.

Sack, B., and V. Wieland. 2000. “Interest-rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence.” *Journal of Economics and Business* 52 (1–2): 205–28.

Sims, C., and T. Zha. 2006. “Were There Regime Shifts in U.S. Monetary Policy?” *American Economic Review* 96 (1): 54–81.

Söderström, U. 2002. “Monetary Policy with Uncertain Parameters.” *Scandinavian Journal of Economics* 104 (1): 5–145.

Staiger, D., J. H. Stock, and M. Watson. 2001. “Prices, Wages and the U.S. NAIRU in the 1990s.” NBER Working Paper No. 8320 (June).

Svensson, L. E. O., and R. Tetlow. 2005. “Optimal Policy Projections.” *International Journal of Central Banking* 1 (3): 177–207.

Taylor, J. B. 1993. “Discretion Versus Policy Rules in Practice.” *Carnegie-Rochester Conference Series on Public Policy* 39: 195–214.

Taylor, J. B., and V. Wieland. 2012. “Surprising Comparative Properties of Monetary Models: Results from a New Model Database.” *Review of Economics and Statistics* 94 (3): 800–816.

Tetlow, R. J., and B. Ironside. 2007. “Real-Time Model Uncertainty in the United States: The Fed, 1996-2003.” *Journal of Money, Credit and Banking* 39 (7): 1533–61.

Tetlow, R., and P. von zur Muehlen. 2001. “Robust Monetary Policy with Misspecified Models: Does Model Uncertainty Always Call for Attenuated Policy?” *Journal of Economic Dynamics and Control* 25 (6–7): 911–49.

Walsh, C. 2004. “Implications of a Changing Economic Structure for the Strategy of Monetary Policy.” In *Monetary Policy and Uncertainty: Adapting to a Changing Economy*, 297–346. Kansas City: Federal Reserve Bank of Kansas City.

———. 2006. “Comment [on ‘Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models’ by Levin, Onatski, Williams, and Williams].” In *NBER Macroeconomics Annual 2005*, ed. M. Gertler and K. Rogoff, 297–308. Cambridge, MA: NBER and MIT Press.

Williams, J. C. 2003. “Simple Rules for Monetary Policy.” In *Economic Review 2003*, 1–12. San Francisco: Federal Reserve Bank of San Francisco.