.

On the Robustness of Simple and Optimal Monetary Policy Rules*

Robert J. Tetlow.
Division of Research and Statistics
Federal Reserve Board

First draft: November 13, 2006; this draft: May 15, 2007.

Abstract

This is a discussion of the paper "Simple versus Optimal Rules as Guides to Policy" by Brock, Durlauf, Nason and Rondina (BDNR) presented in November 2006 at Carnegie-Mellon University under the auspices of the *Carnegie-Rochester Conference Series on Public Policy*. I review the authors' arguments, present a few suggestions for extension and outline where I think at least one strand of the literature should be heading.

- **JEL Classifications:** E37, E5, C6.
- **Keywords:** monetary policy, uncertainty, robustness, minimax, minimax regret, frequency domain

*Contact address: Robert Tetlow, Federal Reserve Board, Washington, D.C. 20551. *Email:* rtetlow@frb.gov. This and other papers can be found at www.roberttetlow.com. I thank Ben McCallum for inviting me to participate and Stephen Durlauf for stimulating dialogue. I also thank Tom Tallarini for helpful remarks and Elizabeth Elzer for assistance with the charts. All remaining errors are mine. The views expressed in this paper are those of the author alone and do not represent those of the Federal Reserve Board or other members of its staff.

1. Introduction

Over the past decade or so, the subject of the appropriate response of monetary policy to uncertainty has resurfaced. The original literature spearheaded by Brainard (1967) has been broadened to include a much wider class of uncertainty in much more elaborate environments. The problem studied by Brainard is now known as *parameter uncertainty*. Other strands of the literature examine different facets of the broader issue, including *data uncertainty* and *model uncertainty*. These strands are complementary but arguably the most lively and analytically difficult is the model uncertainty literature. Besides the type of uncertainty, the literature also bifurcates on approach. In some contributions, uncertainty is taken as parametric, meaning that all the tools of Bayesian decision theory can be applied; in others, uncertainty is taken to be non-parametric so that tools that do not apply probability theory are employed, usually robust control; see, in particular, Gilboa and Schmeidler (1989).

Brock, Durlauf, Nason and Rondina (henceforth, BDNR) take on aggressively the problem of model uncertainty by adopting and adapting the tools from control theory in engineering. In this they are following the territory staked out by Hansen and Sargent in various contributions, especially (2006), as well as others including Sargent (1999), Tetlow and von zur Muehlen (2001), Giannoni (2002), Onatski and Stock (2002) and Onatski (2003). Some but not all of these earlier contributions dwell in the frequency domain as BDNR but none embrace the spectral perspective for analysis as heartily as BDNR do. In addition, BDNR introduce a new criterion for assessing policy rules, minmax regret, which, they argue, deals effectively with the excessively conservative nature of minmax analysis.

In this discussion, I review what BDNR do, add some perspective, and then follow up with a couple of suggestions (which apply more to future work than to the present paper). Finally, I return to the theme of the conference to ask where the broader literature–not just these authors–might head.

2. What BDNR do

BDNR is the sixth paper in a series dating back to 2003 by the first two authors either alone or in conjunction with others.[1] While various authors

have examined model uncertainty from a number of perspectives, Brock and Durlauf are almost unique in their willingness to apply either Bayesian or robust control methods to the problem. The present paper builds on previous contributions, especially Brock, Durlauf and Rondina (2006). Together, the two papers represent a formidable effort, with still more work lurking in the background in appendices. The volume and complexity of the work demonstrates vividly that the authors are after big game.

In this outing, the authors seek to persuade the reader of two arguments: first, the utility of *minmax regret* as a criterion for assessing the performance of candidate monetary policy rules; and, second, the observation that policies that perform "well" in the traditional sense of minimizing some intertemporal quadratic loss function might nonetheless hide unpleasant characteristics in their spectral decompositions. A corollary of the latter argument is that policy advisors should be cognizant of the outcomes of policies by frequency and should present to decision makers the frequency responses of the policies they advocate. The unmistakable conclusion from the applied work they present is that the presumption that so-called targeting rules are superior to simple rules is unwarranted. In a world of model uncertainty, the robustness of targeting rules must be compared on a case-by-case basis against sub-optimal but perhaps more robust alternatives. In my discussion, I shall have nothing to say about minmax regret, not because the subject is unworthy of attention but simply because the argument for robustness in the frequency domain interests me more.

BDNR motivate the attention to the frequency domain on both positive and normative grounds. The positive argument is pure taste: decision makers may simply prefer to damp cycles at some frequencies more than others, even in the absence of model uncertainty. I found this argument less compelling than the normative one: there may be uncertainties associated with either the model or

[1] See Brock, Durlauf and West (2003, 2004), Brock and Durlauf (2004,2005) and Brock, Durlauf and Rondina (2006).

Figure 1

Selected U.S. Macroeconomic Series 1994:Q1 - 2006:Q2

Unemployment rate, Inflation, and the Funds Rate

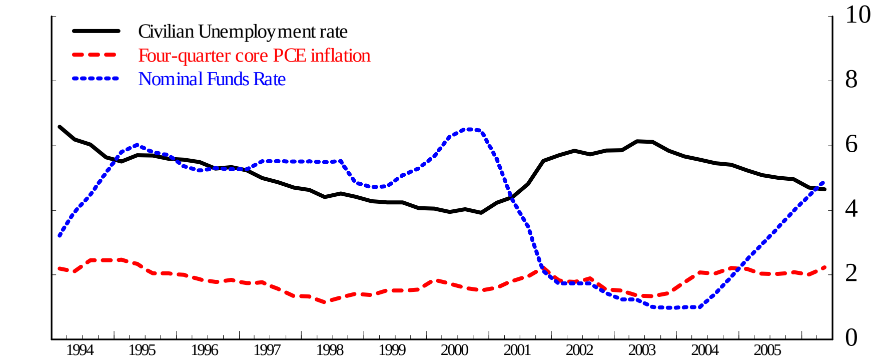

Four-quarter growth in Stock Market and Property Wealth and the Funds Rate

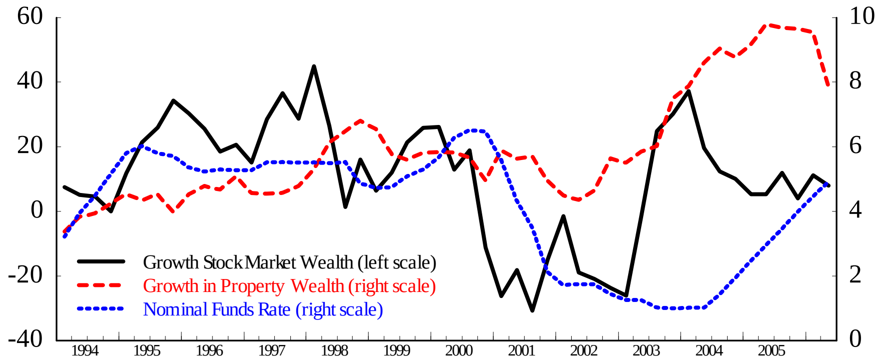

the data that would induce the policy maker to want to avoid model responses in frequency ranges about which he or she knows little. No model is designed to explain phenomena at all frequencies; it stands to reason that policy makers would choose to avoid those frequencies for which the model is not well suited.[2] A real-world example of this is shown in figure 1 above.

The upper panel of the figure shows the federal funds rate (the blue dotted line) along with the four-quarter core PCE inflation rate (the red dashed line) and the unemployment rate (the solid black line) for the United States from 1994:Q1 to 2006:Q2. As Svensson and Tetlow (2005) document, it is the outcomes from policy experiments for these variables that are commonly shown

[2] This is, of course, simply the normative counterpart to the argument of Lucas (1987) on modeling. Onatski and Williams (2003) formulate an argument along these lines.

to Federal Open Market Committee members in the Fed's Bluebook. The figure suggests that the FOMC's conduct of monetary policy over this period was, on the whole, quite successful: inflation stayed close to 2 percent and the unemployment rate remained within a range of 4 to 6 percent. There was a recession in 2001, but by historical standards it was a remarkably mild one. In sum, the low-frequency cycle shown in the upper panel leads one to conclude that policy was conducive to macroeconomic stability.

The bottom panel offers a somewhat different take on the issue. The blue dotted line is the federal funds rate, as before. The black solid line, read off of the left-hand scale, is the four-quarter growth rate in equity prices. The red dashed line, read off of the right-hand scale, is the four-quarter growth rate in property wealth, which is mostly housing market wealth. The Fed's critics point to the asset-price consequences of the monetary policies of the mid-to-late 1990s. Indeed, the black line shows substantial growth in stock prices, on average, relative to GDP, from about 1995 to early 2000, while the red line shows the housing market boom that started in late 2002. The critics split into two camps. One camp suggests that inappropriately smooth funds rate settings in the period through most of the 1990s touched off the stock-market boom.[3] After the bubble burst, the decline in the funds rate to 1 percent in 2003-2004 then led to the housing market boom that is now arguably reversing itself. The other camp argues, to the contrary, that it was the high-frequency response of the Fed to the Asia crisis and Russian debt default in 1998, which manifest itself in sharp reductions of the funds rate, and the subsequent reversal of those reductions in 1999, that brought about the crash of stock prices in 2000-2001. To restate the argument, the critics would say, *either* that ignoring the high-frequency implications on asset prices of business-cycle frequency demand management, *or* fine-tuning the high-frequency financial-market consequences of the Asia crisis, touched off asset pricing cycles that are still ongoing. One does not have to endorse either of these critiques to recognize a *prima facie* argument for worrying about the higher-frequency consequences of the stabilization of low frequency phenomena and *vice versa*.

BDNR present many expressions in their paper; here I shall reprise just two. The first is:

[3] Mussa (2003) argues that it was a (small) mistake for the Fed not to have responded to the emerging stock-market bubble in the late 1990s.

$$f_x^C(\omega) = S^C(\omega) f_x^{NC}(\omega) S^C(\omega)' \tag{2.1}$$

which says that the spectral density of the system under control (superscript C) equals that of the system without control (NC), adjusted by a sensitivity function, $S^C(\omega)$. Then, for a given model, M, and a given control, C, the sensitivity function can be shown to obey:

$$\int_{-\pi}^{\pi} \log[\left|S^C(e^{-i\omega})\right|^2] d\omega = K_B(M, C). \tag{2.2}$$

This is the Bode integral, which is the main building block for where the authors are headed in this paper and in their research agenda more broadly. What equation (2.2) says is that the sensitivities integrate to a constant. The problem is, what constant? Brock, Durlauf and Rondina (2006) show that it depends on the circumstances; there may be different answers for backward-looking models and forward-looking models, univariate models and multivariate ones.

This is an example of a certain uneasiness in the marriage of engineering and economics. First, engineers typically work with either SISO ("single-input, single-output") or MIMO ("multiple-input, multiple-output") models, whereas monetary economists are interested in a single input (that is, one control variable) and several output (target, or state variables). Brock, Durlauf and Rondina (2006) develop, and BDNR apply, some steps forward on this score, but it is fair to say that there is still some ways to go. Second, forward-looking expectations are unique to economics. Here too, some progress had been made. And third, as noted above, the theory from engineering uses a stable "no-control" solution as a benchmark. This is problematic because it is normally the role of the monetary authority to impose a nominal anchor and thus establish the stability of the system. In other words, the system *is not supposed to be stable* in the absence of control.

That said, while it is clear that $f^C(\omega)$ does not exist for a model that is not stable without control, Brock, Durlauf and Rondina (2006) show that K_B is nonetheless calculable. Still, what use a calculable K_B is without the benchmark to which it is applied, is questionable.

In any case, the precise value of K_B will depend on the unstable eigenvalues of the system. These eigenvalues will be nonlinear functions of the parameters of the system. And these parameters, in turn, are subject to considerable mismeasurement because they represent low-frequency dynamics which are known to be difficult to estimate. Thus, the application of equation (2.2) is susceptible to some of the same problems that drive the model uncertainty problem in the first place.

3. A modest proposal

3.1. Rebenchmarking the no-control baseline

Let me propose a simple–and perhaps naïve–work around to the problem of the stability of the system under "no control". Let us express the (simplified) system as follows:

$$x_t = ax_{t-1} + bu_t + \varepsilon_t \tag{3.1}$$

where x is a state variable, u is a control variable, ε is a random disturbance, and a and b are scalar coefficients.[4] Let us assume that $a \geq 1$ so that the system in the absence of control, i.e., $u_t = 0 \; \forall \; t$, is unstable. Now let me conjecture a more-or-less arbitrary controller, $u_t = -fx_{t-1}$. Substitute this into equation (3.1) and define a new controller:

$$\widehat{u}_t = (\widehat{f} - f)x_{t-1} \tag{3.2}$$

so that we have:

$$x_t = (a - bf)x_{t-1} + b\widehat{u}_t + \varepsilon_{t.} \tag{3.3}$$

with $(a - bf) < 1$ by assumption.

What we have done is redefine the application of policy as a *perturbation* to some baseline policy, and in doing so we have rendered stable the formerly no-control baseline model. This has two advantages. First, estimators of stable systems are liable to be more reliable than those that can tolerate instabilities. And second, the interpretation of policy as a perturbation to some other policy seems more natural than a comparison to a fictitious no-control case.

But even if one accepts this argument there remains the question of how one should choose the baseline policy. BDNR spend much of their time comparing three classes of policy rules, a version

[4] The model could easily be generalized to multiple states and multiple (finite) lags–or leads–without any loss of generality.

of the Taylor rule–by which they mean an *ad hoc* policy rule that satisfies the Taylor principle, a restricted optimal policy, and a fully optimal policy. Each of these rules is a commitment strategy. They differ in terms of the number of arguments to the rule, and on whether the rule parameters have been optimized.[5] The fully optimal rule contains all of the states of the model, whereas the restricted optimal (as well as the Taylor rule) restricts feedback to subsets of the state vector. As such, they nest one another nicely and thus one could begin with a baseline policy of, say, "the Taylor rule" and ask how optimization of the parameters of that rule changes performance and the spectrum of the model. Then one can make the jump to the fully optimal rule to see whether the increase in performance that comes from adding states as measured by the loss function comes at the price of some deterioration measured in terms of the spectrum.

That said, there may be other criteria for settling on a base-case feedback rule. To introduce one alternative, let me digress very briefly on the meaning of robustness.

3.2. On robustness

There are at least two classes of robustness, *robust performance* and *robust stability.* Robust performance is the concept that most people think about. It asks whether a rule can be chosen that works "well" in a range of models. Performance is typically measured in terms of the loss using standard loss functions. Assessment of robustness in this sense can be carried out in a parametric (Bayesian) framework, or in a nonparametric (minmax) framework. Robust stability, on the other hand, attempts to find a rule that maximizes the set of models for which the economy is stable. The performance of a robustly stable system, in the sense just described, is something of an afterthought, used to calculate the cost of insuring against instability. Applications of robust stability in economics are few; see, however, Tetlow and von zur Muehlen (2001, 2005). Robust stability is interesting for a several reasons. First, robust stability is worked out using the maximum singular value computations, which operate in the frequency domain, just as BDNR do.[6] As such, they are natural complements. Second, by computing the rule that renders stable

[5] Optimized, that is, in the sense of minimizing some appropriately chosen linear-quadratic Gaussian performance metric.

[6] See, in particular, Dahleh and Diaz-Bobillo (1995) and Zhou, Doyle and Glover (1996), especially chapter 11.

the largest set of models around some reference model, a natural benchmark is established by which one can compare alternative policies both in terms of the spectrum and in some performance metric. Such a policy is the most robust policy from a stability perspective; all other policies can then be compared to the robust stable policy for their performance improvements measured against their stability deteriorations. Third, a corollary of the above is that the kind of tools that BDNR are entertaining can be used to explore the costs and benefits of robustness at various frequencies focussing, perhaps, on those frequencies that were underemphasized in the construction of the original model much in the way that Onatski and Williams (2003) advocate.

4. Looking ahead

The title of this conference is *Mainstream Monetary Policy Analysis circa 2006: Are there reasons for concern?* BDNR point to some areas of concern, arguing that the concept of performance that is used in mainstream monetary policy analysis is overly narrow. I am sympathetic to their point of view. In this section, I add a couple of items to their list.

The first is, *real-time problems*. In my view, given the disagreement among economists regarding the correct workhorse model of the economy, and the associated model uncertainty, it is too much to ask central bankers to select, once and for all, a time-invariate commitment strategy, although that does not diminish the utility of using such rules as guides. An alternative approach is to take seriously the real-time nature of monetary decision making.[7] An example of the immediacy of this problem is demonstrated by Figure 2 below. The figure is an updated version of one that appears in Tetlow and Ironside (2005). It shows the evolution over time in the sacrifice ratio of the Fed staff's FRB/US model of the U.S. economy; that is the loci of simulated values of the unemployment cost of disinflation *by model vintage*. The figure shows that when the FRB/US model went into service in July 1996, the sacrifice ratio was a relatively low 2-1/4 percent, a mainstream estimate at the time.[8] Over time, however, with re-estimations and respecifications of the model, the estimated

[7] I should note that BDNR are very much aware of this issue; they explicitly mention it in the conclusion of their paper.

[8] That such estimates can figure prominently in policy debates is exemplified by the transcript of the July 2-3, 1996, FOMC meeting which quotes then Fed Governor Janet Yellen: "The sacrifice ratio in our new FRB-US model, without credibility effects, is 2.5..." On this basis, she spearheaded a discussion of what long-run target rate

Figure 2

Sacrifice ratio by model vintage
FRB/US Model vintages from 1996:Q3 - 2006:Q3

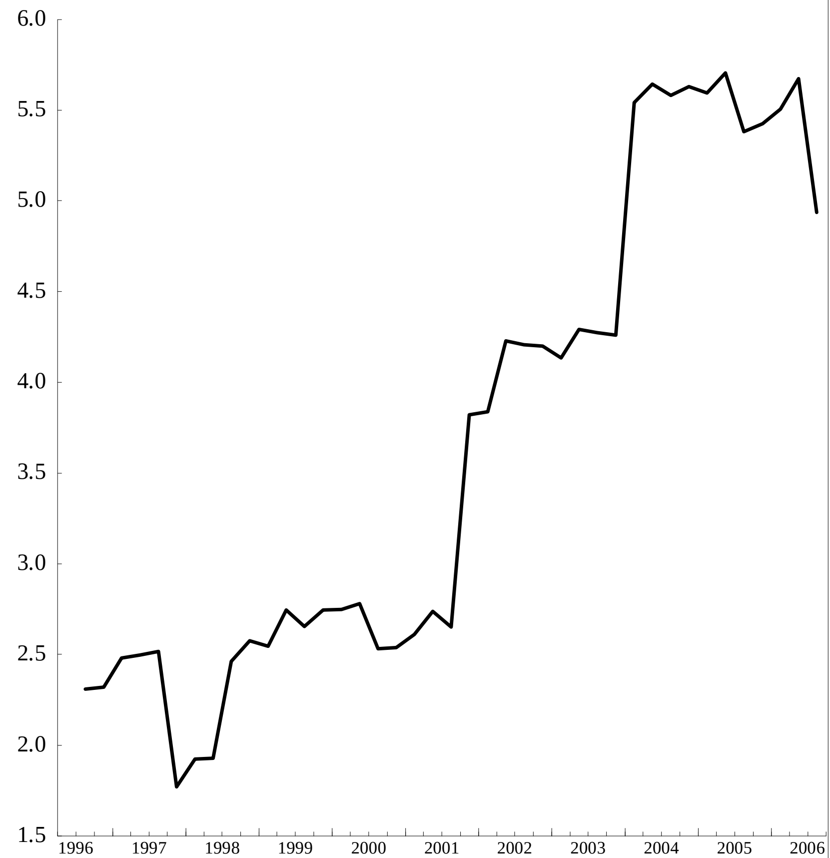

(Cumulative increase in unemployment over 5 years, divided by the change in inflation) sacrifice ratio rose, not quite monotonically, so that it now stands at approximately 5. This result is not an idiosyncrasy of the FRB/US model; one can get qualitatively similar findings from rolling regressions of simple reduced-form Phillips curves with either *ex post* or real-time data. It is, moreover, a finding that arises in other countries as well, as Murchison (2007) acknowledges for the case of Canada.

One may question whether the increase in the sacrifice ratio is real, or evidence of a fundamental misspecification; for present purposes, it does not matter. What matters is models that the Fed staff as policy advisors might have taken to be true, would have turned out, with the benefit of hindsight, to be suspect. In other words, the real-time model uncertainty problem is a real and an important one.

This observation leads directly to another: even as inflation has become less volatile, the proportion of fluctuations in inflation that is predictable has declined; see, e.g., Stock and Watson (2007) and Tulip (2005). Indeed, this is arguably what is being captured by the rising sacrifice ratio; a declining influence of "excess demand" on movements of inflation. A practical consequence of this is that the Bank of Canada—to cite one example—finds that because inflation in Canada has become so stable, not only have traditional reduced-form Phillips curves broken down, but private-sector forecasts of inflation more-or-less mirror what the Bank of Canada shows in its *Monetary Policy Report*.[9] Neither the data, nor private sector expectations render much in the way of useful information for central banks to evaluate their models in real time. Can such a self-confirming equilibrium persist? Or will central banks experience periodic escapes from low-inflation equilibria leading to another Great Inflation?[10] One can make an argument either way. What is hard, I think, to argue against is the case for devising policies that can stabilize the target variables of central banks without extinguishing the information content of the data. The same engineering-cum-control literature that BDNR cull provides some tantalizing hints on how this might be approached, although here again the application to economics is not for the faint of heart.

of inflation the FOMC might like to establish. Yellen is now president of the Federal Reserve Bank of San Francisco, and a FOMC participant.

[9] Murchison (2007) summarizes recent empirical evidence on time variation in the Canadian inflation process. He notes that core CPI inflation in Canada is now well described as white noise.

[10] See, e.g, Sargent,Williams and Zha (2006).

I hope these authors—and others—take up the challenge.

References

Brock, William A., and Steven N. Durlauf (2005) "Local Robustness Analysis: Theory and Application" 29,*Journal of Economic Dynamics and Control*,11 2067-2092.

Brock, William A. and Steven N. Durlauf (2004) "Elements of a Theory of Design Limits to Optimal Policy" 72,*Manchester School* (supplement)1-18

Brock, William A., Steven N. Durlauf and Giacomo Rondina (2006) "Design Limits and Dynamic Policy Analysis" unpublished manuscript, Department of Economics, University of Wisconsin, Madison.

Brock, William A., Steven N. Durlauf, James Nason and Giacomo Rondina (2006) "Simple versus Optimal Rules as Guides to Policy" paper presented to the *Carnegie-Rochester Conference on Public Policy*, November 10-11, 2006, Carnegie-Mellon University, Pittsburgh, PA. forthcoming in *Journal of Monetary Economics*.

Brock, William A., Steven N. Durlauf, and Kenneth D. West (2003) "Policy Evaluation in Uncertain Economic Environments" *Brookings Papers on Economic Activity*,1 235-301.

Brock, William A.; Durlauf, Steven N. and Kenneth D. West (2004) "Model Uncertainty and Policy Evaluation: some theory and empirics" Journal of Econometrics (forthcoming).

Dahleh, Munther A. and Ignacio J. Diaz-Bobillo (1995) *Control of Uncertain Systems: a Linear Programming Approach* (Englewood Cliffs, NJ: Prentice Hall)

Giannoni, Marc .P (2002) "Does Model Uncertainty Justify Caution?: robust optimal monetary policy in a forward-looking model" 6,*Macroeconomic Dynamics,*1: 111-144.

Gilboa, Itzhak and David Schmeidler (1989) "Maxmin Expected Utility with Nonunique Prior" 18,*Journal of Mathematical Economics*,2: 141-153.

Hansen, Lars.P and Thomas J.Sargent (2006) *Robustness* (Princeton Princeton University Press) forthcoming.

Lucas, Robert E. (1987) *Models of Business Cycle Theory* Yrjo Jahnsson Lecture Series (London: Blackwell).

Murchison, Stephen (2007) "The Changing Nature of Inflation in Canada: Recent Bank of Canada Research" paper prepared for the Autumn Central Bank Economists' Meeting at the Bank for International Settlements, October 27-28 , 2005, forthcoming in Bank of Canada Discussion Series.

Mussa, Michael (2003) "Asset Prices and Monetary Policy" in Willam C. Hunter, George G. Kaufman and Michael Pomerleano (eds.) *Asset Price Bubbles: the implications for monetary, regulatory and international policies* (Cambridge: MIT Press).

Onatski, Alexei (2003) "Robust Monetary Policy under Model Uncertainty: incorporating rational expectations" unpublished manuscript, Columbia University.

Onatski, Alexei and James H. Stock (2002) "Robust Monetary Under Uncertainty in a Small Model of the U.S. Economy" 6,*Macroeconomic Dynamics*,1: 85-110.

Onatski, Alexei and Noah Williams (2003) "Modeling Model Uncertainty" *Journal of the European Economic Association*,1: 1087-1122.

Sargent, Thomas J. (1999) "Comment" in John B.Taylor (ed.) Monetary Policy Rules (Chicago: University of Chicago Press and NBER), pp. 144-154.

Sargent, Thomas; Noah Williams and Tao Zha (2005) "Shocks and Government Beliefs: the rise and fall of American Inflation" 96,*American Economic Review*,4 (September), pp. 1193-1224.

Stock, James H. and Mark Watson (2006) "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*,39: 3-33.

Svensson, Lars.E.O. and Robert Tetlow (2005) "Optimum Policy Projections" *International Journal of Central Banking*,1: 177-207.

Tetlow, Robert J. and Brian Ironside (2005) "Real-time Model Uncertainty in the United States: the Fed, 1996-2003" CEPR working paper no. 5305 (November) http://www.cepr.org/pubs/dps/DP5305.asp forthcoming in *Journal of Money, Credit and Banking*.

Tetlow. Robert and Peter 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: 911-949.

Tetlow, Robert and Peter von zur Muehlen (2005) "Robustifying Learnability" Finance and Economics Discussion Series paper no. 2005-58 (August) http://www.federalreserve.gov/pubs/feds/2005/200558/200558pap.pdf

Tulip, Peter (2005) "Has Output Become More Predictable? Changes in Greenbook Forecast Accuracy" Finance and Economics Discussion Series paper no. 2005-31, Federal Reserve Board, (August), http://www.federalreserve.gov/pubs/feds/2005/200531/200531pap.pdf

Zhou, Kemin, John C. Doyle and Keith Glover (1996) *Robust and Optimal Control* (Upper Saddle River, NJ: Prentice Hall).