A Defence of the FOMC*

Martin Ellison
University of Oxford

Thomas J. Sargent
New York University

July 6, 2010

Abstract

We defend the forecasting performance of the FOMC against a criticism of Christina and David Romer (2008) by assuming that the FOMC's forecasts depict a worst-case scenario that it uses to design decisions that are robust to misspecification of the staff's model. We use a simple macro model and a plausible loss function to illustrate how such an interpretation of the FOMC's forecasts can explain the findings of Romer and Romer, including the pattern of differences between FOMC forecasts and forecasts published by the staff of the Federal Reserve System in the Greenbook.

Keywords: forecasting, monetary policy, robustness

JEL classification: C53, E52, E58

*An early version of this paper under the title "What questions are staff and FOMC forecasts supposed to answer?" was thoughtfully discussed by James Bullard at the 10th EABCN Workshop on Uncertainty over the Business Cycle, Frankfurt, March 2009. Craig Burnside and Giorgio Primiceri presented very helpful discussions of the paper at, respectively, the Reserve Bank of New Zealand and Center for International Economic Development Conference to mark the 20th Anniversary of Inflation Targeting in New Zealand, Wellington, 2009 and the NBER Monetary Economics Meeting, New York, March 2010. We thank Marco Bassetto, Christian Matthes, Anna Orlik, and David Romer for helpful comments on an earlier draft and Timothy Besley for early discussions about the problem. We thank Ricardo Mayer for sharing his MATLAB routines to calculate robust optimal policies in hidden Markov models, and Christina and David Romer for making available their dataset of FOMC and staff forecasts. Francisco Barillas provided excellent research assistance in the initial stages of the project. Martin Ellison acknowledges the Hoover Institution for its generous hospitality whilst part of this paper was completed.

1 Introduction

Romer and Romer (2008) assert that forecasts made by the Federal Reserve Open Market Committee (FOMC) detract from the policymaking process in the US and that the FOMC should leave forecasting to the staff of the Federal Reserve System. They base their criticism on an econometric comparison of the accuracy of forecasts published by the FOMC in the *Monetary Policy Report* and forecasts published by the staff of the Federal Reserve System in the *Greenbook*. The staff forecasts are available to FOMC members at the time they produce their own forecasts, so together with presumed superior knowledge of their own preferences, the information advantage lies firmly with the FOMC. Despite this, Romer and Romer (2008) find that:

1. Optimal predictions of inflation and unemployment essentially put zero weight on FOMC forecasts and unit weight on staff forecasts.
2. Staff forecasts have smaller mean squared forecast errors than FOMC forecasts.
3. Statistical and narrative evidence indicates that differences between FOMC and staff forecasts affect actual policy outcomes.

Romer and Romer (2008) use these findings to portray the FOMC as a policymaker that is "not using the information in the staff forecasts effectively" and that "may indeed act on information that is of little or negative value". In their opinion the evidence is sufficiently damning to warrant restructuring the role of the FOMC in monetary policymaking:

> "a more effective division of labor within the Federal Reserve System might be for the staff to present policymakers with policy options and related forecast outcomes, and for policymakers to take those forecasts as given. With this division, the role of the FOMC would be to choose among the suggested alternatives, not to debate the likely outcome of a given policy."

Against these criticisms we provide two defences, one simple, another sophisticated. Our simple defence is based on Figure 1, which plots the FOMC and staff Greenbook forecasts studied by Romer and Romer (2008). In almost all instances, differences between comparable forecasts are miniscule. The only time when notable differences perhaps occur is during the Volcker disinflation of the early 1980s. Indeed, FOMC and staff inflation forecasts never differ by more than 50 basis points after the Volcker disinflation, and there is no significant difference in the accuracies of FOMC and staff forecasts throughout Greenspan's tenure as President of the Federal Reserve. An observer could easily come away from this figure with the impression that, while differences between the forecasts are as statistically significant as Romer and Romer (2008) claim, they are not of enough substantial economic importance to merit concern or discussion.

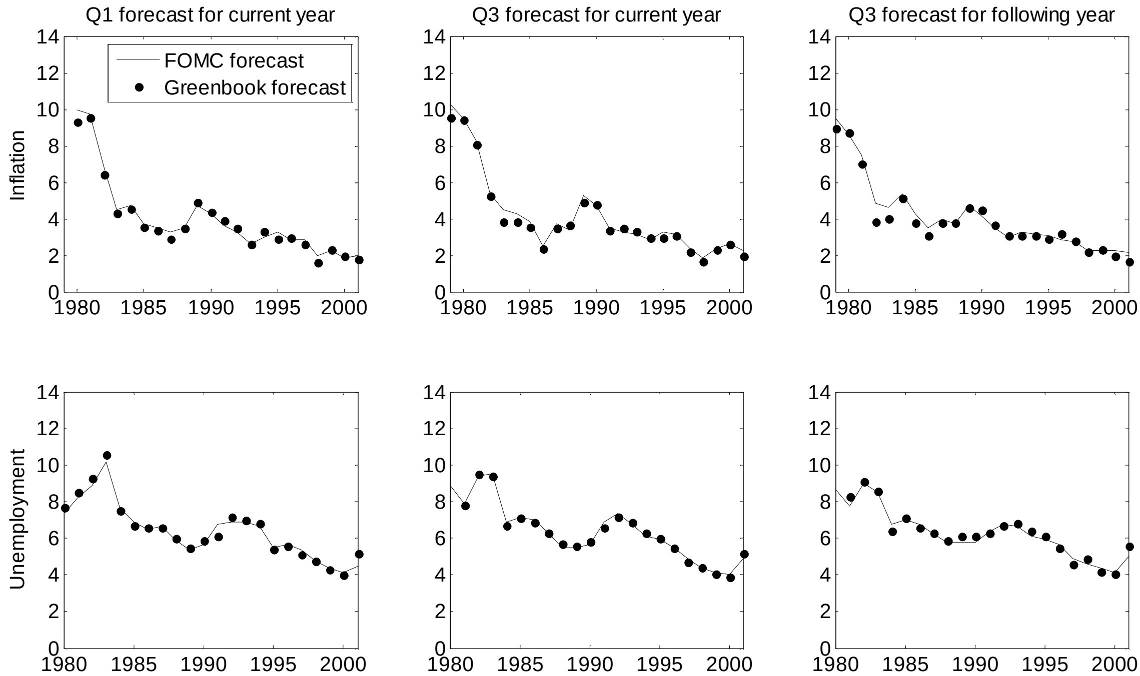

Figure 1: FOMC forecasts and Greenbook forecasts

The remainder of this paper is devoted to our second and more sophisticated defence of the FOMC. Instead of simply dismissing the forecast differences detected by Romer and Romer (2008) as being too small to be of any substantial interest, our second defence sets aside such admittedly understandable impressions gleaned from Figure 1, takes the discrepancies identified by Romer and Romer's statistical findings as something substantial to be accounted for, and fashions an argument based on an assumption that the FOMC and the staff Greenbook forecasts are *answers to different questions*. We pursue this second defence despite Figure 1 because the *pessimism* that Romer and Romer (2008) detect lurking in the FOMC forecasts invites invoking the instrumental pessimism affiliated with robust control theory to explain the qualitative pattern of forecast differences. Romer and Romer's findings provide a laboratory within which we can explore the details of the context-specific pessimism that facilitates constructing robust decisions.

Qualitatively, the criticisms made by Romer and Romer (2008) are unassailable in a world where private agents, policymakers, and researchers share a common model of the probability densities governing economic outcomes. In such a context it is difficult to justify inferior FOMC forecasts. Our defence therefore rests partly on breaking the single probability density assumption by allowing the FOMC to doubt the specification of the model used by the staff to produce its forecasts.[1] In the alternative vision of policymaking adopted in this paper, the staff uses a good macroeconomic model to forecast. But the FOMC suspects that the staff's model is imperfect and wants policies that are robust to specification errors. Those doubts get expressed in the FOMC's forecasts and in its policy choice. To complete our explanation, we assume that the data generating process is in fact described by the staff's model, so that the FOMC's specification fears are, after the fact, 'all in its head'. By assuming that the staff model agrees with the data generating, we drive the forecast error differences in the staff's favour.[2]

To engineer a robust decision rule, the policymaker "exponentially twists" what in our application is the forecasting density of the staff and puts larger probabilities on outcomes that involve inflation and unemployment being away from their targets. Twisting the staff's forecasting probability density in this way results in worst-case scenarios that differ in their severity according to the magnitude of the specification doubts that policymakers wishes to guard against. These worst-case scenarios are a key input to the robust policymaking decision process, because by responding optimally to *them* a policy maker acquires acceptable performance of its loss function when evaluated under each of a large set of possible models, not just under the staff model.[3] Our defence of the FOMC assumes that the forecasts it publishes are exactly these worst-case scenarios.

Furthermore, if the FOMC undertakes robust control, then the worst-case scenarios that it publishes will definitely influence the policy actions actually taken; it is precisely when the worst-case scenario differs from the staff forecast that robust policy calls for the FOMC to take preemptive policy steps. The finding of Romer and Romer (2008) that forecast differences predict monetary policy

[1]A very incomplete list of examples drawn from the growing body of work in macroeconomics that incorporates concerns about robustness of decisions to model misspecifications includes Barlevy (2009), Benigno and Nisticò (2009), Billi (2009), Brock and Durlauf (2005), Brock, Durlauf, and Rondina (2008), Brock, Durlauf, Nason, and Rondina (2007), Caballero and Krishnamurthy (2008), Caballero and Kurlat (2009), Dennis, Leitemo and Söderström (2009a, 2009b), Epstein and Schneider (2008), Giannoni (2002, 2007), Kasa (2002), Leitemo and Söderström (2008a, 2008b), Lewis and Whiteman (2006), Luo and Young (2010), Onatski and Stock (2002), Tetlow (2007), Tetlow and Ironside (2005), Tetlow and von zur Muehlen (2001, 2004, 2009), and Uhlig (2009). See Kasa (2001, 2006) for work focusing on robust filtering and prediction. All of these papers feature some version of max-min expected utility theory, although they differ in terms of the details of how misspecifications are formalised. See Woodford (2006, 2010) for an analysis in which a monetary policy maker trusts its own model but not its knowledge of private agents' expectations. In Woodford's model, equilibrium prices do not reveal private agents' beliefs to the Ramsey planner, while in the model of Karantounias (2009) they do, and that affects the robust policy design problem.

[2]In empirical applications of all models with multiple priors, it is necessary to make *some* assumption about which distribution generates the data. The equating of all subjective distributions with the objective distribution inherent in typical rational expectations models relieves the rational expectations econometrician of the need to make such an independent assumption about the data generating mechanism.

[3]See the ex post Bayesian interpretation of robust decision rules advanced by Hansen and Sargent (2008, chapters 1 and 7).

actions is, therefore, completely compatible with thoughtful application of robust control techniques by the FOMC. In our view, the FOMC may be setting policy rationally to guard against model misspecification by responding to worst-case scenarios rather than reacting to "information that has little or negative value". Our case defending the FOMC rests on whether specification doubts of this type really can quantitatively explain the differences in forecasts published by the FOMC and the staff.

To add content to our defence, we demonstrate how a concern about robustness can explain the broad qualitative features of the results of Romer and Romer (2008) in a simple model of US monetary policy inspired by Primiceri (2006). In our model, the FOMC faces a joint estimation and optimisation problem as it attempts to set appropriate policy whilst simultaneously tracks a hidden state variable in the form of a time-varying NAIRU (Non-Accelerating Inflation Rate of Unemployment). Allowing the policymaker to have specification doubts puts our model within the general class of hidden Markov models discussed by Hansen and Sargent (2007) and Hansen et al. (2010). Accordingly, our policymaker faces a policy design problem in which it doubts not only its model *per se* but also how to use its model to infer the NAIRU. Once the model has been set up, it is straightforward to apply the techniques in Hansen et al. (2010) to show that the type of evidence used by Romer and Romer (2008) to criticise the FOMC can be explained by the assumption that the FOMC answering forecasting questions by reporting the worst-case scenarios that it uses to design a robust policy.

For our defence of the FOMC to be convincing, it should be that the forecasts published by the FOMC are systematically biased towards a worst-case scenario in comparison to the staff forecasts. Romer and Romer (2008) found that FOMC forecasts were on average higher than staff forecasts for inflation and lower than staff forecasts for unemployment. At first sight, the optimistic bias in unemployment forecasts appears at odds with our claim that the FOMC forecasts are worst-case scenarios. However, what constitutes a worst-case scenario is dependent on the staff's approximating model as well as on the policy maker's objective function. In our model, we find that, in plausible regions of the parameter space, the worst-case scenario biases the inflation forecast upwards away from its target and the unemployment forecast downwards towards zero. The intuition for this outcome lies in the dynamics of the model and the way in which the combination of high inflation and low unemployment at the forecast horizon considered by the FOMC signals persistently bad outcomes into the more distant future. We therefore argue that seemingly "pessimistic" inflation forecasts and "optimistic" unemployment forecasts can still be rationalised as a description of the worst-case scenario. Furthermore, the region of the parameter space where this happens is consistent with the FOMC being more concerned about model misspecifications that lead to poor quality inferences about the current state of the economy rather than misspecifications that lead to poor understanding of the state transition dynamics of the system. According to Bullard (2009), it is the concern for accurate tracking of the economy - not a concern for accurate forecasting - that is uppermost in the minds of FOMC members.

The paper is organised as follows. Section 2 describes a simple model of monetary policymaking and shows how it maps into the more general class of hidden Markov models analysed by Hansen et al. (2010). After Section 3 describes the policy maker's filtering problem, the robust policy is derived in Section 4 and a calibrated numerical example is constructed in Section 5 to show how the findings of Romer and Romer (2008) can be explained as an artefact of the FOMC exhibiting a preference for robustness. Section 6 demonstrates how our story can also justify the relative degree of optimism and pessimism seen in actual FOMC and staff forecasts once the FOMC is assumed to be more concerned about tracking hidden states than dynamics of the system *given* those hidden states. A final Section 7 concludes with a discussion of the narrative evidence in support of a division of labour in the Federal Reserve System whereby FOMC forecasts become twists of the staff forecasts.

2 A Model of Monetary Policymaking

The *modus operandi* of monetary policymakers is obscure to outsiders and has to be inferred from the speeches and decisions made by central bankers. Nevertheless, a consensus has emerged that modern monetary policymaking incorporates three key beliefs. First, there is a natural rate of unemployment at which inflation is stable. Second, there is a transmission mechanism through which monetary policy actions affect the economy. Third, monetary policymakers face trade-offs. Indeed, the Financial Times ("King backs job losses to curb inflation") imputed such beliefs to Bank of England Governor Mervyn King on Tuesday April 1, 2008. The FT wrote:

> "The economy needs to slow to the point where there is spare capacity in order to bring inflation under control, Mervyn King, the Bank of England governor, said on Monday. ... Mr King's recognition that the Bank's monetary stance was designed to slow the economy to reinforce its monetary policy committee's inflation-fighting credentials came at an awkward time, he conceded, describing the 'difficult balancing act'."

We adopt a simple model of monetary policymaking that is designed to capture these features in a parsimonious way. The model shares much of the structure proposed by Primiceri (2006) but features an unobserved non-accelerating inflation rate of unemployment (NAIRU) that confronts the monetary policymaker with a joint estimation and decision problem. We assume that the policymaker has an *approximating model* in which the NAIRU u_t^* evolves as:

$$(u_{t+1}^* - u^{**}) = \delta(u_t^* - u^{**}) + \eta_{t+1},$$

where η_{t+1} is an i.i.d. mean zero Gaussian shock and $\left[u_0^* \; u_{-1}^*\right]' \sim \mathcal{N}(\mu(u^*), \Sigma(u^*))$. u^{**} is the steady-state value of the NAIRU. The approximating model further describes inflation π_t and unemployment

U_t as related to the NAIRU and a policy instrument V_t by:

$$\begin{aligned}
\pi_{t+1} &= \pi_t + \gamma_0(U_t - u_t^*) + \gamma_1(U_{t-1} - u_{t-1}^*) + \varepsilon_{\pi t+1}, \\
(U_{t+1} - u_{t+1}^*) &= \rho_1(U_t - u_t^*) + \rho_2(U_{t-1} - u_{t-1}^*) + V_t + \varepsilon_{Ut+1},
\end{aligned}$$

with $\varepsilon_{\pi t+1}$ and ε_{Ut+1} i.i.d. mean zero Gaussian shocks to inflation and unemployment, respectively. The shocks $\left(\eta_{t+1}, \varepsilon_{\pi t+1}, \varepsilon_{Ut+1}\right)$ have standard deviations (c_π, c_U, c_{u^*}). The monetary policymaker's objective is the expected value of:

$$-.5\sum_{t=0}^{\infty}\beta^t \left((\pi_t - \pi^*)^2 + \lambda(U_t - ku_t^*)^2 + \phi(V_t - V_{t-1})^2\right)$$

where λ is the weight placed on unemployment and $k \in (0,1)$ controls whether the policymaker dislikes unemployment or the gap between unemployment and the NAIRU. The parameter ϕ measures the preference for policy smoothing. The monetary policymaker's signal vector at time $t+1$ is:

$$\begin{bmatrix} \tilde{\pi}_{t+1} \\ \tilde{U}_{t+1} \end{bmatrix} = \begin{bmatrix} \pi_{t+1} \\ U_{t+1} \end{bmatrix} + \begin{bmatrix} \vartheta_{\pi t+1} \\ \vartheta_{Ut+1} \end{bmatrix},$$

i.e., it observes noisy measures of inflation and unemployment. The standard deviations of the measurement errors $(\vartheta_{\pi t+1}, \vartheta_{Ut+1})$ are $\left(c_{\tilde{\pi}}, c_{\tilde{U}}\right)$. To map the model into the hidden Markov models of Hansen et al. (2010), let the observed state vector y_t, the unobserved state vector z_t, the control a_t, the shock vector $w_t \sim \mathcal{N}(0, I_5)$, and the signal vector s_t be given by:

$$y_t = \begin{bmatrix} 1 \\ V_{t-1} \end{bmatrix}, \quad z_t = \begin{bmatrix} \pi_t \\ U_t \\ U_{t-1} - u_{t-1}^* \\ u_t^* \end{bmatrix}, \quad a_t = V_t, \quad w_t = \begin{bmatrix} \eta_t \\ \varepsilon_{\pi t} \\ \varepsilon_{Ut} \\ \vartheta_{\pi t} \\ \vartheta_{Ut} \end{bmatrix}, \quad s_t = \begin{bmatrix} \tilde{\pi}_t \\ \tilde{U}_t \end{bmatrix},$$

and write the matrices in the laws of motion and signal equation:

$$\begin{aligned}
y_{t+1} &= A_{11}y_t + A_{12}z_t + B_1 a_t + C_1 w_{t+1}, \\
z_{t+1} &= A_{21}y_t + A_{22}z_t + B_2 a_t + C_2 w_{t+1}, \\
s_{t+1} &= D_1 y_t + D_2 z_t + H a_t + G w_{t+1}
\end{aligned} \tag{1}$$

as

$$A_{11} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad A_{12} = 0_{2\times 4}, \quad B_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad C_1 = 0_{2\times 5},$$

$$A_{21}=\begin{bmatrix}0 & 0\\ (1-\delta)u^{**} & 0\\ 0 & 0\\ (1-\delta)u^{**} & 0\end{bmatrix},\quad A_{22}=\begin{bmatrix}1 & \gamma_0 & \gamma_1 & -\gamma_0\\ 0 & \rho_1 & \rho_2 & (\delta-\rho_1)\\ 0 & 1 & 0 & -1\\ 0 & 0 & 0 & \delta\end{bmatrix},\quad B_2=\begin{bmatrix}0\\1\\0\\0\end{bmatrix},$$

$$C_2=\begin{bmatrix}0 & c_\pi & 0 & 0 & 0\\ c_{u^*} & 0 & c_U & 0 & 0\\ 0 & 0 & 0 & 0 & 0\\ c_{u^*} & 0 & 0 & 0 & 0\end{bmatrix},\quad D_1=\begin{bmatrix}0 & 0\\ (1-\delta)u^{**} & 0\end{bmatrix},\quad D_2=\begin{bmatrix}1 & \gamma_0 & \gamma_1 & -\gamma_0\\ 0 & \rho_1 & \rho_2 & (\delta-\rho_1)\end{bmatrix},$$

$$H=\begin{bmatrix}0\\1\end{bmatrix},\quad G=\begin{bmatrix}0 & c_\pi & 0 & c_{\tilde{\pi}} & 0\\ c_{u^*} & 0 & c_U & 0 & c_{\tilde{U}}\end{bmatrix},$$

where C_1, C_2, G are indicator matrices that pick out the required elements of w_{t+1}. Notice that the matrices in the systematic parts of the signal equation are identical to those in the equation of the unobserved state vector z_{t+1}. The quadratic form in the Hansen et al. (2010) objective function:

$$-\frac{1}{2}\sum_{t=0}^{\infty}\beta^t\begin{bmatrix}a_t\\ y_t\\ z_t\end{bmatrix}'\begin{bmatrix}Q & P_1 & P_2\\ P_1' & R_{11} & R_{12}\\ P_2' & R_{21} & R_{22}\end{bmatrix}\begin{bmatrix}a_t\\ y_t\\ z_t\end{bmatrix}$$

can be expressed as:

$$-\frac{1}{2}\sum_{t=0}^{\infty}\beta^t\begin{bmatrix}V_t\\ 1\\ V_{t-1}\\ \pi_t\\ U_t\\ U_{t-1}-u^*_{t-1}\\ u^*_t\end{bmatrix}'\begin{bmatrix}\phi & 0 & -\phi & 0 & 0 & 0 & 0\\ 0 & \pi^{*2} & 0 & -\pi^* & 0 & 0 & 0\\ -\phi & 0 & \phi & 0 & 0 & 0 & 0\\ 0 & -\pi^* & 0 & 1 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & \lambda & 0 & -\lambda k\\ 0 & 0 & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & -\lambda k & 0 & \lambda k^2\end{bmatrix}\begin{bmatrix}V_t\\ 1\\ V_{t-1}\\ \pi_t\\ U_t\\ U_{t-1}-u^*_{t-1}\\ u^*_t\end{bmatrix}.$$

3 Filtering

Here is how the policymaker would infer the unobserved NAIRU if it had full confidence in the specification of its approximating model. The policymaker observes y_t, has a prior distribution $z_0 \sim \mathcal{N}(\check{z}_0, \Delta_0)$ over the initial values of the unobserved states, and observes a sequence of signals $\{s_{t+1}\}$. With no concerns about model misspecification, the policymaker applies Bayes' law directly to the approximating model (1) to construct a sequence of posterior distributions $z_t \sim \mathcal{N}(\check{z}_t, \Delta_t)$ for $t \geq 1$, where the sufficient statistics $\check{z}_t = E(z_t\,|y_t, \ldots, y_0)$ and $\Delta_t = E\left[(z_t-\check{z}_t)(z_t-\check{z}_t)'\,|y_t, \ldots, y_0\right]$

satisfy the recursions:

$$\begin{aligned}
y_{t+1} &= A_{11}y_t + A_{12}\check{z}_t + B_1 a_t + C_1 w_{t+1} + A_{12}(z_t - \check{z}_t), \\
\check{z}_{t+1} &= A_{21}y_t + A_{22}\check{z}_t + B_2 a_t + K_2(\Delta_t)Gw_{t+1} + K_2(\Delta_t)D_2(\check{z}_t - z_t), \\
\Delta_{t+1} &= \mathcal{C}(\Delta_t).
\end{aligned} \tag{2}$$

$K_2(\Delta_t)$ and (Δ_t) satisfy the Kalman filtering equations:

$$\begin{aligned}
K_2(\Delta) &= (A_{22}\Delta D_2' + C_2 G')(D_2 \Delta D_2' + GG')^{-1}, \\
C(\Delta) &\equiv A_{22}\Delta A_{22}' + C_2 C_2' - K_2(A_{22}\Delta D_2' + C_2 G')',
\end{aligned} \tag{3}$$

and the policymaker's information set at t can be represented as $(y_t, \check{z}_t, \Delta_t)$.

4 Robust Policy and Worst-Case Scenarios

Now we allow the policymaker to doubt the specification of its approximating model. Hansen and Sargent (2007) and Hansen et al. (2010) show how to compute a decision rule for a_t that is robust to possible misspecifications of (i) the approximating model (1) defining the joint distribution of $(y_{t+1}\ z_{t+1})'$ conditional on values of $(y_\tau\ z_\tau)'$ for $\tau \leq t$, and (ii) the probability distribution of the unknown state z_t conditional on the history of signals s_τ for $\tau \leq t$ that comes from applying the ordinary Kalman filter (2)-(3) to the approximating model. Following the steps in Hansen et al. (2010), we begin by letting primes represent next period values to ease notation and noting that the law of motion for $(y, z, \check{z}, \Delta)$ can be written in terms of the state variables $(y, \check{z}, \Delta)$ as:

$$\begin{aligned}
y' &= A_{11}y + A_{12}\check{z} + B_1 a + C_1 w' + A_{12}(z - \check{z}), \\
z' &= A_{21}y + A_{22}\check{z} + B_2 a + C_2 w' + A_{22}(z - \check{z}), \\
\check{z}' &= A_{21}y + A_{22}\check{z} + B_2 a + K_2(\Delta)Gw' + K_2(\Delta)D_2(z - \check{z}), \\
\Delta' &= \mathcal{C}(\Delta),
\end{aligned}$$

where $w' \sim \mathcal{N}(0, I)$ and $z - \check{z} \sim \mathcal{N}(0, \Delta)$.[4] To represent misspecification in the dynamics of the approximating model, the policymaker replaces the distributions of w' and $z - \check{z}$ by distorted distributions $w' \sim \mathcal{N}(\tilde{v}, \Sigma)$ and $z - \check{z} \sim \mathcal{N}(u, \Gamma)$ that potentially feed back on state variables. The distorted distributions allow perturbations to the dynamics of the approximating model. At this point, Hansen et al. (2010) derive robust policy as the outcome of a two-player zero-sum game in which the policymaker chooses an action a to maximise its objective whilst a fictitious agent chooses perturbations w' and $z - \check{z}$ to minimise that same objective. Hansen and Sargent (2007, p. 33) show

[4] Notice how the approximating model includes the law of motion for $(\check{z}, \Delta)$ as dictated by Bayes' law.

that a modified certainty equivalence principle holds in this setting, so instead of directly analysing the full stochastic game it is sufficient to solve a deterministic game in which the policymaker chooses an action a and the minimising agent chooses the *mean* distortions $\tilde{v}$ and u.[5] The appropriate law of motion for the deterministic game has shocks replaced by distorted means:

$$\begin{aligned} y' &= A_{11}y + A_{12}\check{z} + B_1 a + C_1\tilde{v} + A_{12}u, \\ z' &= A_{21}y + A_{22}\check{z} + B_2 a + C_2\tilde{v} + A_{22}u, \\ \check{z}' &= A_{21}y + A_{22}\check{z} + B_2 a + K_2(\Delta)G\tilde{v} + K_2(\Delta)D_2 u, \\ \Delta' &= \mathcal{C}(\Delta), \end{aligned} \tag{4}$$

where $\tilde{v}$ and u are treated as under the control of the minimising agent and allowed to feed back on state variables $(z, \check{z}, \Delta)$. For a quadratic continuation value function $W(y, \check{z}, \Delta, z)$ and one-period return function $\tilde{U}(y, \check{z}, z - \check{z}, a)$, the policymaker chooses an action a and accompanying mean distortions $\tilde{v}$ and u by solving:[6]

$$\max_a \min_u \left[\tilde{U}(y, \check{z}, z - \check{z}, a) + \theta_2 \frac{u'\Delta^{-1}u}{2} + \min_{\tilde{v}} \left(\beta W(y', \check{z}', \Delta', z') + \theta_1 \frac{\tilde{v}'\tilde{v}}{2} \right) \right], \tag{5}$$

where the optimisation is subject to the laws of motion (4) and the minimising agent faces penalties $\theta_1(\tilde{v}'\tilde{v})/2$ and $\theta_2(u'\Delta^{-1}u)/2$ on its choices of $\tilde{v}$ and u.[7] The penalties are on the entropy contributions

[5] This is a consequence of the return function being quadratic, the transition law being linear, and distributions of random shocks and the prior for z_0 being Gaussian. The minimising decision player *increases* shock covariance matrices as well as means, but the certainty equivalence result allows us to compute the mean distortions by solving the purely deterministic game. The omitted stochastic terms affect constants in the value functions but not decision rules.

[6] This is game I of Hansen et. al. (2010), which corresponds to recursions (20)-(21) of Hansen and Sargent (2007). These pertain to a situation in which the decision maker conditions continuation values on hidden state variables.

[7] The continuation value function W has form:

$$W(y, \check{z}, \Delta, z) = -\frac{1}{2} \begin{bmatrix} y \\ z \\ \check{z} \end{bmatrix}' \Omega(\Delta) \begin{bmatrix} y \\ z \\ \check{z} \end{bmatrix} - \omega,$$

and is computed as the fixed-point of:

$$W(y, \check{z}, \Delta, z) = U(y, z, a) + \min_v \left\{ \beta W(y', \check{z}', \Delta', z') + \theta_1 \frac{v'v}{2} \right\},$$

where $'$s denote next period values and the law of motion is modified in the following way to condition on z:

$$\begin{bmatrix} y' \\ z' \\ \check{z}' \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} & 0 \\ A_{21} & A_{22} & 0 \\ A_{21} & K_2(\Delta)D_2 & A_{22} - K_2(\Delta)D_2 \end{bmatrix} \begin{bmatrix} y \\ z \\ \check{z} \end{bmatrix} - \begin{bmatrix} B_1 \\ B_2 \\ B_2 \end{bmatrix} F(\Delta) \begin{bmatrix} y \\ \check{z} \end{bmatrix} + \begin{bmatrix} C_1 \\ C_2 \\ K_2(\Delta)G \end{bmatrix} v,$$

together with:

$$\Delta' = \mathcal{C}(\Delta).$$

Here v is the distorted mean of w' conditioned on $(y, z, \check{z}, \Delta)$, while $\tilde{v}$ is the distorted mean of w' emerging from (5) conditional on $(y, \check{z}, \Delta)$.

$(\tilde{v}'\tilde{v})/2$ and $(u'\Delta^{-1}u)/2$ incurred when the minimising agent distorts the means of w' and z. The degree to which the minimising agent is constrained depends on the positive multipliers θ_1 and θ_2, with lower values giving more scope for the minimising agent to perturb the approximating model. When $\theta_1 = \theta_2 = +\infty$, the policy maker trusts his model. Then the minimizing values of u and $\tilde{v}$ are both zero and problem (5) becomes an ordinary Bellman equation. The lower is the value of θ_1, the more the policymaker distrusts its approximating model of the dynamics of the state, given the current state; the lower is the value of θ_2, the more the policy maker distrusts its current probability distribution over the hidden state z. We discuss the impact of the penalty parameters θ_1 and θ_2 in the context of our model in Section 6.

The max min problem faced by the policymaker can be solved using standard techniques from linear-quadratic control. The solution is for the policymaker to follow a feedback rule:

$$a = -\begin{bmatrix} F_y & F_z \end{bmatrix} \begin{bmatrix} y \\ \check{z} \end{bmatrix}, \tag{6}$$

and for the minimising agent to choose distorted conditional means according to:

$$\tilde{v} = -\begin{bmatrix} K_y & K_z \end{bmatrix} \begin{bmatrix} y \\ \check{z} \end{bmatrix}, \tag{7}$$

$$u = -\begin{bmatrix} L_y & L_z \end{bmatrix} \begin{bmatrix} y \\ \check{z} \end{bmatrix}. \tag{8}$$

The feedback rule (6) prescribes a robust policy for a policymaker concerned that its approximating model may be misspecified. In this paper, we assume that the staff of the Federal Reserve system produces Greenbook forecasts by using the approximating model of the economy. Of course, the staff's forecast have to assume *some* decision rule for the monetary authority. To make its forecast under the approximating model, we endow the staff with the decision rule that the FOMC ultimately chooses, namely, the robust decision rule. Under this interpretation, the staff produce forecasts believing that decisions and outcomes are governed by the first-order vector stochastic difference equation:

$$\begin{bmatrix} y_{t+1} \\ z_{t+1} \\ \check{z}_{t+1} \end{bmatrix} = \begin{bmatrix} A_{11} - B_1 F_y & A_{12} & -B_1 F_z \\ A_{21} - B_2 F_y & A_{22} & -B_2 F_z \\ A_{21} - B_2 F_y & K_2(\Delta) D_2 & A_{22} - B_2 F_z - K_2(\Delta) D_2 L_z \end{bmatrix} \begin{bmatrix} y_t \\ z_t \\ \check{z}_t \end{bmatrix} + \begin{bmatrix} C_1 \\ C_2 \\ K_2(\Delta) G \end{bmatrix} w_{t+1}, \tag{9}$$

and one-period ahead staff forecasts are:

$$\begin{aligned} E\left[y' \,| y, \check{z}\right] &= (A_{11} - B_1 F_y) y + (A_{12} - B_1 F_z)\check{z}, \\ E\left[z' \,| y, \check{z}\right] &= (A_{21} - B_2 F_y) y + (A_{22} - B_2 F_z)\check{z}. \end{aligned} \tag{10}$$

j-period ahead staff forecasts are obtained by iterating forward (10) and raising the appropriate transition matrices to the jth power. The j-period ahead staff forecasts under the approximating model are denoted by:

$$\begin{aligned} E\left[y^j \left| y, \check{z}\right.\right] &= A_{11}(j)y + A_{12}(j)\check{z}, \\ E\left[z^j \left| y, \check{z}\right.\right] &= A_{21}(j)y + A_{22}(j)\check{z}. \end{aligned} \tag{11}$$

The staff forecasts use the approximating model to estimate how the states in the economy will evolve, given the robust policy chosen by the policymaker. However, that robust policy rule is not a 'best response' to the staff's forecasts but instead to the worst-case law of motion associated with the worst-case forecasts implicit in the max min problem described above. To exhibit these worst-case forecasts, we return to the policymaker's max min problem. The forecasts underpinning its solution are worst-case scenarios that explicitly incorporate the actions of the minimising agent, as imagined by the policymaker as a way to cope with specification doubts. The worst-case one-step ahead forecasts that result are twisted by the actions (7)-(8) of the minimising agent:[8]

$$\begin{aligned} \hat{E}\left[y' \left| y, \check{z}\right.\right] &= (A_{11} - B_1F_y - C_1K_y - A_{12}L_y)y + (A_{12} - B_1F_z - C_2K_z - A_{12}L_z)\check{z}, \\ \hat{E}\left[z' \left| y, \check{z}\right.\right] &= (A_{21} - B_2F_y - K_2(\Delta)GK_y - K_2(\Delta)D_2L_y)y \\ &\quad +(A_{22} - B_2F_z - K_2(\Delta)GK_z - K_2(\Delta)D_2L_z)\check{z}, \end{aligned}$$

with multi-step worst-case forecasts given by forward iteration as before:

$$\begin{aligned} \hat{E}\left[y^j \left| y, \check{z}\right.\right] &= \hat{A}_{11}(j)y + \hat{A}_{12}(j)\check{z}, \\ \hat{E}\left[z^j \left| y, \check{z}\right.\right] &= \hat{A}_{21}(j)y + \hat{A}_{22}(j)\check{z}. \end{aligned} \tag{12}$$

We interpret the forecasts published by the FOMC in the Monetary Policy Report to be these worst-case forecasts. The gap between (11) and (12) is our theory of the differential prediction errors analysed by Romer and Romer (2008).

5 A Calibrated Example

Our defence of the FOMC relies on the ability of our model to rationalise the findings of Romer and Romer (2008). In this section we present our arguments via a numerical example. The parameter values for the numerical example are presented in Table 1, which with two exceptions are taken

[8] In addition to twisting forecast means, the minimizing agent also increases conditional variances as described in footnote 5. That implies that worst-case 'fan charts' would be wider than those produced under the approximating model. We do not pursue this observation about fan charts here because we think that to do so in an informative way would require expanding the features about which there is uncertainty by treating the coefficients in our state space model (1) as hidden state variables too.

from the OLS estimates in the third column of Table 1 in Primiceri (2006). The first exception is that we set a lower value of δ and a higher value of c_u, so fluctuations in the NAIRU play less of a role in determining unemployment.[9] The second exception is that we set a lower k to induce the policymaker to move unemployment away from its estimate of NAIRU. The parameters $c_{\tilde{\pi}}$ and $c_{\tilde{u}}$ calibrating measurement error volatilities have small values. We set the initial prior over $\left(u_0^*, u_{-1}^*\right)$ to have a mean of $(6,6)$ and a covariance equal to the steady state implied by the Kalman filter. We examine the model at first for $\theta_1 = \theta_2 = 200$. In Section 6, we explore how aspects of our results depend on our setting of the parameter θ_1 governing fear of misspecified state transition dynamics and parameter θ_2 governing fear of a misspecified posterior distribution over hidden states.

γ_0	-1.02	β	0.99	θ_2	200
γ_1	0.93	π^*	2	c_{u^*}	$\sqrt{0.02}$
ρ_1	1.756	k	0.2	c_π	1.04
ρ_2	-0.779	λ	1	c_u	1
δ	0.95	ϕ	475	$c_{\tilde{\pi}}$	$\sqrt{0.1}$
u^{**}	6	θ_1	200	$c_{\tilde{u}}$	$\sqrt{0.1}$

Table 1: Parameter values

The behaviour of the calibrated model is illustrated by a representative simulation in Figure 2. The solid line in the first part of each panel is a simulated time path for inflation and unemployment, so at $t = 10$ the economy is just exiting a recession with inflation low and unemployment high. The solid and dotted lines in the second part of each panel are forecasts for $t > 10$ using information available up to and including period $t = 10$. The solid lines, labelled 'staff forecast', are derived under the approximating model (11) and predict rapid returns of inflation and unemployment to their steady-state values. The dotted lines, labelled 'FOMC forecast', are worst-case scenarios (12) in which both inflation and unemployment overshoot and the returns to steady state are prolonged and oscillatory.

[9] The rationale for this will be discussed in footnote 10 of Section 6.

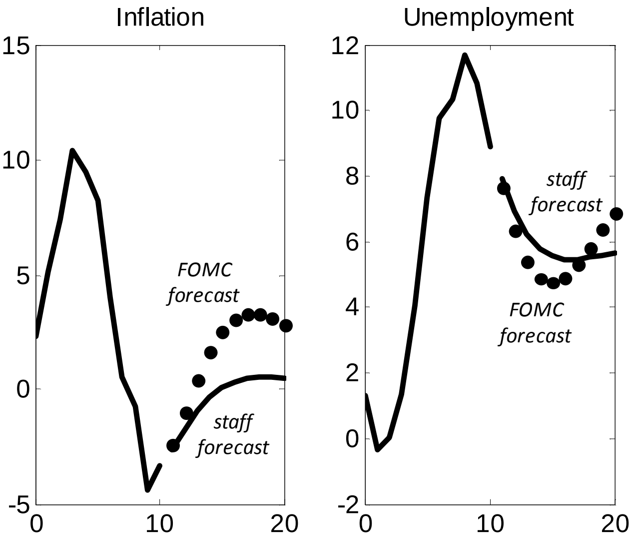

Figure 2: Simulation of the model with $\theta_1 = \theta_2 = 200$; staff and FOMC forecasts at $t = 10$

We now use the numerical example to defend the FOMC against the criticisms made by Romer and Romer (2008). In the introduction, we reported how three specific findings led them to their conclusions, so to refute their claims we show that each finding is consistent with our interpretation of FOMC forecasts as twists of the staff forecasts.

1. *Optimal predictions of inflation and unemployment essentially put zero weight on FOMC forecasts and unit weight on staff forecasts.*

The first finding of Romer and Romer (2008) follows immediately from our interpretation of FOMC forecasts as forecasts under its worst-case model and from our assumption that in fact the staff's model accurately approximates the data generating process. In our model, the FOMC forecasts contain no information over and above that in the staff forecasts under the staff's approximating model, so by definition the optimal predictions of inflation and unemployment under the approximating model should put zero weight on FOMC forecasts and unit weight on staff forecasts. To make this more concrete, we replicate the econometric analysis of Romer and Romer (2008) but with data simulated from our calibrated model. They estimate a regression of the form $X_t = a + bS_t + cP_t + e_t$, where X_t is inflation or unemployment and S_t and P_t are the relevant FOMC and staff forecasts. We simulate the model for $n = 68$ periods and consider FOMC and staff forecasts at the horizon of four periods ahead. Table 2 shows the results; it is directly comparable to Table 1 in Romer and Romer (2008).

	Constant	Staff forecast	FOMC forecast	R^2
Inflation	-0.83 (0.44)	1.22 (0.42)	-0.16 (0.24)	0.38
Unemployment	0.52 (1.28)	1.02 (0.45)	-0.15 (0.26)	0.17

Table 2: Role of staff and FOMC forecasts in predicting actual values for simulated data with $\theta_1 = \theta_2 = 200$

As expected, the optimal predictions of inflation and unemployment put close to zero weight on the FOMC forecast and close to unit weight on the staff forecast. In this particular simulation, there is over-weighting of the staff forecast relative to the FOMC forecasts, a result also obtained and stressed by Romer and Romer (2008).

2. *Staff forecasts have smaller mean squared forecast errors than FOMC forecasts.*

We can explain the second finding of Romer and Romer (2008) as an artefact of a misguided attempt to run a horse race between two forecasts that answer different questions. In our model, the FOMC forecasts come from a worst-case model used for planning purposes that, by construction, will not minimise mean squared forecast errors under the approximating model. With the two forecasts answering different questions, it is inappropriate to measure their performance against a common mean squared error criterion that by definition favours the staff forecast when the approximating model actually comes closer to governing the data. Calculations using simulated data from our calibrated model confirm this; in a representative simulation the mean squared errors of staff forecast are 3.36 for inflation and 2.80 for unemployment, which compare favourably to the mean squared errors of 4.26 for inflation and 3.34 for unemployment made by FOMC forecasts.

3. *There is statistical and narrative evidence to suggest that differences between FOMC and staff forecasts affect actual policy outcomes.*

The statistical evidence that differences in forecasts affect policy decisions is based on the correlation between forecast differences and the Romer and Romer (2004) measure of monetary policy shocks. The main finding in Romer and Romer (2008) is that contractionary monetary policy shocks are associated with the FOMC inflation forecast being above that of the staff. We refute using this finding to criticise the FOMC by pointing out that it obtains in our model. The problem lies in the way Romer and Romer (2004) estimate a series of monetary policy shocks by regressing the intended federal funds rate on Greenbook forecasts to arrive at a "series for monetary policy shocks that should be free of both endogenous and anticipatory actions". The shock series that results is purged of information in the staff forecast but there is no guarantee that it will be exogenous with respect to the FOMC forecast. Quite the contrary, robust policy expressly requires the policymaker to take particular policy actions at times when worst-case forecasts differ from forecasts from the approximating model. The 'shocks' identified by the Romer and Romer (2004) procedure are then by construction correlated with the differences in forecasts. To see this mechanism in action in our model, we apply the Romer and Romer (2004) procedure to simulated data to identify a measure of monetary policy shocks M_t that is orthogonal to the staff forecasts. We then follow Romer and Romer (2008) by estimating a regression of the form $M_t = a + b(S_t - P_t) + e_t$, where $S_t - P_t$ is the difference between FOMC and staff forecasts. The results appear in Table 3 in the same format as those for actual data in Table 2 of Romer and Romer (2008).

Constant	Inflation	Unemployment	R^2
-0.07 (0.02)	0.023 (0.009)		0.08
-0.05 (0.02)		-0.030 (0.015)	0.06

Table 3: Role of forecast differences in predicting monetary policy shocks for simulated data with $\theta_1 = \theta_2 = 200$

The difference between FOMC and staff forecasts correlates with monetary policy shocks in simulated data as expected. The strongest finding is that monetary policy contractions are associated with the FOMC inflation forecast being above that of the staff, which mirrors the results of Romer and Romer (2008). It is worth stressing that this result is a misleading consequence of incorrect identification of monetary policy shocks by the Romer and Romer (2004) procedure. If we identify shocks by regressing policy actions on *both* staff and FOMC forecasts then we are unlikely to find a significant correlation between shocks and forecast differences. For simulated data this is certainly the case.

The narrative evidence presented by Romer and Romer (2008) centres on the transcripts of three FOMC meetings at which policy actions appear to be rationalised by the differences between FOMC and staff forecasts. At the meetings in July 1979 and February 1982, the FOMC inflation forecast was well above that of the staff, and there was a substantial contractionary policy shock. At the meeting in February 1991 the situation was reversed, with the FOMC inflation forecast well below the staff's so Romer and Romer argue that there was a substantial expansionary shock. Even putting aside the difficulty of correctly identifying policy shocks, it is not clear how much weight one should put on the narrative evidence. For example, Romer and Romer (2008) quote Mr Mayo in July 1979 as saying that "Although the staff forecast is a reasonable one, I find myself a little more pessimistic. I am concerned about both the likelihood of less real growth and more inflation" and Mr Boehe as saying in February 1991 that "I think the staff forecast, while well thought out, is on the rosy side ... I'd rather err on the side of too much stimulus at this point than too little". Whilst the interpretation of narrative evidence such as this is debatable, we see both these quotes as consistent with our view that the FOMC forms worst-case scenarios.

6 Average Forecast Differences

The findings of Romer and Romer (2008) are consistent with our interpretation of FOMC forecasts as worst-case scenarios that inform robust policy designed to confront specification doubts. As such, we have already provided a full defence of the FOMC against their criticisms. We can go further though, because our characterisation of policymaking has sharp predictions about the relationship between FOMC and staff forecasts. In particular, the FOMC forecast should be systematically biased towards the worst-case scenario. Romer and Romer (2008) report that the FOMC inflation forecasts is on average 13 basis points above the corresponding staff forecast. For unemployment, the FOMC forecast is on average 6 basis points below that of the staff.[10] At first sight the combination of pessimistic inflation forecasts and optimistic unemployment forecasts appears difficult to reconcile with our idea of FOMC forecasts as worst-case scenarios. However, what constitutes a worst-case scenario depends on the objective function and the approximating model of the policymaker. A worst-case scenario is dynamic and time-varying, so it is generally inappropriate to equate average forecast difference with pessimism or optimism. Instead, we need to return to the model and calculate the average forecast differences it implies at different horizons. We do this now and ask whether our model can explain the differences in forecasts found by Romer and Romer (2008).

To calculate the average forecast differences at any horizon, it is sufficient to apply forecasting equations (11) and (12) to the steady-state values of $y, z, \check{z}$ of the stochastic difference equation for decisions and outcomes (9). It is a feature of models of linear-quadratic-Gaussian robust control that the policymaker's problem can become impossible if its specification doubts specification become too big.[11] In the language of Whittle (2002), the θ's can be pushed to levels that set off a 'neurotic

[10] The average forecast differences in inflation and unemployment are arguably small, as discussed in the context of Figure 1. Furthermore, during the Greenspan era the FOMC inflation forecast is on average only 2 basis points above the staff inflation forecast.

[11] Technically, what happens if the θs are set low enough is that the min-max problem becomes ill-defined because the minimising player has been granted so much power to harm the maximising player that the objective function gets driven to $-\infty$ under any policy that the maximising player can choose. Setting the θs that low amounts to asking for

breakdown' as specification doubts completely overwhelm the ability of the decision maker to choose. To see how much this affects our calibrated model, Figure 3 documents the range of θ_1 and θ_2 values below which breakdown occurs. Any explanation of the average forecast differences in Romer and Romer (2008) needs to avoid calibrations that are in the breakdown region.[12]

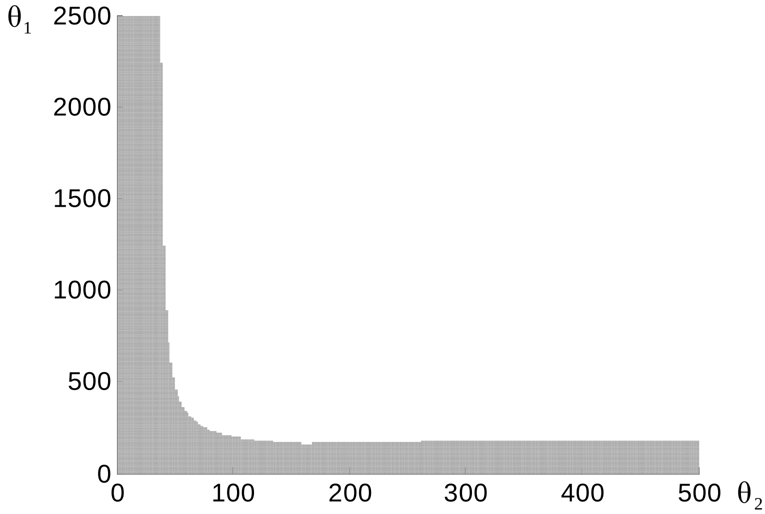

Figure 3: Region of neurotic breakdown

The average difference between FOMC and staff inflation forecasts is shown in Figure 4 as a function of θ_1 and θ_2. The rest of the calibration is as before. Forecast differences are shown via a contour map so, for example, the 0.2 contour depicts the locus of θ_1 and θ_2 values for which the four-period ahead inflation forecast of the FOMC is on average 20 basis points higher than the corresponding staff forecast. According to the figure, the largest forecast differences occur near the region of neurotic breakdown at the point where the policymaker has the largest permissible specification doubts. Conversely, if θ_1 and θ_2 are large then the policymaker has a lot of confidence in its approximating model and the difference between the worst-case scenario and the forecast from the model is small. The contour map shows that FOMC inflation forecasts are on average higher than staff inflation forecasts for any values of θ_1 and θ_2 outside the region of neurotic breakdown. This means it is easy to reconcile our model with the Romer and Romer (2008) finding of FOMC inflation forecasts being on average 13 basis points higher than those of the staff.

robustness over a larger set of models than it is feasible to attain. See Hansen and Sargent (2008, chapter 8) for a discussion of breakdown points.

[12] This requirement explains our decision in Section 5 to calibrate δ such that fluctuations in the NAIRU play less of a role in determining unemployment. If δ is set so that the NAIRU has near unit root behaviour as in Primiceri (2006), then the region of neurotic breakdown is very large and it is difficult for the model to generate noticeable differences in FOMC and staff forecasts.

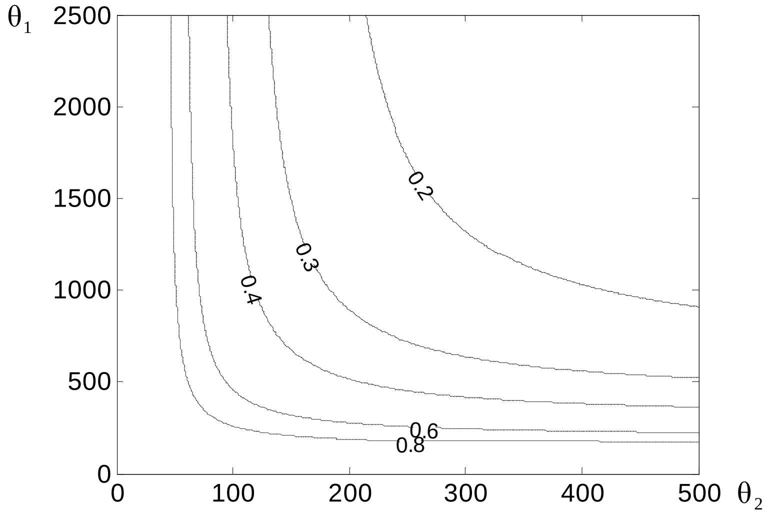

Figure 4: Average difference between FOMC and staff inflation forecasts four periods ahead

The average difference between FOMC and staff unemployment forecasts is shown in Figure 5. This time the contour map is non-monotonic and whether the FOMC forecasts are higher or lower than the staff forecasts depends on the values of θ_1 and θ_2. An early indication that this might be the case was present in Figure 2 where the FOMC unemployment forecast first drops below and then rises above the staff forecast. This is caused by the dynamic nature of the $\max\min$ problem underpinning the worst-case scenario. In the simulated example, the initial low FOMC forecast of unemployment is a valid description of the worst-case because it signals large fluctuations in inflation and unemployment in the future. A similar mechanism drives the average forecast differences in the model, so for some regions of the parameter space, the worst-case scenario for unemployment appears to be 'optimistic' initially and then 'pessimistic'. The non-monotonicity of the contour map means we can rationalise the average optimism of the FOMC in unemployment forecasts with the average pessimism of FOMC inflation forecasts.

Values of $\theta_1 = 76383$ and $\theta_2 = 238$ give a match between our model and the results of Romer and Romer (2008), and represent only a modest concern about misspecification on the part of the FOMC. To see just how modest, we interpret θ_1 and θ_2 using the detection errors suggested by Hansen et al. (2002), Anderson et al. (2003), and Hansen and Sargent (2008, ch. 9). A detection-error probability is the probability that an econometrician makes an incorrect inference about whether observed data is generated by the approximating model or the worst-case model. A detection error probability attains its maximum possible value of .5 when the approximating and worst-case models are identical (i.e., when $\theta_1 = \theta_2 = +\infty$). Detection-error probabilities approach .5 from below when θ_1 and θ_2 both approach $+\infty$ and the worst-case model becomes very close to the approximating model. The detection-error probability in our model is 0.326 in a sample of 100 quarters with $\theta_1 = 76383$ and $\theta_2 = 238$, which indicates that almost a third of the time the detection test would select the wrong model, indicating that the models are difficult to distinguish statistically. Therefore, we interpret our calibrated pair of θ's as implying only modest fears of model misspecification.

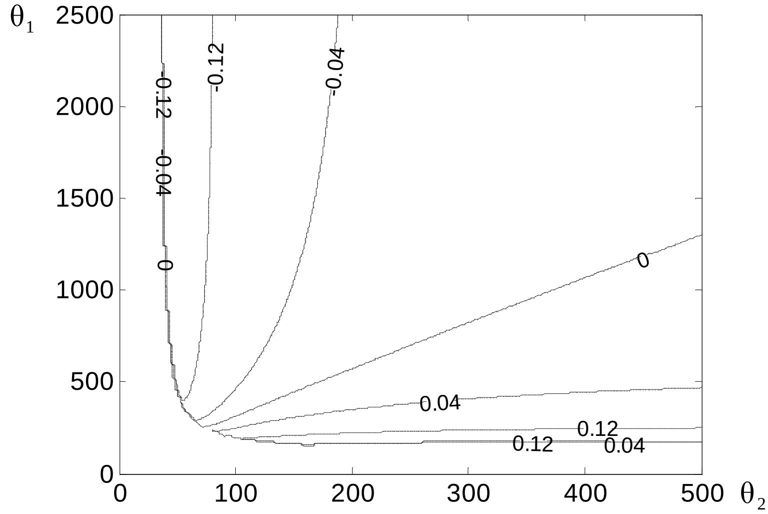

Figure 5: Average difference between FOMC and staff unemployment forecasts four periods ahead

The need to tilt the calibration towards a large θ_1 and a small θ_2 helps identify the nature of specification doubts held by the FOMC. A large value of θ_1 in the min max problem (5) suggests that the FOMC is not overly concerned with the staff's ability to forecast the future state of the economy, provided that it can accurately estimate the current state. This is further apparent when separate detection-error probabilities are constructed for θ_1 and θ_2 following the sequential procedure of Hansen et al. (2010).[13] The detection-error probability for θ_1 is 0.498, which implies that the FOMC has almost no concerns over its ability to forecast the future state of the economy conditional on the current state. Instead, almost all the overall detection-error probability comes from θ_2, which expresses the FOMC's doubting its view of the current state. The small value of θ_2 means that the FOMC *is* concerned about its ability to infer the current state of the economy. In the parlance of monetary policymakers, this translates into the FOMC being more worried about 'tracking' than 'forecasting'. This result resonates with the discussion by President James Bullard of the Federal Reserve Bank of St. Louis (2009) on an earlier version of this paper. In it he claimed that "Forecasting is tracking ... Most of the focus in policy discussion concerns today's state vector ... Further out is normally slow mean reversion ... Through experience, forecasters learned that the near random walk

[13] Hansen et al. (2010) calculate a detection-error probability for θ_1 by pretending that y, z are both observable and then calculating detection-error probabilities for a system with an observed state vector using the approach of Hansen et al. (2002) and Hansen and Sargent (2008, ch. 9).

model works best".

Intuition for why worries about tracking lead the FOMC to make optimistic unemployment forecasts is contained in equations (7) and (8) for the worst-case conditional mean $\tilde{v}$ of the shock vector w_{t+1} of the state and signal dynamics and the worst-case conditional mean u of the hidden state reconstruction error $z - \check{z}$. The lack of concern for forecasting implies that K_y, K_z and $\tilde{v}$ in equation (7) are very small, so the worst case is almost completely driven by the distortion u of equation (8). u is responsible for the optimism in unemployment forecasts already discussed, but implies mild *pessimism* in the FOMC's forecast for the natural rate of unemployment (on average 0.3 basis points of pessimism when $\theta_1 = 76383$ and $\theta_2 = 238$). The worst case is therefore characterised by a pessimistic forecast for the natural rate of unemployment and an optimistic forecast for unemployment. That u distorts unemployment and its natural rate in opposite directions is confirmed in the numerical values of L_y and L_z, which define the distortion $A_{22}u = -A_{22}L_y y - A_{22}L_z\check{z}$ to the state vector z' in equation (4) as:

$$A_{22}u = \begin{bmatrix} -0.0055 & -0.0568 \\ 0.0106 & 0.0695 \\ -0.0063 & 0.0397 \\ 0.0114 & -0.0098 \end{bmatrix} y + \begin{bmatrix} 0.0132 & -0.0424 & 0.0370 & 0.0429 \\ -0.0130 & 0.0505 & -0.0423 & -0.0487 \\ -0.0084 & 0.0301 & -0.0255 & -0.0300 \\ 0.0030 & -0.0081 & 0.0074 & 0.0092 \end{bmatrix} \check{z}.$$

The coefficients on the second and fourth rows respectively determine the worst-case distortions to unemployment and the natural rate of unemployment. They have opposite signs in all but the first column, so unemployment and its natural rate are distorted in opposite directions. Distorting in this way is a feature of the worst case because it makes it more difficult for the policymaker to infer the hidden state variables of the system, in particular the natural rate of unemployment. To see this more clearly, we calculate the variance-covariance matrix of $z - \check{z}$ under the approximating and worst-case models following Hansen et al. (2010):[14]

$$\Delta = \begin{bmatrix} 0.0924 & -0.0011 & -0.0074 & 0.0027 \\ -0.0011 & 0.0929 & 0.0107 & 0.0009 \\ -0.0074 & 0.0107 & 0.2594 & -0.1828 \\ 0.0027 & 0.0009 & -0.1828 & 0.1952 \end{bmatrix},$$

[14] Equation (23) of Hansen et al. (2010) defines the worst-case conditional variance of $z - \check{z}$ by:

$$\Delta_{worst\ case} = \left(\Delta^{-1} - \frac{1}{\theta_2} R_{22} - \frac{1}{\theta_2} \begin{bmatrix} 0 & I & 0 \end{bmatrix} \Omega(\Delta) \begin{bmatrix} 0 \\ I \\ 0 \end{bmatrix} \right)^{-1},$$

provided the matrix on the right hand side is positive definite. $\Omega(\Delta)$ is the matrix in the continuation value function in footnote 7.

$$\Delta_{worst\ case} = \begin{bmatrix} 0.0927 & -0.0016 & -0.0072 & 0.0030 \\ -0.0016 & 0.0959 & 0.0097 & -0.0000 \\ -0.0072 & 0.0097 & 0.2600 & -0.1826 \\ 0.0030 & -0.0000 & -0.1826 & 0.1959 \end{bmatrix}.$$

The worst-case variance-covariance matrix has larger values along its leading diagonal, reflecting the greater difficulty the policymaker has inferring all the hidden state variables. Furthermore, distorting unemployment and the natural rate of unemployment in opposite directions creates negative covariance between the inference errors for these variables (element (2,4) is mildly negative in the worst case). From this we have the final intuition for why the worst-case unemployment forecast is optimistic. In the worst case, the natural rate of unemployment is distorted pessimistically whilst unemployment itself is distorted optimistically. With unemployment and its natural rate moving in opposite directions, it becomes harder for the policymaker to infer the hidden state vector z, the variance-covariance matrix of $z - \check{z}$ deteriorates, and it is more difficult for the policymaker to control the economy.

7 Conclusions

In our story, contrary to Romer and Romer (2008), policymakers do "use the information in the staff forecasts effectively" and do not "act on information that is of little or negative value". The model with specification doubts is consistent with all the findings of Romer and Romer (2008) and interprets the average difference between forecasts as what we regard as a reasonable response of the FOMC to doubts about the specification of its model.

Whether individuals within the Federal Reserve System see themselves as dividing up policymaking tasks in the way we propose is open to debate. The cautious communications strategies adopted by policymakers make it unlikely that they will openly describe their forecasts in terms of adjustments that express their doubts about their model. Nevertheless, some policymakers have gone on the record with arguments that support our view. For example, on 4th January 2008 Forbes Magazine ("Kohn says Fed operating with diverse views, not just strong chairman") reported on a discussion of Romer and Romer (2008) given by former Federal Reserve Monetary Affairs Director Vincent Reinhart at the American Economic Association meetings in New Orleans. Forbes wrote:

> "However, former Fed staffer Vincent Reinhart said while it may look as if 'the FOMC's contribution to the monetary policy process is to reduce forecast accuracy,' they are not there primarily to be forecasters. Instead, they exist in a political system and have to be held accountable for the outcomes of their decisions. 'They can be bad forecasters and good policymakers,' Reinhart said, 'if the diversity of views about the outlook informs

their policy choice.'"

This paper also contributes to a recent literature due to Elliott et al. (2006) that uses forecast biases to identify the objective functions of the people making those forecasts. For example, in this vein Capistrán (2008) argues that the systematic biases periodically appearing in Greenbook forecasts can be rationalised by assuming that the staff of the Federal Reserve System has a time-varying and asymmetric forecasting objective. Our approach is arguably more disciplined than this because the objective function that we posit has a structural interpretation as expressing a decision maker's doubts about the specification of its model. We believe that our approach can deliver insights into the mindset that policymakers do or should have. Another exciting development is the dataset of individual FOMC member forecasts put together by Romer (2009) and recently analysed by Bhattacharjee and Gelain (2010). This is likely to prove a rich seam for future research.

References

[1] Anderson, E., Hansen, L. and T.J. Sargent, 2003, A quartet of semigroups for model specification, robustness, prices of risk, and model detection, *Journal of the European Economic Association* 1(1), 68–123.

[2] Barlevy, G., 2009, Policymaking under uncertainty: Gradualism and robustness, *Economic Perspectives* QII, Federal Reserve Bank of Chicago, 38-55.

[3] Benigno, P. and S. Nisticò, 2009, International Portfolio Allocation under Model Uncertainty, *NBER Working Paper* No. 14734.

[4] Bhattacharjee, A. and P. Gelain, 2010, Monetary Policy under Model Uncertainty: A Bayesian Analysis of FOMC Forecasts, *mimeo*, University of St Andrews

[5] Billi, R.M., 2009, Optimal Inflation for the U.S. Economy, *Federal Reserve Bank of Kansas City Working Paper* No. 07-03 (revised).

[6] Brock, W.A. and S.N. Durlauf, 2005, Local robustness analysis: Theory and application, *Journal of Economic Dynamics and Control* 29, 2067-92.

[7] Brock, W.A., Durlauf, S.N. and G. Rondina, 2008, Frequency-Specific Effects of Stabilization Policies, *American Economic Review* 98, 241-5.

[8] Brock, W.A., Durlauf, S.N., Nason, J.M. and G. Rondina, 2007, Simple versus optimal rules as a guide to policy, *Journal of Monetary Economics* 54, 1372-96.

[9] Bullard, J., 2009, Discussion of Ellison and Sargent: What questions are staff and FOMC forecasts supposed to answer?, 10th EABCN Workshop on Uncertainty over the Business Cycle, Frankfurt.

[10] Caballero, R.J. and A. Krishnamurthy, 2008, Collective Risk Management in a Flight to Quality Episode, *Journal of Finance* LXIII, 2195-230.

[11] Caballero, R.J. and P. Kurlat, 2009. The "Surprising" Origin and Nature of Financial Crises: A Macroeconomic Policy Proposal, *mimeo*, Prepared for the Jackson Hole WY Symposium on Financial Stability and Macroeconomic Policy, August 2009.

[12] Capistrán, C., 2008, Bias in Federal Reserve inflation forecasts: Is the Federal Reserve irrational or just cautious?, *Journal of Monetary Economics* 55, 1415-27.

[13] Dennis, R., Leitemo, K and U. Söderström, 2009a, Methods for robust control, *Journal of Economic Dynamics and Control* 33, 1604-16.

[14] Dennis, R., Leitemo, K and U. Söderström, 2009b, Monetary policy in a small open economy with a preference for robustness, *mimeo*, Federal Reserve Bank of San Francisco.

[15] Elliott, G., Komunjer, I., and A. Timmermann, 2005, Estimation and Testing of Forecast Rationality under Flexible Loss, *Review of Economic Studies* 72, 1107–25.

[16] Epstein, L.G. and M. Schneider, 2008, Ambiguity, Information Quality, and Asset Pricing, *Journal of Finance* 63, pp. 197–228.

[17] Giannoni, M.P., 2002, Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy In A Forward-Looking Model, *Macroeconomic Dynamics* 6, 111-44.

[18] Giannoni, M.P., 2007, Robust optimal monetary policy in a forward-looking model with parameter and shock uncertainty, *Journal of Applied Econometrics* 22, 179-213.

[19] Hansen, L. P. and T.J. Sargent, 2007, Recursive robust estimation and control without commitment, *Journal of Economic Theory*, 136, 1-27.

[20] Hansen, L.P. and T.J. Sargent, 2008, *Robustness*, Princeton, New Jersey, Princeton University Press.

[21] Hansen, L.P., Mayer, R. and T.J. Sargent, 2010, Robust Hidden Markov LQG Problems, *Journal of Economic Dynamics and Control*, in press.

[22] Hansen, L.P., Sargent, T.J. and N.E. Wang, 2002, Robust permanent income and pricing with filtering, *Macroeconomic Dynamics* 6, 40–84.

[23] Karantounias, A.G., (with Lars Peter Hansen and Thomas J. Sargent), 2009, Managing expectations and fiscal policy, *mimeo*, Federal Reserve Bank of Atlanta.

[24] Kasa, K., 2001, A robust Hansen-Sargent prediction formula, *Economics Letters* 71, 43-8.

[25] Kasa, K., 2002, Model Uncertainty, Robust Policies, And The Value Of Commitment, *Macroeconomic Dynamics* 6, 145-66.

[26] Kasa, K., 2006, Robustness and Information Processing, *Review of Economic Dynamics* 9, 1-33.

[27] Leitemo, K. and U. Söderström, 2008a, Robust monetary policy in a small open economy, *Journal of Economic Dynamics and Control* 32, 3218-52.

[28] Leitemo, K. and U. Söderström, 2008b, Robust monetary policy in the New Keynesian framework, *Macroeconomic Dynamics* 12, 126-35.

[29] Lewis, K.F. and C.H. Whiteman, 2006, Robustifying Shiller: Do Stock Prices Move Enough To Be Justified By Subsequent Changes In Dividends?, *mimeo*, University of Iowa.

[30] Luo, Y. and E.R. Young, 2010, Risk-sensitive Consumption and Savings under Rational Inattention, *American Economic Journal: Macroeconomics*, forthcoming.

[31] Onatski, A. and J.H. Stock, 2002, Robust Monetary Policy Under Model Uncertainty In A Small Model of the U.S. Economy, *Macroeconomic Dynamics* 6, 85-110.

[32] Primiceri, G.E., 2006, Why Inflation Rose and Fell: Policy-makers' Beliefs and U. S. Postwar Stabilization Policy, *Quarterly Journal of Economics* 121, 867-901.

[33] Romer, D.H., 2009, A New Data Set on Monetary Policy: The Economic Forecasts of Individual Members of the FOMC, *NBER Working Paper* No. 15208.

[34] Romer, C.D. and D.H. Romer, 2004, A New Measure of Monetary Shocks: Derivation and Implications, *American Economic Review* 94, 1055-84.

[35] Romer, C.D. and D.H. Romer, 2008, The FOMC versus the Staff: Where Can Monetary Policymakers Add Value?, *American Economic Review* 98, 230-235.

[36] Tetlow, R.J., 2007, On the robustness of simple and optimal monetary policy rules, *Journal of Monetary Economics* 54, 1397-405.

[37] Tetlow, R.J. and B. Ironside, 2005, Real-Time Model Uncertainty in the United States: The Fed, 1996-2003, *Journal of Money, Credit and Banking* 39, 1533-61.

[38] Tetlow, R.J. 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, 911-49.

[39] Tetlow, R.J. and P. von zur Muehlen, 2004, Avoiding Nash Inflation: Bayesian and Robust Responses to Model Uncertainty, *Review of Economic Dynamics* 7, 869-99.

[40] Tetlow, R.J. and P. von zur Muehlen, 2009, Robustifying Learnability, *Journal of Economic Dynamics and Control* 33, 296-316.

[41] Uhlig, H., 2009, A model of a systemic bank run, *NBER Working Paper* No. 15072.

[42] Whittle, P., 2002, Risk Sensitivity, a Strangely Pervasive Concept, *Macroeconomic Dynamics* 6, 5-18.

[43] Woodford, M., 2006, An Example of Robustly Optimal Monetary Policy with Near-Rational Expectations, *Journal of the European Economic Association* 4, 386-95.

[44] Woodford, M., 2010, Robustly Optimal Monetary Policy With Near-Rational Expectations, *American Economic Review* 100, 274-303.