INSURANCE POLICIES FOR MONETARY POLICY IN THE EURO AREA

Keith Kuester
Federal Reserve Bank of Philadelphia

Volker Wieland
Goethe University of Frankfurt

Abstract
In this paper, we aim to design a monetary policy for the euro area that is robust to the high degree of model uncertainty at the start of monetary union and allows for learning about model probabilities. To this end, we compare and ultimately combine Bayesian and worst-case analysis using four reference models estimated with pre–European Monetary Union (EMU) synthetic data. We start by computing the cost of insurance against model uncertainty, that is, the relative performance of worst-case or minimax policy versus Bayesian policy. While maximum insurance comes at moderate costs, we highlight three shortcomings of this worst-case insurance policy: (i) prior beliefs that would rationalize it from a Bayesian perspective indicate that such insurance is strongly oriented towards the model with highest baseline losses; (ii) the minimax policy is not as tolerant towards small perturbations of policy parameters as the Bayesian policy; and (iii) the minimax policy offers no avenue for incorporating posterior model probabilities derived from data available since monetary union. Thus, we propose preferences for robust policy design that reflect a mixture of the Bayesian and minimax approaches. We show how the incoming EMU data may then be used to update model probabilities, and investigate the implications for policy. (JEL: E52, E58, E61)

1. Introduction

At the start of European monetary union (EMU) central bankers were faced with tremendous uncertainty concerning the functioning of the new euro area economy. Macroeconomic models that would provide reliable measurements

of the area-wide effects of monetary policy were not available. In his contribution to the European Central Bank's (ECB's) first conference on "Monetary Policy-Making under Uncertainty" on 3 December 1999, the ECB's first President Willem Duisenberg described the ECB's assessment of the situation as follows:

> We at the ECB are committed to developing and maintaining a set of tools that are useful for analyzing the euro area economy and examining the implications for future inflation. This is, however, not a trivial task given the large uncertainties that we are facing due to the establishment of a multi-country monetary union. Not only can we expect some of the historical relationships to change due to this shift in regime, but also, in many cases, there is a lack of comparable and cross-country data series that can be used to estimate such relationships.

At the same occasion, ECB chief economist Otmar Issing pointed to research on robust policy rules for the U.S. economy as an example to be followed:

> Given the degree of model uncertainty, central bankers highly welcome the recent academic research on the robustness of monetary policy rules across a suite of different models.

In the years thereafter, the ECB developed a suite of euro area models estimated with synthetic data from the pre-EMU period and made these public in its working paper series, and researchers around the world further developed alternative approaches to robust policy design.[1]

In this paper, we aim to design a monetary policy for the euro area that is robust to the range of uncertainty spanned by the first generation of euro area models estimated at the ECB but that also allows for learning from the new euro area data. First, we document some of the differences between these models. In particular, we show that simple rules optimized for one specific model are not robust, that is, their performance may deteriorate substantially if implemented in one of the other models.[2] Then, following recent research on robust policy for the U.S. economy by Levin, Wieland, and Williams (2003) and Levin and Williams

---

1.   Several researchers have proposed various forms of worst-case analysis to insure against perturbations of single reference models of the economy. Recent examples include Sargent (1999), Hansen and Sargent (2002) , Giannoni (2002, 2007), Onatski and Stock (2002), Onatski and Williams (2003) Tetlow and von zur Muehlen (2001), Tetlow and von zur Muehlen (2004), and Zakovic, Rustem, and Wieland (2007). An early contribution is von zur Muehlen (1982). Other researchers have focused on comparing policy rules across a limited set of well-studied reference models, such as Levin, Wieland, and Williams (1999, 2003), Levin and Williams (2003), Coenen (2007), Adalid et al. (2005) and Kuester and Wieland (2005). Early contributions in this line of research are Becker et al. (1986) and McCallum (1988).

2.   This step was accomplished in parallel by Adalid et al. (2005) and Kuester and Wieland (2005) (our earlier working paper version).

(2003) we compare traditional Bayesian decision-making with minimax analysis that insures against worst-case scenarios.

In characterizing the Bayesian policy, we use equal model probabilities as prior beliefs of the decision-maker. In principle, one could attempt to compute model probabilities using an identical historical data set.[3] However, as pointed out by President Duisenberg, such data are artificial aggregates obtained from national economies with differential monetary policies and fixed but adjustable exchange rates. Thus, an agnostic approach regarding the relative probabilities of models estimated with historical, synthetic data would have seemed appropriate at the start of EMU. Interestingly, the Bayesian policy with flat priors is much more robust than the corner solutions represented by the rules optimized for each specific model. We then compute the minimax policy that delivers maximum insurance against worst-case outcomes. Expected performance under the minimax policy does not deteriorate much relative to the Bayesian policy with flat priors. In other words, the cost of maximal insurance is moderate.

Whereas Levin and Williams (2003) follow earlier work in assessing the costs of insurance in terms of percentage differences in losses, we propose to use instead an absolute measure that is directly interpretable in terms of economic significance. We construct this measure by translating the absolute increase in losses when moving from the first-best model-specific rule to a robust candidate rule into a counterfactual, equivalent increase in inflation volatility. We call this the Implied Inflation (Variability) Premium, in short, the IIP.[4]

Next, we subject our finding regarding the low cost of insurance by means of minimax policy to further scrutiny. First, we take up Sims's (2001) criticism that minimax policies are equivalent to Bayesian policies with unreasonable priors, then we investigate the fault-tolerance of minimax versus Bayesian policies. To assess the Sims critique we construct the priors that would rationalize our minimax policies from a Bayesian perspective. We confirm that these priors may be viewed as excessively tilted towards one model. In other words, the minimax policy is not very far away from one of the corner solutions that were found to lack robustness in the evaluation of model-specific rules. This weakness of the minimax policy is uncovered by the fault-tolerance criteria.[5] We check whether policy is

---

3. The model developers all used historical, synthetic EMU data but different releases and detrending methods.

4. Findings concerning the cost of insurance by means of the IIP for the euro area differ in an important aspect from Levin and Williams's results with three models of the U.S. economy. Focusing on percentage changes of losses these authors claim that costs of insurance may be prohibitively high when policymakers do not have a dual objective, namely, inflation and output stabilization. Using the IIP measure we find that the central bank can insure against model uncertainty at moderate costs even if the objective function does not contain any output stabilization objective. This is of particular importance in light of the Maastricht treaty that established price stability as the primary objective of the ECB.

5. We use this concept in a different way than Levin and Williams (2003), as discussed in Section 4.

fault-tolerant by testing whether minor perturbations in policy parameters have only a minor impact on the performance of a given policy rule in different models. In other words, we test whether the policymaker's "trembling hand" may cause a significant deterioration of policy outcomes under model uncertainty or not. We find that the Bayesian policy with flat priors is fault-tolerant but not the minimax policy. Finally, a third weakness of the minimax rule is that it gives no room for learning about model probabilities using the true euro area data that have become available since the construction of the models.

We move beyond the approaches investigated in the recent literature by proposing to use a mixture of Bayesian and minimax analysis in the design of monetary policy. In this manner we aim to exploit the seemingly low costs of insurance via minimax while overcoming the shortcomings of minimax policies exposed in the preceding analysis. To this end, we propose to use so-called ambiguity-averse preferences for policy design. An intermediate degree of ambiguity aversion uses the model probabilities according to the Bayesian approach to decision-making but gives extra weight to the worst uncertain outcomes as in the minimax approach (see also Brock, Durlauf, and West 2003). We then compute posterior model probabilities from euro area data since 1999 and discuss how these probabilities can be incorporated in policy design.

To accomplish this analysis of robustness and learning under model uncertainty a number of technical assumptions and modeling choices are unavoidable. For example, we choose to focus on the appropriate choice of the systematic component of a monetary policy reaction function, rather than on discretionary policy actions. To this end, we analyze the effect of policies under commitment to a simple rule. Furthermore, we assume that the private sector forms expectations that are correctly based on the central bank's policy rule and the respective model economy. Finally, the robust rules we design remain fairly close to the currently prevailing policy regime (i.e., an estimated rule), such that a change in rules would be "modest" following the suggestion of Leeper and Zha (2003).[6]

The remainder of the paper is structured as follows. In Section 2 we compare the first generation of euro area models developed at the ECB and show that they offer different policy conclusions. Section 3 presents Bayesian and minimax approaches to model uncertainty. The costs associated with full insurance against the worst case are quantified relative to a flat prior Bayesian policy. Section 4 investigates the fault-tolerance of both policies and derives prior beliefs that would render the Bayesian and minimax policies equivalent. Section 5 proposes a compromise approach in terms of intermediate ambiguity aversion that would

---

6.   Leeper and Zha (2003) present an analysis in which they focus on the effect of deviations from a policy rule (policy interventions) on private sector expectations formation when the private sector cannot observe the monetary regime, a complication from which we abstract. Yet, as in their paper, our policy (rule) change is explicitly designed to be reasonably "modest," that is, it does not imply a marked change of the overall regime.

bring together Bayesian and minimax analysis. Section 6 provides an empirical application with learning about model probabilities following the start of the European Monetary Union. Section 7 concludes. Further details regarding the different models, methods, and sensitivity studies are provided in the appendices.

2. The ECB's First-Generation Euro Area Models Differ in Their Policy Implications

The first generation of estimated macroeconomic models for the euro area that were developed at the ECB has been made available to the public in studies by Fagan, Henry, and Mestre (2005) (for details see ECB Working Paper No. 42, January 2001), Coenen and Wieland (2005) (for an earlier version see ECB Working Paper No. 30, September 2000), and Smets and Wouters (2003) (for an earlier version see ECB Working Paper No. 171, August 2002). These models were estimated with pre-EMU synthetic data at a quarterly frequency. Creating these data required constructing and aggregating comparable measures from individual euro area member countries for the years prior to the start of EMU in 1999. ECB staff accomplished this effort and made the data available to outside researchers along with ECB Working Paper No. 42 (Fagan, Henry, and Mestre 2001).

Each of the models that originate in these papers exhibits long-run monetary neutrality as well as short-run nominal inertia. As a result, monetary policy has short-run real effects and the central bank has the ability to stabilize output and inflation fluctuations. In other dimensions, the models differ, for example, in the degree of forward-looking expectation formation, the extent of optimizing behavior by economic agents, and in terms of magnitude, scope, and parameter estimates.[7] The area-wide model presented in Fagan et al. (2001) is a nonlinear model that is used as an element of the ECB's forecasting process. For the purposes of this study, we use the linearized[8] version of this model derived by Dieppe, Kuester, and McAdam (2004) and referred to as the AW model in the following. Expectation formation in this model is largely backward-looking. The two models of Coenen and Wieland (2005) are much smaller than the AW model but incorporate forward-looking expectations. The CW-T variant with Taylor-style staggered wage contracting (see Taylor 1980) exhibits less nominal rigidity than the CW-F variant with Fuhrer–Moore style contracts (see Fuhrer and Moore 1995). Finally,

---

7. An earlier working paper version of this paper provides a more detailed description of the respective models (see Kuester and Wieland 2005, Appendix A).

8. For a detailed description how exogenous variables in the non-linear model were accommodated in the linearization of the AW model and for a description of the data underlying the estimation of the variance-covariance matrix of shocks for the AW model, please refer to Dieppe, Kuester, and McAdam (2005, Appendices A–C). Extensive benchmarking against the non-linear AW model has been conducted. The results show that the linearized model's behavior resembles closely the behavior of the non-linear version.

the Smets and Wouters (2003) model (SW) most completely embodies recent advances in modeling optimizing behavior of economic agents.

Comparing the four models we have found that they exhibit quantitatively important differences with regard to output and inflation dynamics as well as monetary policy effectiveness. These differences imply substantial model uncertainty faced by ECB policymakers at the start of the EMU. They confirm the perceptions expressed by the quotes of then ECB President Willem Duisenberg and Chief Economist Otmar Issing in the introduction to this paper.

2.1. Model Implications for Euro Area Output and Inflation Dynamics

To investigate the differences in the models' implications for euro area output and inflation dynamics we have derived model-dependent autocorrelation and impulse response functions. The autocorrelation functions are computed based on the model-specific estimated variance–covariance matrix of all structural shocks. The impulse response functions describe the response of output and inflation to a particular monetary policy shock. Furthermore, we have also compared the unconditional second moments of inflation, output, and interest rates across models. Along with these measures we have derived proper empirical benchmarks. We proceed to summarize a few results of the model comparison. For further details the reader is referred to Appendix A.

The comparison of model dynamics depends importantly on the assumed rule for monetary policy, that is, the rule for setting the short-term nominal interest rate, r_t. We consider two alternatives. First, we use a common empirical benchmark rule in each of the models. Then we consider the model-specific rules that were estimated along with the other model equations. The empirical benchmark rule is taken from Gerdesmeier and Roffia (2004). We adjust their monthly estimates to a quarterly frequency

$$r_t = 0.87^3 r_{t-1} + (1 - 0.87^3)(1.93\pi_t + 0.28y_t)$$
$$= 0.66r_{t-1} + 0.66\pi_t + 0.10y_t. \tag{1}$$

Here, r_t is the quarterly nominal interest rate (annualized), π_t is the year-on-year inflation rate and y_t is the output gap. Monthly data from 1985 to 2002 were used in estimation.

The autocorrelation functions indicate that the four models exhibit quite different output and inflation dynamics. Under the Gerdesmeier–Roffia rule inflation dynamics in the CW-T models die out within a year. In the SW model inflation persistence is a bit higher, particularly under the model-specific rule. The CW-F model generates longer-lasting swings in inflation. The AW model stands out with the highest degree of inflation persistence. Output dynamics are more persistent

in the SW and CW-T models. In spite of these differences the model-dependent autocorrelation functions still fall inside a 95% confidence band implied by the data. Using model-specific rules we obtain broadly similar results.

With regard to the unconditional variances, we find that the CW-F model induces the highest variance for inflation among the four models—more than twice as much as the SW model—under the Gerdesmeier–Roffia rule. Furthermore, all models generate a standard deviation of the output gap that is larger than the mean value observed in the data.

The comparison of impulse responses following a positive interest rate shock provides an indication of the uncertainty about policy effectiveness. If the ECB sets interest rates according to the Gerdesmeier–Roffia rule, the AW model again generates the highest degree of inflation persistence. The trough is reached more than five years after the shock. Similarly, output returns only very gradually to baseline. The SW model marks the smallest degree of output and inflation persistence. The CW-F and CW-T models lie in between the others. As expected, Taylor contracts generate less inflation persistence than Fuhrer–Moore contracts. Model specific rules imply greater effects of policy shocks than under the Gerdesmeier–Roffia rule, particularly in the SW model.

In sum we find that each model fits certain stylized facts in the data. Some models fit the cross-correlations better than the standard deviations, others are better at reproducing the stylized response of the economy to a monetary shock than to fit the former moments. Thus, each model may be seen as a reasonable description of the economy although none of the models is the perfect model of the economy. The ordering of models may differ depending on the metric used in estimation and in evaluating the fit. The models exhibit quantitatively important differences concerning the effectiveness of monetary policy, as well as the intrinsic persistence and the unconditional variability of output and inflation. Thus, euro area policymakers are faced with considerable model uncertainty.[9]

2.2. Model Implications for Euro Area Monetary Policy Rules

As noted previously a simple outcome-based Taylor-style policy rule as defined by equation (1) describes euro area interest rate behavior quite well. Such a rule

---

9. Of course, one could give these results a more negative interpretation. Because all models "fail" the empirical evidence presented in a certain dimension, some readers may be tempted to dismiss all of these models for giving policy advice and would rather postpone analysis until a considerably "better" model of the economy has emerged. We have three objections to this view. First, actual policymaking is characterized by a heterogeneity of views regarding the evolution of the economy and also regarding the metric that tells a good model from a bad model. Second, and perhaps most importantly, the fact that all four models are used at the ECB suggests that they all deserve positive weight in policy assessments. Finally, in policy practice it is not possible to discard all models of the economy and postpone policymaking until models arrive that meet greater consensus. It is this very need to use the existing models to inform our views despite their differences which appears to be at the heart of policymaking under model uncertainty.

has three response parameters that we denote by (ρ, α, β) in the following. These parameters multiply the lagged interested rate r_{t-1}, the year-on-year inflation rate π_t, and the output gap y_t, all in deviations from steady state:

$$r_t = \rho r_{t-1} + \alpha \pi_t + \beta y_t. \tag{2}$$

We now proceed to ask the question how well such a rule would perform in stabilizing output and inflation dynamics in the euro area models under consideration. We focus our analysis on this class of rules, not only because of their empirical success in explaining ECB interest rate choices, but also because research on the U.S. economy referred to in the quote of Otmar Issing in the Introduction emphasizes the robustness properties of such rules. For example, simple outcome-based rules of this form tend to be more robust than rules with more states (see Levin, Wieland, and Williams 1999) and forecast-based rules with longer horizons (see Levin, Wieland, and Williams 2003). Further improvements in stabilization performance that can be achieved with forecast-based rules are hampered by the need for a forecasting model that may be incorrect and introduce additional uncertainty (see Levin et al. 2003 and Coenen 2007).

To evaluate policy performance we use the weighted sum of unconditional variances of inflation, the output gap, and the change in the interest rate. This measure is consistent with a standard quadratic loss function:[10]

$$\mathcal{L}_m = \mathrm{Var}(\pi) + \lambda_y \mathrm{Var}(y) + \lambda_{\Delta r} \mathrm{Var}(\Delta r), \quad m \in \mathcal{M}. \tag{3}$$

$\mathrm{Var}(\cdot)$ denotes the unconditional variance, and $\mathcal{M} = \{\text{CW-F, CW-T, SW, AW}\}$ is the model space. The parameter $\lambda_y \geq 0$ determines the policymaker's preference for reducing output variability around potential relative to curbing inflation variability. The weight $\lambda_{\Delta r} > 0$ introduces a preference for restraining the variability of changes to nominal interest rates. For the remainder of this paper we use values of $\lambda_y \in \{0, 0.5, 1\}$ that cover the range from strict to flexible inflation targeting and a value for $\lambda_{\Delta r} = 0.5$.[11]

---

10.   See Levin et al. (1999, 2003), Coenen (2007), and others in this literature. We follow the common practice in the literature in examining only diagonal cases of the loss function, which is also the most common representation of welfare when obtaining quadratic loss functions on the basis of approximating the welfare function from simple micro-founded models (see Woodford 2003). Furthermore, such a loss function is consistent with the objectives of many central banks. This does not mean, however, that with a diagonal loss function policymakers would not take cross-moments into account. They minimize losses subject to the law of motion implied by the models, which typically imply that endogenous variables are correlated. Finally, we also note that more complicated micro-founded models such as the Smets and Wouters (2003) model provide quadratic approximations of the welfare of a representative consumer that incorporate additional variables.

11.   This value for $\lambda_{\Delta r}$ ensures that interest rate variability does not stray very far from what is observed empirically for the euro area, so all policy rules which we consider are close to the current policy regime. A change of policy rule should thus be modest enough to render the models still a viable representation of the economy. This is related to the modesty concept in Leeper and Zha (2003) albeit we assume that private sector expectations incorporate the true monetary rule at each point in time. We have conducted further sensitivity studies with $\lambda_{\Delta r} \in \{0.1, 1\}$ which are reported in the Appendix.

Our analysis is subject to a number of caveats. First, all models are linearized around their steady state. Monetary policy therefore does not have an effect on the average level of inflation and output in the economy, nor is such an effect accounted for in the computation of losses. Second, in computing losses we abstract from any transitional dynamics when moving from an initial stochastic steady state obtained under one interest rate rule to the equilibrium under the optimized rule.[12] Although we do not assess the importance of these two caveats, they might partly be alleviated by our explicit focus on policy rules that are close to the current monetary regime in terms of interest rate behavior. A third caveat that applies to our analysis is that we understand a model as being given by a set of equations plus one particular parametrization. All four models are used at the ECB and, as commonly occurs, have been estimated using different methodologies. The parametrization used in this paper represents the favorite parametrization that the group of modelers involved in building the model provided. Still, different parameterizations may have different welfare and optimal policy implications.[13]

With these caveats in mind, we proceed to assess the performance and robustness of optimal simple policy rules as defined by equation (2). We conduct this assessment assuming that the policymaker credibly commits to following the rule. The best simple rule is obtained by choosing the response parameters, (ρ, α, β), in equation (2) so as to minimize the loss defined by equation (3) for each model.[14] Table 1 summarizes the parameters of the simple rules optimized for each model for the three different values of the weight on output variability in the loss function, λ_y.[15]

The best rule in the CW-F model features moderate to strong interest rate smoothing (ρ between 0.8 and 0.9) and strong feedback to inflation (α between 0.7 and 0.8). The best CW-T rule implies a bit more interest smoothing (greater ρ)

---

12. We also assume that monetary policy is perfectly credible and that the private sector immediately adjusts its expectations formation process to account fully for the actual monetary policy rule. The work by Erceg and Levin (2003) suggests that allowing for a slow adaptation of private sector beliefs to a change in the rule will lead to slow adjustment of the economy's law of motion to the new stochastic steady state with a potential bearing on the ranking of different rules in terms of welfare.

13. The results of Levin et al. (2006) appear to suggest though that model uncertainty (in their case the specification of the labor market) entails much greater differences in optimal simple policy rules (be they model-specific or robustified) than parameter uncertainty. Similar conclusions emerge when Levin and Williams (2003) compare their results to the findings of the literature.

14. An alternative would be to use the first-best policy under commitment as a benchmark. We have not done so because we are interested in measuring the cost of insurance implied by a simple rule (chosen due to a preference for robustness concerning model uncertainty), rather than being interested in the benefits that may be possible from first-best policy compared to simple rules in any specific model. For this purpose the model-optimized simple rule is the proper benchmark. As to the potential benefits of first-best policies compared to optimized three-parameter rules, those are moderate for models with rational expectations but can be substantial for models with primarily backward-looking expectations (see Dieppe et al. 2005).

15. The moments reported here are unconditional second moments when the policymaker commits to using the respective rule. We abstract from analyzing the standard deviations on any path during a transition phase when switching from any current rule to the rule given in Table 1.

TABLE 1. Optimized simple rules.

Model	$\lambda_y = 0$			$\lambda_y = 0.5$			$\lambda_y = 1$		
	ρ	α	β	ρ	α	β	ρ	α	β
CW-F	0.9	0.8	0.4	0.8	0.7	0.6	0.8	0.7	0.8
CW-T	1.0	0.3	0.1	0.8	0.2	0.6	0.8	0.2	0.8
SW	1.0	0.4	0.0	1.0	0.2	0.8	1.0	0.2	1.3
AW	0.6	0.5	0.5	0.4	0.6	1.3	0.4	0.5	1.7

Note: Parameters of optimal simple rules for each model. Shown is the case $\lambda_{\Delta r} = 0.5$ for loss function (3). Results using a lower or a higher weight on interest rate changes, $\lambda_{\Delta r} \in \{0.1, 1\}$, are presented in Appendix E.

but smaller response coefficients on inflation due to the lower degree of inflation persistence in that model. The best SW policy consistently uses the highest weight on the lagged interest rate ($\rho = 1$) essentially implementing a first-difference rule. This finding confirms earlier research showing that the prominence of rational expectations and forward-looking behavior in models with optimizing agents renders such a policy strategy optimal. In sharp contrast, the rule optimized in the AW model exhibits much lower values of ρ ranging between 0.4 and 0.6. Optimal AW policy also features a strong feedback to the output gap as a proxy for future inflation (see Dieppe et al. 2005).

The characteristics of optimal policy rules found in our model space are broadly representative of the range of characteristics of optimal simple rules found in the literature; Table B.1 in Appendix B reports an overview of four studies for the U.S., the euro area, and Canada, respectively. Exceptions concern the stronger response to inflation in those rules when output receives no weight in the policymakers' loss function.

Furthermore, it is of interest to compare the optimized simple rules from the different models with the estimated ECB policy rule taken from Gerdesmeier and Roffia (2004) given by equation (1). There is no perfect match. However, the estimated coefficients on the lagged interest rate and inflation ($\rho = 0.66, \alpha = 0.66$) are close to the values that are optimal in the AW model ($\rho = 0.6, \alpha = 0.5$) when the central bank's preference weight on output variability is zero ($\lambda_y = 0$). The estimated coefficient on the output gap is quite a bit smaller in the estimated rule than in the AW model (at $\beta = 0.1$). Such low values occur only for the SW model and the CW-T model when $\lambda_y = 0$.

Next we investigate the robustness of simple rules, that is, we evaluate the performance of rules optimized for one model in the other three models. Such a comparison reveals that the euro area models considered have quite different policy implications. Table 2 reports the percentage increase in loss when using a rule optimized for Model X in Model Y relative to using the rule that is optimal for Y; see the entries in square brackets. We confirm earlier results in the literature. Rules with a high degree of interest rate smoothing such as those optimized in the SW model generate substantial losses and even explosiveness in models with

TABLE 2. Robustness of rules optimized for a specific model.

| | CW-F rule evaluated in | | | | | | CW-T rule evaluated in | | | | | |
| | CW-T | | SW | | AW | | CW-F | | SW | | AW | |
λ_y	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]
0.0	.15	[31]	.16	[53]	.56	[126]	1.13	[161]	.03	[9]	∞	[∞]
0.5	.10	[12]	.11	[17]	.79	[82]	1.03	[103]	.09	[15]	.34	[29]
1.0	.13	[12]	.19	[27]	.93	[71]	1.57	[145]	.13	[18]	.55	[36]

| | SW rule evaluated in | | | | | | AW rule evaluated in | | | | | |
| | CW-F | | CW-T | | AW | | CW-F | | CW-T | | SW | |
λ_y	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]	IIP	[%$\mathcal{L}$]
0.0	2.06	[360]	.04	[8]	∞	[∞]	.63	[78]	.12	[24]	.23	[78]
0.5	.72	[66]	.16	[19]	2.06	[324]	1.50	[167]	.12	[15]	.28	[49]
1.0	1.18	[101]	.27	[26]	2.01	[220]	ME	[ME]	ME	[ME]	.36	[54]

Notes: IIP-implied inflation (variability) premium relative to first-best simple rule for each model (in percentage points). In square brackets %$\mathcal{L}$-percentage losses relative to first-best simple rule for each model (in %). The notation ∞ indicates that the implemented rule results in instability; the notation ME indicates that the implemented rule results in multiple equilibria. Shown is the case $\lambda_{\Delta r} = 0.5$ for loss function (3). Results using a lower or a higher weight on interest rate changes, $\lambda_{\Delta r} \in \{0.1, 1\}$, are presented in Appendix E.

a significant backward-looking component such as the AW model. On the other hand, policy designed for models with strong intrinsic persistence such as AW may not be active enough to anchor expectations if agents indeed were more forward-looking. For example, for $\lambda_y = 1.0$, the best AW policy would imply indeterminacy in CW-F and CW-T. Otherwise percentage losses are largest in AW and CW-F, reaching up to 360%.

The percentage losses reported in Table 2, which have been commonly used in the recent literature, may overemphasize the extent of model uncertainty in particular when the baseline loss is rather low or when the loss function heavily penalizes small deteriorations in economic outcomes. Similarly, substantial changes in economic outcomes may be de-emphasized when the baseline loss is already large.

As a remedy to this drawback of relative percentage losses we introduce the implied premium on inflation variability or implied inflation premium, the IIP. This premium measures the increase in the standard deviation of inflation relative to the outcome under the best simple rule that is necessary to match the loss under the alternative policy. In other words, we attribute the deterioration in loss that results from a lack of robustness entirely to inflation variability keeping the standard deviation of output and interest rates at the benchmark level. The benefit of this premium is that units are intuitive (percentage points of the standard deviation of the inflation rate) and interpretable on an economic scale. The corresponding premia are reported in Table 2 under the heading IIP.

The implied inflation variability premia shed new light on the comparison of rules across models. With no weight on the output gap the SW model exhibits

low baseline losses. Using the CW-F policy in the SW model implies an increase in loss of 53% (upper left panel of Table 2), which seems sizeable. The IIP of 0.16 percentage points on annual inflation (an increase from 0.82 to 0.98 percentage points, say), however, is rather small. Similarly, using the CW-F policy in AW generates an increase in losses of 126%, which appears prohibitive. The implied increase of 0.56 percentage points in inflation variability is still sizeable but does not appear extreme. As these examples show, relative losses may be misleading. We will therefore report the IIP where appropriate in the remainder of the paper.

3. Two Perspectives on Policy Robustness: Bayes or Minimax

3.1. *Bayesian Policy Minimizes Expected Loss Given Model Probabilities*

A natural first step in the search for a rule that performs more consistently across models than the model-specific rules discussed in the preceding section is to take a Bayesian perspective, as recommended by Levin, Wieland, and Williams (2003). A more robust rule can be found by minimizing a probability-weighted loss function. The Bayesian loss is

$$\mathcal{L}^B = \min_{(\rho,\alpha,\beta)} E_{\mathcal{M}}\{\mathcal{L}_m\} = \min_{(\rho,\alpha,\beta)} \sum_{m\in\mathcal{M}} p_m \mathcal{L}_m, \tag{4}$$

where p_m are the policymaker's priors as to model m. As a start, we consider flat priors $p_m = 1/|\mathcal{M}|$. This may be an appropriate prior when there is uncertainty surrounding the relative fit of models, for example when policymakers do not consider the data informative about model probabilities or when it is not fully agreed against which metrics or with which data series to evaluate the models. Given that all four models considered in this paper were estimated with pre-EMU synthetic data and recognizing the concerns expressed by ECB President Willem Duisenberg as quoted in the introduction to this paper, it seems very reasonable to us to use flat priors as our point of departure. Of course, the data that have become available subsequent to the start of EMU may be used by the ECB to learn about the likelihood of different models. We study such a learning process and its implications for robust policy design in Section 6 of this paper.

The performance of the Bayesian policy in terms of implied inflation variability premia is reported in Table 3.

Considering all four models in policy design in this manner helps avoid extreme outcomes such as indeterminacy and multiple equilibria in forward-looking models or explosiveness in backward-looking models. Furthermore the increase in inflation variability that would match the increase in loss relative to the

TABLE 3. Flat Bayesian priors versus model-specific simple rules.

λ_y	Optimal rule			Inflation (variability) premium and [%$\mathcal{L}$]			
	ρ	α	β	CW-F	CW-T	SW	AW
0.0	0.8	0.6	0.4	.07 [7]	.10 [21]	.16 [53]	.12 [22]
0.5	0.7	0.7	0.8	.09 [7]	.08 [9]	.14 [22]	.26 [21]
1.0	0.7	0.8	1.1	.11 [7]	.12 [11]	.21 [29]	.32 [19]

Notes: Optimal policy rule parameters and IIP relative to the simple rules optimized for each model (in percentage points). Shown is the case $\lambda_{\Delta r} = 0.5$ for loss function (3). For comparison with the literature, we also report the percentage increase in losses relative to the simple rules optimized for each model (in square brackets). Results using a lower or a higher weight on interest rate changes, $\lambda_{\Delta r} \in \{0.1, 1\}$, are presented in Appendix E.

optimized model-specific rules, that is the IIP, reaches at most 0.32 percentage points on the standard deviation of annual inflation. Because the CW-F model features the highest benchmark losses of all models (Table A.1 in the appendix), it implicitly plays a greater role in the Bayesian optimization than the other models given equal probabilities. The Bayesian rule with flat priors looks much like the best CW-F policy albeit with less interest rate smoothing. The reduction in ρ serves to help performance in the AW model.

To allow for a comparison with other findings in the literature, Table 3 also reports the percentage increase in losses. These results serve to illustrate the advantage of using the IIP measure. Percentage losses of more than 50% in the SW model when using the Bayesian flat prior policy may seem to suggest that insurance against model-uncertainty is prohibitively costly. Once measured on an economically interpretable scale, however, the opposite conclusion emerges: Policymakers in the euro area can insure against model uncertainty at reasonable cost in each of the models of the economy.

3.2. Minimax Policy Ensures Against the Worst-Case Model

Bayesian optimization with flat priors is one possible approach when policymakers are fairly agnostic about which among a set of models is most appropriate as a tool for policy design. However, policymakers may well consider it impossible to define any sensible prior probability distributions over the model space. This perspective is typically termed *Knightian uncertainty*. Such an agnostic policymaker could instead ask how to best insure herself against worst-case scenarios, namely, worst-case models. This question can be answered by minimax analysis. The minimax optimization problem arises from a game between a policymaker who attempts to minimize loss and nature that chooses a model from the model space so as to maximize loss,

$$\mathcal{L}^M = \min_{(\rho,\alpha,\beta)} \max_{m \in \mathcal{M}} \mathcal{L}_m. \tag{5}$$

TABLE 4. Minimax policy relative to model-specific rules.

	Optimal rule			Inflation premium				Loss premium	
λ_y	ρ	α	β	CW-F	CW-T	SW	AW	$\Delta\mathcal{L}^{worst}$	$\Delta\mathcal{L}^{expect}$
0.0	0.9	0.8	0.5	0	.15	.16	.57	7	17
0.5	0.8	0.7	0.6	0	.10	.11	.73	7	9
1.0	0.8	0.7	0.8	0	.13	.20	.77	7	6

Notes: Optimal policy rule parameters and IIP relative to first-best simple rule for each model (in percentage points). Shown is the case $\lambda_{\Delta r} = 0.5$ for loss function (3). $\Delta\mathcal{L}^{worst}$-percentage reduction of worst-case loss relative to worst outcome under flat Bayesian priors; $\Delta\mathcal{L}^{expect}$-percentage increase in expected loss relative to Bayesian policy (flat priors). Results using a lower or a higher weight on interest rate changes, $\lambda_{\Delta r} \in \{0.1, 1\}$, are presented in Appendix E.

Thus, no model probabilities need to be specified to derive the minimax or full insurance policy. Appendix C describes in more detail the algorithm we use to solve the minimax problem.[16]

The policy rule coefficients and IIP premia implied by the minimax policy are summarized in Table 4.

The minimax policy is close to the best rule for the CW-F model, that is, the model with the highest baseline loss. The IIPs premia are relatively modest for the CW-T and SW models but increase with the AW model, adding up to 0.77 percentage points to the standard deviation of the annual inflation rate. The deterioration in stabilization performance compared to model-specific rules in the AW model is noticeably stronger than under the Bayesian policy with flat priors (Table 3).

Overall, the costs for insuring against model uncertainty nevertheless seem moderate given the considerable range of output and inflation dynamics in the model space. Expected loss relative to flat Bayesian priors increases by 6–17%. Similarly, the gains from insurance relative to Bayesian policy in the worst case scenario are moderate. Because the CW-F model already influences the flat prior Bayesian policy quite heavily, worst-case losses are only reduced by 7% under the minimax relative to the Bayesian policy. Finally, we note that the "insurance premium," that is, the cost of an increase in expected loss for a given reduction of worst-case loss, decreases as policymaker preferences assign a greater weight to the output gap.

4. Three Arguments Against Minimax

In light of the preceding analysis the minimax policy appears to be a very attractive choice as a robust monetary policy rule for the euro area. It achieves

---

16. Most papers employing minimax analysis consider perturbations around single reference models of the economy (e.g., Onatski and Williams 2003). Given the diversity of the models used in our paper any reference model meant to nest all the models would be untractable and it would be difficult to control all scenarios which the policymaker considers are reasonable. The latter is automatically ensured in our analysis which builds on a discrete model space.

effective insurance against the worst-case model and its dynamics. This insurance is obtained at surprisingly low cost in terms of expected performance deterioration relative to a flat-prior Bayesian policy. However, three arguments against the minimax approach deserve further consideration: (i) the minimax approach may generate policy recommendations that a Bayesian decision-maker would choose only if the probability of the model with the highest losses is overwhelming; (ii) to the extent that minimax leans strongly towards a single worst-case model, small perturbations of the minimax policy parameters may inherit the lack of robustness of the rules optimized for that model (i.e., minimax may not be fault-tolerant); and (iii) the minimax policy offers no avenue for incorporating posterior model probabilities derived from data available since the start of EMU. We postpone the consideration of posterior model probabilities to Section 6 and consider here the first two arguments against minimax.

4.1. Minimax Policy is Replicated by Bayesian Priors Concentrated on a Subset of Models

As noted by Sims (2001) and others, minimax analysis also has a Bayesian interpretation. There exists a set of "priors" that, if used along with a Bayesian objective function, would yield the same robust policy rule as the one obtained using a minimax objective. The major difference to truly Bayesian priors is that these implied priors are an outcome of the optimal decision and not a primitive for the decision problem. Therefore they need not be rational from a Bayesian perspective. Instead they inform the policymaker about which model underlies the worst-case scenarios?[17] Sims has criticized minimax for often implying economically unreasonable priors once it is considered from a Bayesian perspective. We already guard against unreasonable priors by restricting the model space to models that fit European data and have been developed for the purpose of evaluating euro area monetary policy. However, minimax policies may be driven primarily by just one of the models in the model space. Thus, we now proceed to extract the priors that would rationalize the euro area minimax policy derived in the preceding section from the Bayesian perspective. The procedure involved is discussed in Appendix D. The implied priors are shown in Table 5.

Indeed, these priors indicate that model risk in the space we consider lies primarily with the CW-F model. Only if the policymakers assign a high weight to output stability does the AW also play a small role in the determination of the

---

17. An alternative interpretation of minimax analysis exists. Instead of choosing Bayesian priors to support the minimax robust policy, one can also transform the loss-function so that a decision under a Bayesian objective coincides with the decision under a minimax objective. This allows us to interpret minimax robust decisions as decisions of a Bayesian policymaker with an infinite degree of risk aversion. Adam (2004) clarifies the link between Bayesian decision-making and minimax decision-making.

TABLE 5. Minimax implied priors.

λ_y	CW-F	CW-T	SW	AW
0.0	1.000	0	0	0
0.5	0.989	0	0	0.011
1.0	0.961	0	0	0.039

Notes: Bayesian priors backing the flat minimax solution. Shown is the case $\lambda_{\Delta r} = 0.5$ for loss function (3). Results using a lower or a higher weight on interest rate changes, $\lambda_{\Delta r} \in \{0.1, 1\}$, are presented in Appendix E.

minimax policy. The large weight given to the CW-F model is consistent with our earlier finding that inflation in that model (under the rule optimized for the model) is substantially more volatile than in the AW model (1.6 vs. 0.9 percentage points). The output gap, however, is somewhat more volatile in the AW model (2.0 relative to 1.4 percentage points). Thus, as the preference parameter on the output gap, λ_y, increases this volatility difference gains more importance in the minimax calculus and the AW model receives some (albeit small) weight in the priors implied by the minimax policy. Neither the CW-T nor the SW model ever appear among the worst case losses. Thus, we conclude that Sims's (2001) critique of minimax applies to the euro area minimax policy we have derived.

4.2. Minimax Policy Is Not as Fault Tolerant as Bayesian Policy

Given that the minimax approach leans strongly towards a single worst-case model, the question arises whether small perturbations of that policy, perhaps due to the "trembling hand" of the policymaker, might not inherit the lack of robustness of the rules optimized for this specific model. Levin and Williams (2003) have proposed fault tolerance as a diagnostic tool to assess robustness relative to small perturbations. They suggest plotting variations in the policy rule parameters against percentage increases in losses for each model. The procedure implies varying one parameter at a time while keeping the others fixed at the values implied by the best simple rule for each model. A *model* is deemed fault tolerant if deviations from the best rule for that specific model do not trigger a steep increase in losses in that model. A drawback of this procedure is that comparisons across models are not ceteris paribus in terms of the policy parameters. For example, when varying the smoothing parameter, ρ, policy rules used in the different models also have different responses to inflation and the output gap, α and β. Thus, we decided to use the concept of fault tolerance somewhat differently. We ask whether *policy* is fault tolerant, meaning in particular whether minor perturbations in policy parameters have only a minor impact on performance across models.

Our findings for the Bayesian flat prior policy are displayed in Figure 1. Each panel considers implied inflation premia across the four models when varying one parameter of the Bayesian rule at a time, while keeping the others fixed.

FIGURE 1. Fault tolerance for flat Bayesian priors. For each of the models, a line traces out the IIP (in basis points) under the flat Bayesian policy as the specified parameter of the policy rule is varied, holding the other two parameters fixed at their respective values under flat Bayesian priors. The left (right) column of panels pertains to the preferences for $\lambda_y = 0$ (1) and $\lambda_{\Delta r} = 0.5$. A vertical line marks the optimal Bayesian parameter values.
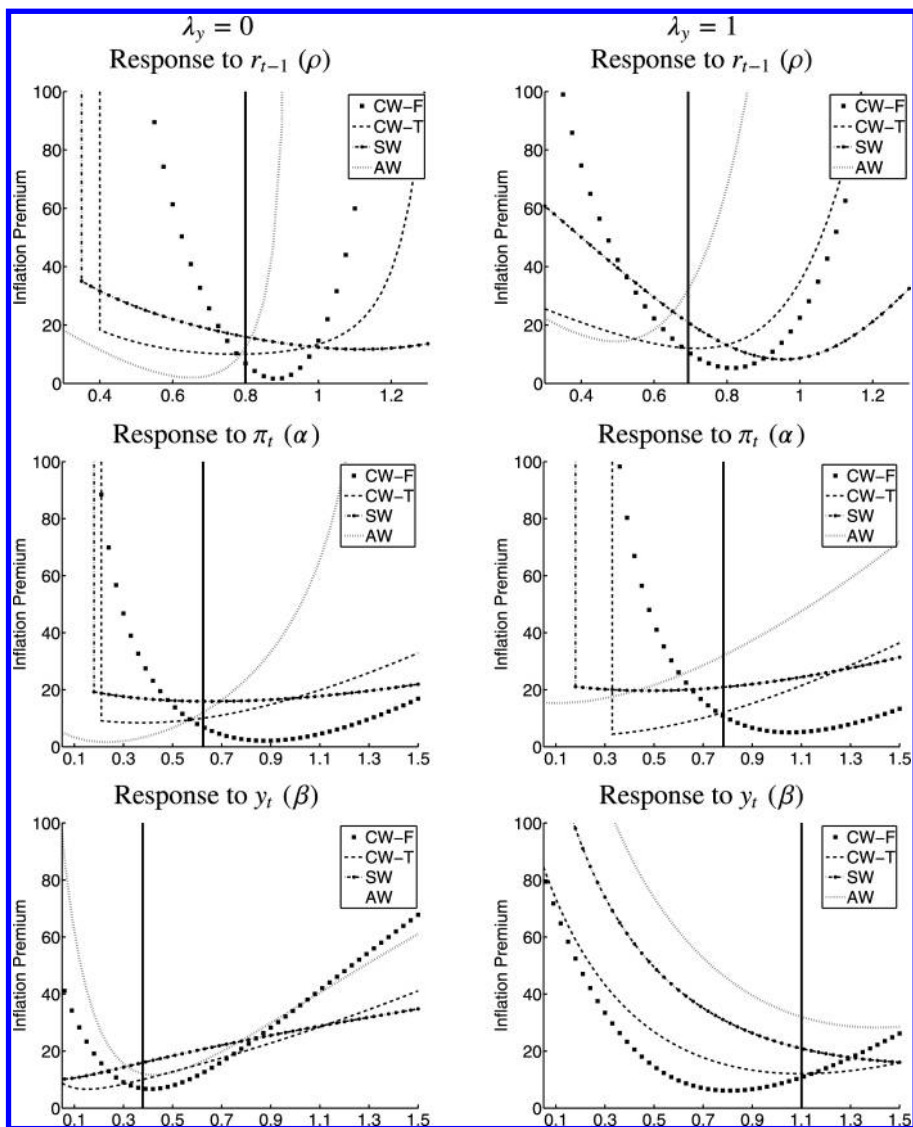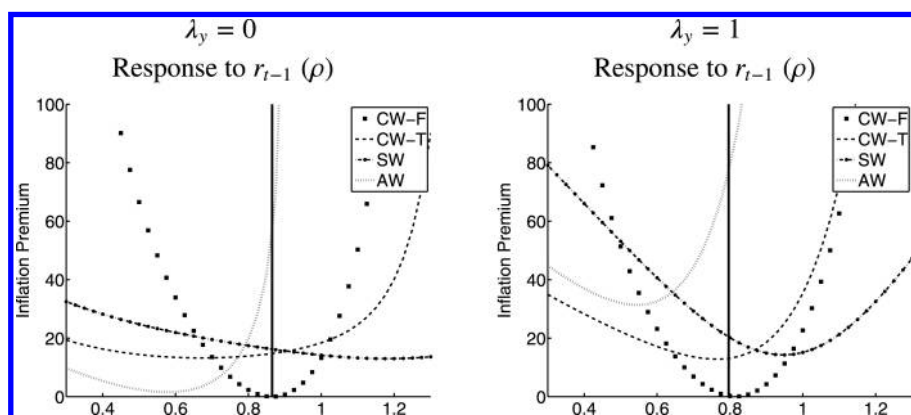
FIGURE 2. Fault intolerance for minimax policy. For each of the models, a line traces out the IIP (in basis points) under the full insurance policy as the persistence parameter of the policy rule is varied, holding the other two parameters fixed at their respective values under the full insurance policy. The left (right) panel pertains to the preferences for $\lambda_y = 0$ (1) and $\lambda_{\Delta r} = 0.5$. A vertical line marks the optimal minimax parameter value.

The respective columns (of three panels each) are associated with values for the preference parameter λ_y of zero and one. Policy is fairly tolerant to deviations in α and β. IIPs do not rise fast in the neighborhood around the optimal parameter values. The admissible range of variations is smaller for ρ but still permits a safety margin for a policymaker with "trembling hands." We conclude that Bayesian policy under flat priors is fault tolerant.

Interestingly, the minimax policy loses the fault tolerance property enjoyed by the Bayesian rule with flat priors. Figure 2 displays fault tolerance plots for the parameter on the lagged interest rate, ρ, in the minimax policy.[18] Minor changes in policy can lead to a strong increase in the implied inflation variability premia for the AW model, because this model does not respond well to a high degree of interest-rate smoothing (see also Dieppe et al. 2005). Minimax monetary policy only just avoids this in the AW model but not with a wide enough security margin. Thus, "trembling hands" in the conduct of policy could have a strong negative impact when implementing a policy that attempts to guard against worst-case model uncertainty.

5. A Compromise: Intermediate Ambiguity Aversion or "Bayes Meets Minimax"

Having shown the Bayesian priors associated with the euro area minimax policy and having documented its lack of fault tolerance, we judge the pure minimax

---

18. The minimax policy is similarly fault intolerant with respect to the other parameters of the rule. Plots are available upon request.

approach less useful for euro area monetary policymakers. However, rather than abandoning the minimax approach completely, we instead suggest a compromise. We propose to take some prior information (or preferences across models) into account and balance this prior (or preference) against the worst uncertain outcome. Following Epstein and Wang (1994) such preferences may be formalized as follows:

$$\mathcal{L}^A = \min_{(\rho,\alpha,\beta)} \left\{ (1-e) \sum_{m\in\mathcal{M}} p_m \mathcal{L}_m + e \max_{m\in\mathcal{M}} \mathcal{L}_m \right\}, \tag{6}$$

where $e \in [0, 1]$ indexes the degree of desired insurance against worst case outcomes. These preferences are also called *ambiguity averse* (see Brock, Durlauf, and West 2003) because the decision-maker places extra weight on the worst uncertain outcome.[19] At the one extreme, $e = 0$, these preferences amount to minimizing expected losses from a Bayesian perspective. At the other extreme, $e = 1$, these preferences deliver the minimax policy displayed in Table 4. The value of e may be chosen by the policymaker according to his desire for insurance. The prior model probabilities, p_m, may either be estimated from available data or set to reflect the policymaker's convictions regarding appropriate modeling.

In this section, we present three examples of ambiguity-averse policies. We first consider a policymaker who has flat prior model probabilities. In addition, this policymaker has a desire for insurance against model uncertainty that leads her to choose a value of $e = 0.5$ for the ambiguity aversion parameter. The policy pursued by this decision-maker corresponds to the flat prior Bayesian policy with additional weight given to the (endogenously determined) worst-case scenario.

The first panel of Table 6 (labeled Flat priors) reports the policy rules and IIP premia for the ambiguity averse policy with flat Bayesian priors. The resulting policy coefficients and performance measures reflect the policymaker's priors together with the extra insurance against the worst-case scenario. The worst-case scenario continues to be represented primarily by the CW-F model, as indicated by the implied priors from a purely Bayesian perspective (shown in the first panel of Table 7).

The loss in performance due to additional insurance against the worst case is slightly greater for the CW-F and CW-T models but smaller for the SW and particularly the AW model. Thus, the policy rule constitutes a mix of the prescriptions by the minimax policy and the Bayesian policy. As a result, the ambiguity-averse policy with flat priors is considerably more fault-tolerant than the minimax policy,

---

19.   The preferences are a special case of Epstein and Wang's (1994) "ϵ-contamination" preferences. Strictly speaking, also the minimax preferences reflect ambiguity aversion, of course.

TABLE 6. Ambiguity-averse policy.

	Optimal rule			Inflation (variability) premium			
λ_y	ρ	α	β	CW-F	CW-T	SW	AW
Flat priors							
0.0	0.8	0.7	0.4	.02	.19	.17	.24
0.5	0.8	0.8	0.7	.02	.09	.13	.45
1.0	0.8	0.8	0.9	.02	.13	.21	.56
AW-leaning							
0.0	0.8	0.8	0.5	.06	.14	.18	.11
0.5	0.6	1.0	0.9	.10	.14	.19	.20
1.0	0.6	1.1	1.3	.12	.21	.30	.23
SW-leaning							
0.0	0.9	0.7	0.4	.00	.13	.15	.69
0.5	0.8	0.7	0.7	.01	.09	.09	.74
1.0	0.8	0.7	0.9	.02	.11	.15	.79

Notes: Optimal policy rule parameters and IIP relative to first-best simple rule for each model (in percentage points). Shown is the case $\lambda_{\Delta r} = 0.5$ for loss function (3). Three ambiguity-averse solutions are shown: one using flat model priors, one giving extra weight only to the AW, and one giving extra weight only to the SW model. Parameters for the ambiguity averse solution are $p_m = 0.25$ for the former, and $p_{AW} \approx 1$ and $p_{SW} \approx 1$, respectively, for the latter. In each of the cases, the policymaker gives a weight of $e = 0.5$ to the worst-case scenario. Results using a lower or a higher weight on interest rate changes, $\lambda_{\Delta r} \in \{0.1, 1\}$, are presented in Appendix E.

as evident from a comparison of the fault-tolerance charts in Figure 3 with those of Figure 2.

The second example illustrating the use of the ambiguity averse preferences concerns a policymaker who strongly favors the AW model. The AW model follows a long and well-established tradition of macroeconomic modeling at central banks. Thus, it would not be surprising if some members of the Governing Council of the European Central Bank would favor policy analysis conducted on the basis of this model. The AW model gives little weight to forward-looking private sector expectations. Dynamic adjustment mechanisms are defined sufficiently flexible to incorporate the empirical output and inflation persistence in the endogenous dynamics of the model. In the following we consider a policymaker who, instead

TABLE 7. Ambiguity-averse implied priors.

	Flat priors				AW-leaning				SW-leaning			
λ_y	CW-F	CW-T	SW	AW	CW-F	CW-T	SW	AW	CW-F	CW-T	SW	AW
0.0	.625	.125	.125	.125	.5	0	0	.5	.500	0	.5	0
0.5	.625	.125	.125	.125	.5	0	0	.5	.478	0	.5	.022
1.0	.625	.125	.125	.125	.5	0	0	.5	.459	0	.5	.041

Notes: Bayesian priors backing the ambiguity averse solution. Shown is the case $\lambda_{\Delta r} = 0.5$ for loss function (3). Three ambiguity-averse solutions are shown: one using flat model priors, one giving extra weight only to the AW, and one giving extra weight only to the SW model. Parameters for the ambiguity averse solution are $p_m = 0.25$ for the former, and $p_{AW} \approx 1$ and $p_{SW} \approx 1$, respectively, for the latter. In each of the cases, the policymaker gives a weight of $e = 0.5$ to the worst-case scenario. Results using a lower or a higher weight on interest rate changes, $\lambda_{\Delta r} \in \{0.1, 1\}$, are presented in Appendix E.
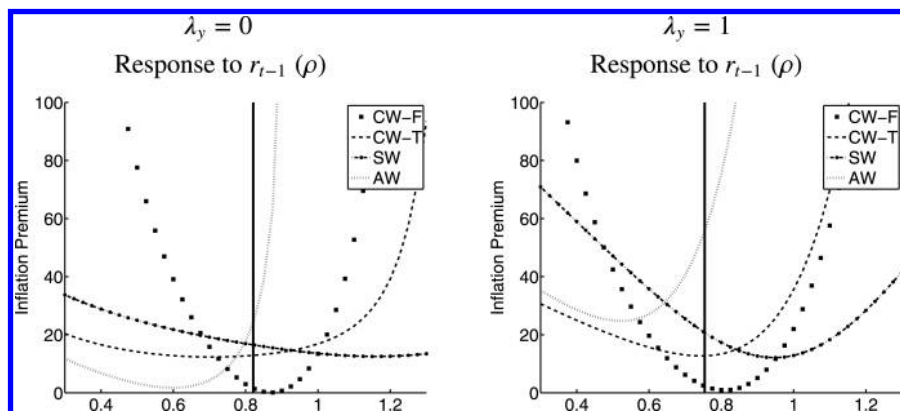
FIGURE 3. Fault tolerance of ambiguity-averse policy with flat priors. For each of the models, a line traces out the IIP (in basis points) under the ambiguity-averse policy as the persistence parameter (l) of the policy rule is varied, holding the other two parameters fixed at their respective values under the full insurance policy. The left (right) panel pertains to the preferences for $\lambda_y = 0$ (1) and $\lambda_{\Delta r} = 0.5$. A vertical line marks the optimal ambiguity averse parameter value. Shown is the ambiguity-averse policy with flat priors ($p_m = 0.25$ for all models), giving weight $e = 0.5$ to the worst-case loss.

of assigning equal weight to all models, has a profound preference for this modeling approach. The policymaker puts almost exclusive weight on the AW model (i.e., $p_{AW} \approx 1$, and $p_m \approx 0$ for all other models). However, this policymaker is sufficiently open-minded to take into account worst-case scenarios from the complete model space. To reflect this open-mindedness we set the ambiguity aversion parameter again to $e = 0.5$.

The performance measures reflect the compromise made by the policymaker. Performance in the more forward-looking models, CW-F, CW-T, and SW, deteriorates relative to minimax outcomes while the performance in the AW model is greatly improved (see panel "AW-leaning" in Table 6). Policy exhibits less interest-rate smoothing and reacts more to the output gap than under minimax. Thus, it has incorporated some of the policy features that were optimal under the AW model (Table 1). As can be inferred from the implied priors that would rationalize this ambiguity averse policy from a Bayesian perspective, the CW-F model still has a strong influence on the worst-case (see the second panel in Table 7). Yet, by construction, the AW model receives more weight relative to the minimax implied priors. The other two models (CW-T or SW) do not influence the worst-case scenarios as indicated by the implied priors.

The AW model and the SW model may be characterized to be at opposite ends of macro-modeling paradigms in use at central banks. In contrast to the former, the SW model emphasizes microeconomic foundations in terms of rational optimizing and forward-looking behavior by households and firms. It

shares these features with many recent macroeconomic models. Thus, we now consider a policymaker who strongly favors the SW paradigm. The results are shown in the third block of Table 6. The "SW-leaning" ambiguity-averse policy exhibits slightly more interest-rate smoothing than the minimax policy. Thus, the SW-leaning ambiguity-averse policy partially embodies some features that were optimal for the SW model (Table 1). As a result, performance in the more forward-looking SW model and in the CW-T model improves relative to minimax outcomes whereas performance in the backward-looking AW model deteriorates.

The implied priors for the SW-leaning ambiguity-averse policy (shown in the third panel of Table 7) illustrate that the CW-F model still influences the worst-case scenarios quite heavily. Interestingly, however, this example also illustrates the endogeneity of the implied priors given to the worst-case scenarios. The AW model has somewhat more influence on the worst-case scenarios for $\lambda_y > 0$ as indicated by the implied priors of the SW-leaning ambiguity-averse policy.

Concluding this section, we review the fault tolerance of the ambiguity-averse policies that lean towards the AW or SW model. The first row of Figure 4 reports the fault tolerance charts under the AW-leaning ambiguity-averse policy for two alternative weights on output stabilization (left panel: $\lambda_y = 0$ and right panel: $\lambda_y = 1$). Relative to the minimax and the Bayesian policy "AW-leaning" ambiguity-averse policy moves towards features of the optimal policy for the AW model. Yet the AW model was instrumental in determining fault intolerance for the minimax and fault tolerance for the Bayesian policies. As a result, the "AW-leaning" ambiguity-averse policy is more fault-tolerant than its Bayesian or minimax counterparts. Interestingly, this is not the case under the "SW-leaning" policy, which moves the ambiguity-averse policy farther away from the prescriptions tolerated by the AW model (see the panels in the second row of Figure 4). Because the ambiguity-averse policy is explicitly designed to bring some priors to bear on the worst-case analysis, the fault-tolerance of this policy depends on which model is given additional weight in the analysis relative to minimax. Ambiguity-averse policy enjoys greater fault tolerance than minimax if flat priors are used. Alternatively, if a specific model receives extra weight, the ambiguity-averse policy exhibits greater fault tolerance within that specific model but possibly lower fault tolerance with respect to other models.

6. ECB Learning About Model Probabilities Since EMU

An ambiguity-averse policy would allow ECB policymakers to take into account model probabilities and insure against worst-case scenarios, while largely retaining the fault-tolerance of the Bayesian policy with flat priors. In addition, it would allow the ECB to use the new data after the start of monetary union to update
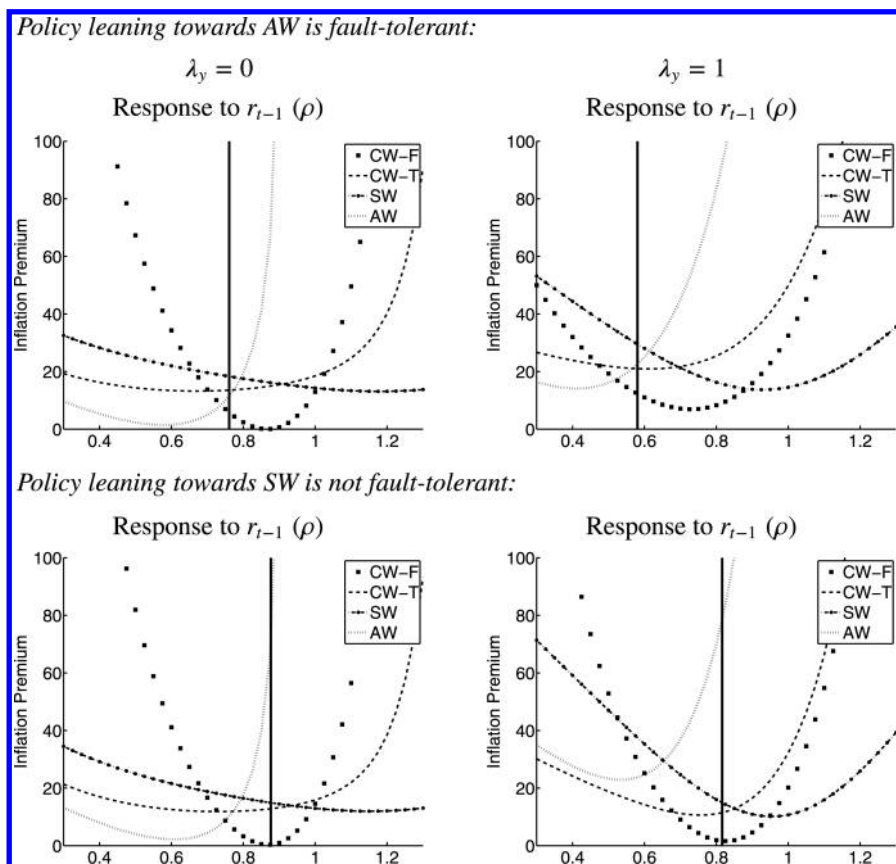
FIGURE 4. Fault tolerance of AW/SW leaning ambiguity-averse policy ($e = 0.5$). For each of the models, a line traces out the IIP (in basis points) under the ambiguity-averse policy as the persistence parameter of the policy rule is varied, holding the other two parameters fixed at their respective values under the full insurance policy. The left (right) panel pertains to the preferences for $\lambda_y = 0$ (1) and $\lambda_{\Delta r} = 0.5$. A vertical line marks the optimal ambiguity-averse parameter value. Shown are two cases giving weight $e = 0.5$ to the worst-case loss: ambiguity-averse policy leaning towards the AW model (top row) and ambiguity-averse policy leaning towards the SW model (bottom row). Respectively, p_{AW} and $p_{SW} \approx 1$.

model probabilities and incorporate them in policy design. We start this section by asking how posterior model probabilities have evolved since 1999. This learning process reveals new information about the likelihood of alternative models of monetary policy transmission in the newly established euro area economy. Second, using this information, we ask whether at different points in time the policy advice derived from a robust policy rule would have notably differed from the results presented in the previous sections.

To be able to compute the likelihood of the four models over time, we make three simplifying assumptions.[20] First, we limit ourselves to the information contained in three time series: output, the short-term interest rate, and consumer price inflation. These series are shared by all models. All of these series have been pre-filtered as described in the notes to Figure A.1 in Appendix A. Second, we leave the model structure as specified by the respective model builders. Thus, we do not change the modeling assumptions over time as information comes in, nor do we change the values of the structural parameters. Third, and closely related, we assume that the ECB followed one and the same policy rule over the entire EMU period. In particular, we use the rule estimated by Gerdesmeier and Roffia (2004) see equation (1). For coherence we use the same policy rule to close each of the models.[21]

Starting with flat priors at the beginning of 1999:Q1 we update posterior model probabilities at the end of each quarter up to 2005:Q4. We are interested in the posterior model probability of model i conditional on these data up to period T, $p(M_i \mid Y^T)$. Let $p(M_i)$ be model priors at the beginning of 1999:Q1. Let $p(Y^T \mid M_i)$ be the likelihood of model i.[22] We assume that the shocks in all four models are drawn from a multi-variate normal distribution. By Bayes's law the model posterior is given by

$$p(M_i \mid Y^T) = \frac{p(Y^T \mid M_i) \cdot p(M_i)}{\sum_{j=1}^{\|\mathcal{M}\|} p(Y^T \mid M_j) \cdot p(M_j)}.$$

The top panel of Figure 5 shows the evolution of the posterior model probabilities over time. Not surprisingly, model probabilities vary over time as new data arrive. However, by 2005:Q4 the CW-T model is much more likely than the other models. Then, the probability assigned to the CW-T model stands at 70%. The other three models carry a posterior probability of about 10% each. Interestingly, the learning process over time implies that models such as the AW or CW-F model that previously influenced the worst-case scenario receive much lower probability weights.

---

20.   The methodology we make use of has certain limitations. In particular, the models' parameter estimates may reflect information held at different points in time and the data are revised (ex post) data. In addition, besides depending on the measure of fit, the relative ranking of the models can depend on the precise sample chosen, on the period chosen to initialize the Kalman filter, and on the choices made for de-trending the data. This sensitivity is well known in the literature (e.g., Levin et al. 2006). We nevertheless view this exercise as valuable as it provides a guide to the stability of our robust policy prescriptions over time.

21.   The data used are from the AWM data set. We use synthetic data from 1984:Q1 to 1998:Q4 to initialize the Kalman filter for each model. Yet only data from 1999:Q1 onwards are used to evaluate the likelihood. The policy rule is appended by a monetary policy shock the standard deviation of which is chosen on the basis of the historical residuals from the Gerdesmeier–Roffia rule.

22.   This formulation makes it clear that we continue to abstract from parameter uncertainty in this exercise. A model is understood to be characterized by its structural equations, the law of motion of shocks and a set of parameters estimated or calibrated by the respective modeler group.
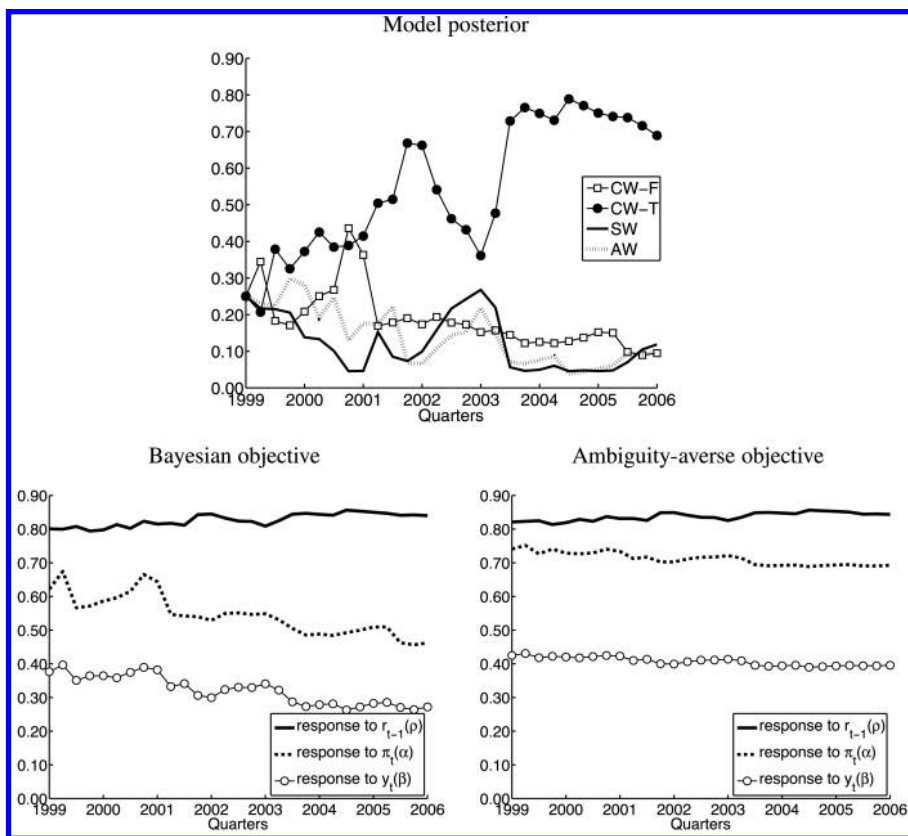
FIGURE 5. Data-driven model posteriors and robust policy rules over time. The top panel shows model posterior probabilities over time. For each of the models, a line traces out the model probability at the beginning of each quarter. Monetary policy is conducted using the Gerdesmeier and Roffia (2004) rule; see equation (1). The policy rule is appended by a monetary policy shock the standard deviation of which is chosen on the basis of the historical residuals from the Gerdesmeier–Roffia rule. Posteriors are computed starting with flat priors at the beginning of 1999:Q1 and using inflation, interest rates, and the output gap to compute the model-likelihood. In each model, data from 1984:Q1 to 1998:Q4 are used to initialize the Kalman-filter. The other two panels show optimal (robust) policy rule coefficients for two objective functions when the data-driven model probabilities are used. The panels show the case of $\lambda_{\Delta r} = 0.5$ and an exclusive weight on inflation stabilization. The bottom left panel shows results for the Bayesian objective. The bottom right panel shows results for the ambiguity-averse objective with data-driven model priors. The latter assumes that $e = 50\%$ of the weight is given to the respective worst-case scenario and the remaining weight is given to expected losses, where expectations are taken using the model probabilities at the respective point in time.

Next, we ask to what extent the prescriptions for rules that are robust to model uncertainty have changed over time. Of course, the pure minimax policy does not change with incoming data. The reason is that all models remain in the model space (i.e., they retain strictly positive weight throughout the entire

sample) and their parameters do not change. By contrast, the Bayesian and the ambiguity-averse policies can incorporate the revised model probabilities obtained with new data. To illustrate this evolution of robust policy, we focus on policies derived from policymaker preferences with zero weight on output stabilization, $\lambda_y = 0$. Our motivation is the hierarchical objective assigned to the ECB by the Maastricht treaty—first comes the pursuit of price stability, and conditional on that the support of other EU objectives.

The evolution of the policy rule coefficients, (ρ, α, β), consistent with Bayesian policy is shown in the bottom left panel of Figure 5. Despite the gap opening between model likelihoods over the EMU period, the Bayesian policy rule prescriptions change only quite slowly. Interest rate smoothing increases slightly over time, and the weight given to inflation and the output gap in the policy rule declines a bit. The bottom right panel shows the parameters of ambiguity-averse policy when $e = 50\%$ of the weight is given to the worst-case and the remaining weight given to expected losses. As in the case of the Bayesian policy, expectations are computed using the time-dependent model probabilities at the respective point in time. Again the policymaker preference is such that no weight is placed on output stabilization, $\lambda_y = 0$. As can be seen from the right-hand panel, the prescriptions regarding robust policy remain valid even though the posterior probabilities associated with the models that tend to drive worst-case scenarios (CW-F and AW) decline over time: under model uncertainty, policymakers will favor a rule with some but not excessive interest rate smoothing and non-negligible weight on the output gap in the rule—even if stabilizing output is not an explicit policy objective of the central bank.

7. Conclusions

Considering four alternative models of the euro area that have been developed and used at the ECB, we have found that maximum insurance against model uncertainty via minimax policy comes at moderate costs in terms of lower expected performance. However, the implied priors that would rationalize the minimax policy from a Bayesian perspective indicate that minimax-type insurance is strongly oriented towards a single model with high baseline losses. As a consequence, the minimax policy is not as tolerant towards small perturbations of policy parameters as the Bayesian policy rule. Thus, we propose to strike a compromise and use so-called ambiguity-averse preferences for policy design. These preferences allow the specification of prior model probabilities but also give extra weight to the worst uncertain scenarios. Given flat prior model probabilities, these policies inherit the fault tolerance of the Bayesian policies, while guaranteeing some extra insurance against worst-case losses.

The implied priors and the implied inflation variability premium that we have computed are intended to help policymakers in refining their views on alternative reference models and in reconsidering their desired degree of insurance. For

practical application, an iterative procedure using implied priors and calculating costs of alternative degrees of insurance should be a useful tool for policy design.

We also showed how policymakers can learn about the likelihood of alternative models of the euro area from data since the start of EMU. We have found that whereas data-driven posterior model probabilities evolved over time, the characteristics of robust policy rules did not change very much: These rules are not overly persistent and, at the same time, they also give some weight to stabilizing output even if the central bank is only concerned with inflation stabilization.

For future work, it would be of interest to extend the analysis to allow for the emergence of new models and parameterizations in the central bank's learning process. Re-estimation of macroeconomic models used in policy analysis typically requires a major vetting effort. In the type of analysis presented in the preceding section, we would treat a new parametrization of a given model as a new model and then investigate to what extent its posterior model probabilities increase relative to older models with new data.

A full analysis of the implications of parameter uncertainty and learning within each model is beyond the scope of this paper. Under multiplicative parameter uncertainty the principle of certainty equivalence does not hold and standard linear-quadratic solution methods cannot be used anymore. For an analysis of Brainard-style conservatism and the countervailing experimental activism under learning we refer the reader to Wieland (2000, 2006), who provides a quantitative assessment in smaller backward-looking and partially forward-looking models, and to Sack (2000), who studies Brainard-style conservatism in a VAR. Minimax policies with multiplicative parameter uncertainty in a small model are studied by Zakovic, Rustem, and Wieland (2007).

Another issue to be investigated in the future concerns the uncertainty regarding the unobserved variables determining the long-run model equilibrium. The uncertainty we have considered by means of the linearized benchmark models concerned the dynamics around a given steady state. In the future, it would be interesting to consider also uncertainty about the steady state itself. This will mean including uncertainty with regard to such variables as the natural rates of interest and unemployment, and the level of potential output.

Appendix A: The Models and Model Uncertainty

To provide a sense of the differences between the four models we compare empirical and model-based autocorrelation functions and impulse responses. In deriving these measures we consider two alternative assumptions regarding the policy rule implemented by the ECB. The ECB's rule is either set equal in all models to the empirically estimated rule by Gerdesmeier and Roffia (2004) or it is set differently in each model and equal to the specific rule estimated along with the other equations in that model. Figure A.1 reports the autocorrelation functions.
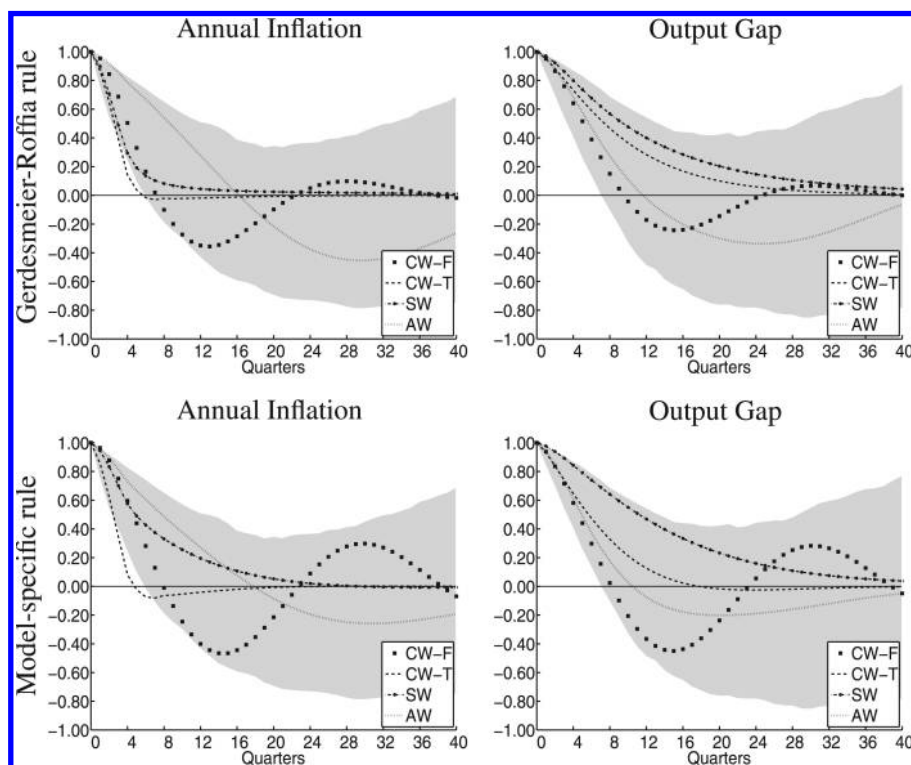
FIGURE A.1. Autocorrelation functions. Shown are the autocorrelations of inflation and the output gap under two sets of policy rules. The first row shows autocorrelations when the Gerdesmeier and Roffia (2004) rule is implemented in all models; see equation (1). The second row shows autocorrelations when in each model a model-specific rule is implemented. For the SW model the rule estimated in Smets and Wouters (2003) is used. For the CW-F and CW-T we use data on output, inflation, and interest rates to estimate a rule of the form of equation (2) by maximum likelihood from 1984:Q1 to 2005:Q4, taking all parameters apart from the parameters of the monetary rule as given. For the AW model, a standard Taylor rule with $\rho = 0$, $\alpha = 1.5$, and $\beta = 0.5$ is used. Shaded areas mark 95% confidence intervals for the autocorrelation in the data. Autocorrelation bounds have been obtained by bootstrapping from an estimated trivariate VAR(1). The data are from the AWM data set and range from 1984:Q1 to 2005:Q4. Log output is linearly detrended over this sample. Year-on-year CPI inflation and the quarterly interest rate (annualized) have been linearly de-trended until 1998:Q4 to account for the convergence period prior to EMU.

As a basis for gauging the autocorrelations empirically, the figure also shows 95% confidence bands for the autocorrelation function implied by the actual data (a shaded area). These bands have been obtained by bootstrapping from a trivariate first-order vector autoregression (VAR) in inflation, the output gap and the nominal rate; the data have been de-trended prior to running the VAR; see the notes to Figure A.1 for details. The four models imply quite different dynamics thereby spanning a significant range of model uncertainty. In the CW-T

and SW models inflation dynamics largely die out within 4 to 6 quarters, whereas the CW-F model generates somewhat longer-lasting dynamics. The AW model stands out with the highest degree of inflation persistence. Output is strongly serially correlated in all models, as it is in the data, but shows fairly different dynamic patterns. In spite of these differences all four models' autocorrelations fall within the 95% confidence intervals—a few periods in the serial correlation pattern of the CW-T model apart.

The second row of Figure A.1 shows autocorrelation functions when model-specific rules are implemented instead.[23] This means that now not only do the model dynamics bear on the difference in model behavior but also the difference in rules. The models continue to show dynamics with substantial differences which are, however, by and large covered by the autocorrelation bands.

To illustrate the extent of uncertainty about policy effectiveness, we report impulse response functions for inflation and the output gap in Figure A.2. We simulate a 100 basis point shock to the interest rate (annualized) assuming that monetary policy subsequently follows the Gerdesmeier–Roffia rule (first row) or the model-specific policy rules (second row). In order to allow the reader to benchmark the models against outside evidence, we also show symmetric 95% confidence bands obtained from estimating a structural vector autoregression (SVAR) for the euro area with the often-used Cholesky identification assumption for a monetary policy shock; for details see the notes to Figure A.2. These bands provide a rough guide to the magnitude of the effect of a monetary innovation.[24]

Starting with the common policy rule (first row), again, the AW model generates the highest degree of endogenous persistence with inflation reaching its trough more than 5 years after the shock. Similarly, output returns only very gradually to baseline in this model. The SW model marks the smallest degree of output and inflation persistence. The CW-F and CW-T models lie in between in terms of persistence of the dynamics in response to a policy shock. As expected, Taylor-contracts generate less inflation persistence than Fuhrer–Moore contracts but output is more persistent in the CW-T model. The responses to a monetary shock under the common rule do not fully coincide with the response to the monetary shock identified in the VAR, yet overall magnitudes are reasonable for

---

23. For the SW model, we use the rule estimated in Smets and Wouters (2003) . For the CW-F and CW-T we use data on output, inflation, and interest rates to estimate model-specific rules of the form $r_t = \rho r_{t-1} + \alpha \pi_t + \beta y_t$ by maximum likelihood from 1984:Q1 to 2005:Q4, taking all parameters apart from the parameters of the monetary rule as given. The results are for CW-F: $\rho = 0.88$, $\alpha = 0.12$, $\beta = 0.22$; for CW-T: $\rho = 0.82$, $\alpha = 0.19$, $\beta = 0.18$. For the AW, we use the original Taylor rule, which was used intensively in the calibration of the model, namely, $r_t = 1.5\pi_t + 0.5y_t$.

24. In judging the models' responses against the identified responses in the VAR we ask the reader to bear in mind that none of the models were constructed to directly fit this evidence. As a consequence, strictly speaking, none of the models satisfy the identification assumptions in the structural VAR as inflation responds instantaneously to a monetary innovation in all four models.
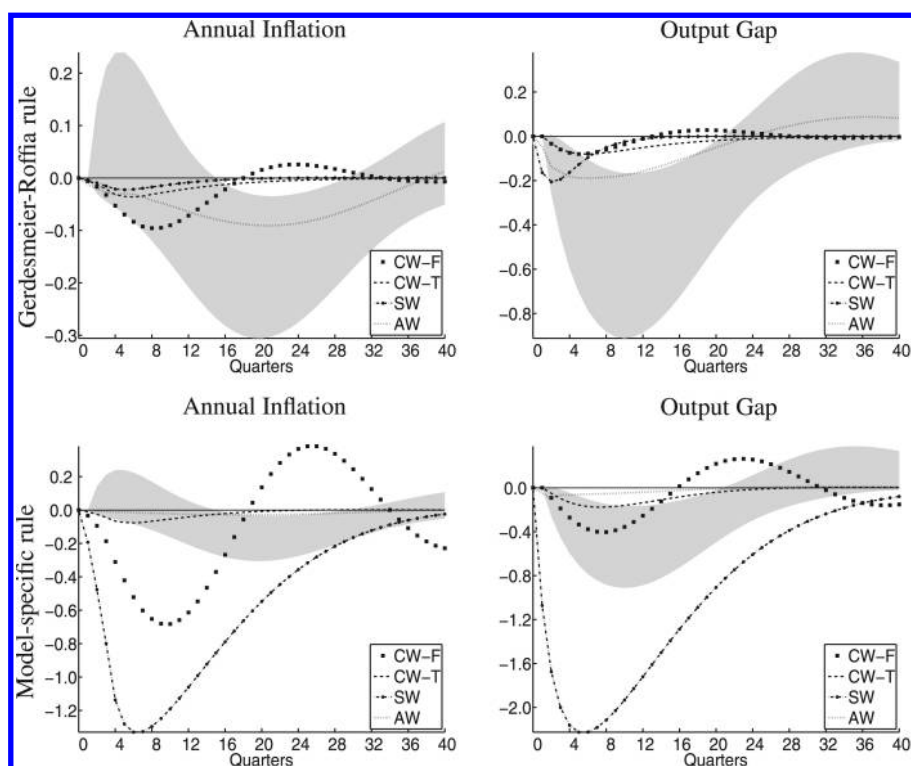
FIGURE A.2. Impulse responses. Impulse responses to a 100bps temporary unanticipated rise in the nominal short term rate (in percentage points annualized). The first line shows responses when, in all models, monetary policy follows the Gerdesmeier and Roffia (2004) rule, which is given by equation (1). All variables are measured in percentage points. The second line shows impulse-responses when in each model monetary policy follows a model-specific rule. For the SW model the rule estimated in Smets and Wouters (2003) is used. For the CW-F and CW-T model the rule was estimated by maximum likelihood from 1984:Q1 to 2005:Q4 taking the remaining parameter values of the model as given. For the AW model, the original Taylor rule is used, namely, $r_t = 1.5\pi_t + 0.5y_t$. Also shown (as shaded areas) are 95% confidence bands from an identified VAR(1) in CPI inflation, output, and interest rates in log levels, estimated on the sample 1984:Q1 to 2005:Q4. The VAR includes a constant, a time-trend, and a broken trend (breaking after 1998:Q4) to account for the nominal convergence prior to EMU.

the inflation response—if at the lower end of the spectrum for most of the models. On the other hand, the model-specific rules illustrate that all models assign a powerful role to monetary policy: The form of systematic monetary policy has a strong impact on the behavior of each model economy.[25]

---

25. On the basis of Figure A.2 one might therefore conjecture that once the policy rule was parameterized with that target in mind, each model could be made to fit the identified impulse responses better than Figure A.2 suggests.

TABLE A.1. Standard deviations.

	Data		Gerdesmeier–Roffia rule				Model-specific rule			
	Avg. std. dev.	[95%]	CW-F	CW-T	SW	AW	CW-F	CW-T	SW	AW
π	0.7	[.4–1.1]	2.2	1.1	0.8	1.3	2.7	1.0	1.1	1.1
y	1.4	[.8–2.3]	2.0	1.9	2.1	2.5	1.8	1.6	2.6	2.1
Δr	0.5	[.4– .6]	0.9	0.5	0.4	0.4	0.5	0.3	0.6	0.5

Notes: Unconditional standard deviations of the target variables in the four models when the policymaker commits himself to using either the Gerdesmeier and Roffia rule (1) or a model-specific rule. All variables are measured in percentage points. Inflation is in annual terms and the change in the nominal rate, Δr, is quarterly annualized. For the model-specific rules, the SW model uses the rule estimated by Smets and Wouters (2003); the CW-F and CW-T models use a rule which was estimated by maximum likelihood taking the other parameter values of these models as given and using data on output, CPI inflation, and interest rates from 1984:Q1 to 2005:Q4. For the AW model, a standard Taylor rule is used, $r_t = 1.5\pi_t + 0.5y_t$. Also shown are the mean standard deviations in the de-trended data along with 95% confidence intervals (in square brackets) extracted from bootstrapping a VAR(1), which was estimated on the de-trended data from 1984:Q1 to 2005:Q4 as in Figure A.1.

Table A.1 compares the unconditional second moments of inflation, the output gap, and the change in the nominal interest rate, π, y, and Δr, in these models under the two sets of rules to moments measured in the data. Standard deviations and bounds for the data are again obtained by bootstrapping from a first-order VAR.[26]

Under the Gerdesmeier–Roffia rule the CW-F model induces the highest variance for inflation among the four models—more than twice as much as the SW model—and consequently it induces a rather high variability of interest rate increments. All models generate a standard deviation of the output gap that is larger than the mean value observed in the data. Also under the model-specific rules the CW-F model shows the largest degree of inflation volatility. Under this set of rules the SW model shows more output volatility than the AW.

Summarizing, this analysis highlights three important points for this paper. First, each model fits certain stylized facts in the data. Some models fit the cross-correlations better than the standard deviations, others are better at reproducing the stylized response of the economy to a monetary shock than to fit the former moments. Each model thus can be seen as a reasonable description of the economy. Yet, none of the models is *the* perfect model of the economy: The ordering of models may differ depending on the metric used in estimation and in evaluating the fit.[27] Second, despite all models being reasonable to a certain extent, they

---

26. All models are linearized around their respective steady state/balanced growth path values. We use the Andersen–Moore algorithm to compute the equilibrium law of motion and the doubling algorithm to obtain the variances of endogenous variables for a given parametrization of the monetary policy rule and given the remaining parameters of the model. The standard deviations reported thus are unconditional standard deviations of the variables in the stochastic steady state. Because the models are linearized, the second moments are accurate up to second order only.

27. Interestingly, for example, the model doing best with respect to the response of inflation and output to a monetary innovation in Figure A.2 when closed by a model-specific rule is the model which did worst when comparing the autocorrelation of inflation implied by the model to the data in Figure A.1: the CW-T model.

all show substantial differences when it comes to the implied effectiveness of monetary policy, to intrinsic persistence, and to the implied standard deviations of key variables. As a result of the former two points there is, third, considerable model uncertainty for monetary policy in the euro area.

Appendix B: Optimized Simple Rules from the Literature

This Appendix presents coefficients for optimized simple Taylor-style rules as reported in other studies. It is meant to enable the reader to gauge how representative the set of models used in this study is.

The optimized simple rules reported in Table B.1 were taken from studies available in the literature. The coefficients shown belong to the rules reported in these studies that came closest to our case of $\lambda_{\Delta r} = 0.5$. Also we constrain ourselves to rules that fall within the class of rules covered in our paper. The notes to Table B.1 provide further information.

TABLE B.1. Optimized simple rules in the literature.

	$\lambda_y = 0$			$\lambda_y \geq 1$			
	ρ	α	β	ρ	α	β	Model
U.S.							
Levin et al. (1999)[a]	1.0	0.6	0.1	0.8	0.2	0.7	FM, Fuhrer and Moore (1995).
	1.1	3.7	1.0	1.0	0.2	1.3	MSR, Orphanides and Wieland (1998).
	0.9	1.4	0.1	1.0	0.1	1.2	FRB, Brayton et al. (1997).
	0.9	2.8	0.6	1.0	0.1	1.5	TMCM, Taylor (1993).
Levin and Williams (2003)	1.5	6.4	0	1.5	1.0	2.5	New Keynesian benchmark model.
	1.0	0.7	0	0.9	0.6	0.8	Fuhrer (2000).
	0.4	2.3	0.9	0.1	2.4	1.5	Rudebusch and Svensson (1999).
	0.4	1.7	0.3	0.6	1.2	1.1	Bayesian robust rule (above models).
	0.3	2.2	0.4	0.9	0.6	0.8	Minimax robust rule (above models).
Euro area							
Adalid et al. (2005)[b]	0.3	6.5	3.7	0.3	6.5	4.3	Angelini et al. (2004).
	0.5	3.8	2.3	0.6	3.6	3.3	Bayesian robust rule (four models).
Canada							
Cote et al.(2004)[c]	0	2-4	.5-2	–	–	–	MULTIMOD, NAOMI, QPM.
	0.9	.1-.2	0-.1	–	–	–	LPM, M1-VECM.

[a] The coefficient values were read off Fig. 6.3 in Levin et al. (1999). In this figure, the parameters refer to a zero weight on output $\lambda_y = 0$ through to an exclusive weight on the output stabilization objective. The numbers reported here in column $\lambda_y \geq 1$ refer to the latter case.

[b] The study uses three of the models employed here (CW-T, SW, AW) plus a backward-looking three-country model along the lines of Rudebusch and Svensson (1999) used at the Banca d'Italia. The study uses a lower weight on interest rate smoothing than the current paper.

[c] The study compares the robustness of optimized simple rules coming from five models of the Canadian economy in 12 models of the Canadian economy. The loss function does not contain an interest rate smoothing term, so $\lambda_{\Delta r} = 0$. λ_y is set to 0.25. Reported is the range of parameters for the rules proposed, grouping the models into two sets of models which deliver similar prescriptions for the rule.

We also report the rules used by Taylor (1999), which are by and large representative for optimal rules studied in the literature. In nine macroeconomic models for the U.S. Taylor evaluates the following five rules of the general form: $r_t = \rho r_{t-1} + \alpha \pi_t + \beta y_t$: $[\rho, \alpha, \beta] \in \{[1.0, 3.0, 0.8], [1.0, 1.2, 1.0], [0.0, 1.5, 0.5], [0.0, 1.5, 1.0], [1.3, 1.2, 0.06]\}$.

Appendix C: The Minimax Algorithm

We employ the discrete minimax algorithm as implemented in MATLAB 6.1 as `fminimax`. The algorithm is based on Brayton et al. (1979). We briefly introduce the idea behind the algorithm.

Quasi-Newton methods solve constrained optimization problems stepwise, finding the optimal step $s\Delta_k$ starting from a point x_k. To illustrate this, consider the equality constrained problem

$$\min_{\mathbf{x}} f(\mathbf{x}), \quad \text{subj. to: } \mathbf{g}(\mathbf{x}) = \mathbf{0}. \tag{C.1}$$

Let $L(\mathbf{x}, \boldsymbol{\lambda})$ be the corresponding Lagrangian. The Kuhn–Tucker conditions state that there exists a non-negative vector $\boldsymbol{\lambda}^*$ such that at the optimum point $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ the constraints are met and

$$\mathbf{h}(\mathbf{x}^*, \boldsymbol{\lambda}^*) := \left.\frac{\partial f}{\partial \mathbf{x}}\right|_{\mathbf{x}^*} + \left.\frac{\partial \mathbf{g}'}{\partial \mathbf{x}}\right|_{\mathbf{x}^*} \boldsymbol{\lambda}^* = \mathbf{0}.$$

Expanding around $\mathbf{x}_k := \mathbf{x}^* - \boldsymbol{\Delta}_k$, where $\boldsymbol{\Delta}_k$ is "small,"

$$\mathbf{0} = \mathbf{h}(\mathbf{x}^*, \boldsymbol{\lambda}^*) \simeq \mathbf{h}(\mathbf{x}_k, \boldsymbol{\lambda}^*) + f_{\mathbf{xx}}|_{\mathbf{x}_k} \boldsymbol{\Delta}_k + \sum \lambda_i^* \, g_{i\mathbf{xx}}|_{\mathbf{x}_k} \boldsymbol{\Delta}_k,$$

$$\mathbf{0} = \mathbf{g}(\mathbf{x}^*) \simeq \mathbf{g}(\mathbf{x}_k) + \left.\frac{\partial \mathbf{g}}{\partial \mathbf{x}'}\right|_{\mathbf{x}_k} \boldsymbol{\Delta}_k. \tag{C.2}$$

System (C.2) is the same system of equations as the first order conditions derived by solving the following quadratic subproblem (QSP), provided Δ_k is small enough so $\boldsymbol{\lambda}_k^{qsp} \simeq \boldsymbol{\lambda}^*$, where $\boldsymbol{\lambda}^{qsp}$ are the multipliers of the QSP

$$\min_{\boldsymbol{\Delta}_k} \quad \frac{1}{2} \boldsymbol{\Delta}_k' \mathbf{H} \boldsymbol{\Delta}_k + \left.\frac{\partial f}{\partial \mathbf{x}'}\right|_{\mathbf{x}_k} \boldsymbol{\Delta}_k,$$

$$\text{subj. to:} \quad \mathbf{g}(\mathbf{x}_k) + \left.\frac{\partial \mathbf{g}}{\partial \mathbf{x}'}\right|_{\mathbf{x}_k} \boldsymbol{\Delta}_k = \mathbf{0}, \tag{C.3}$$

where $\mathbf{H} := L_{\mathbf{xx}}|_{\mathbf{x}_k, \boldsymbol{\lambda}^*}$. An optimum is found by a sequence of updates $\mathbf{x}_{k+1} = \mathbf{x}_k + s\boldsymbol{\Delta}_k$, and $\boldsymbol{\lambda}_{k+1} = (1-s)\boldsymbol{\lambda}_k + s\boldsymbol{\lambda}_k^{qsp}$. Stepsize s is determined on the basis of a problem-specific merit function.

The minimax problem

$$\min_{\mathbf{x}} \max_{j} \quad f_j(\mathbf{x}), \quad \text{subj. to } \mathbf{g}(\mathbf{x}) \leq \mathbf{0}, \tag{C.4}$$

where some of the constraints in $\mathbf{g}(\cdot)$ may be strict equality constraints, can be restated as a non-linear optimization problem

$$\min_{\mathbf{x}, \gamma} \gamma, \quad \text{subj. to: } \mathbf{g}(\mathbf{x}) \leq \mathbf{0} \text{ and } f_j(\mathbf{x}) \leq \gamma \; \forall j. \tag{C.5}$$

Let $\mathbf{y} := (\gamma, \mathbf{x}')'$, $f(\mathbf{y}) := \gamma$ and stack the constraints $f_j(\mathbf{x}) \leq \gamma$ as $\mathbf{b}(\mathbf{y}) \leq \mathbf{0}$. The Lagrangian for condition (C.5) then reads as $L(\mathbf{y}, \lambda) = f(\mathbf{y}) + \lambda_1' \mathbf{g}(\mathbf{y}) + \lambda_2' \mathbf{b}(\mathbf{y})$. Analogous to condition (C.3), step $\Delta_k^{(y)}$ is found by solving the quadratic subproblem using active set methods (let $\tilde{\mathbf{b}}$ mark the binding constraints).

$$\min_{\Delta_k^{(y)}} \ \Delta_k^{(\gamma)} + \frac{1}{2} \Delta_k^{(\mathbf{x})'} \mathbf{H_x} \Delta_k^{(\mathbf{x})},$$

$$\text{subj. to:} \quad \mathbf{g}(\mathbf{y}_k) + \left.\frac{\partial \mathbf{g}}{\partial \mathbf{y}'}\right|_{\mathbf{y}_k} \Delta_k^{(\mathbf{y})} = \mathbf{0}, \tag{C.6}$$

$$\tilde{\mathbf{b}}(\mathbf{y}_k) + \left.\frac{\partial \tilde{\mathbf{b}}}{\partial \mathbf{y}'}\right|_{\mathbf{y}_k} \Delta_k^{(\mathbf{y})} = \mathbf{0},$$

where $H_\mathbf{x} := L_{\mathbf{xx}}|_{\mathbf{x}_k, \lambda_k}$. For the minimax problem any second (cross) derivative of L involving γ is zero. In `fminimax`, the step-size s in updating $\mathbf{y}_{k+1} := \mathbf{y}_k + s\Delta_k^{(\mathbf{y})}$ has to lead to an improvement in either of two merit functions. Let $\mathcal{E}$ be the set of strict equality constraints in $\mathbf{g}$. The merit functions evaluated at a guess $\mathbf{y}(s) = \mathbf{y}_k + s\Delta_k^{(\mathbf{y})}$ are of the form

$$P(\mathbf{y}, \boldsymbol{\mu}) = \text{Loss}(\mathbf{y}) + \sum_{i \in \mathcal{E}} \mu_{1,i} |g_i(\mathbf{x})| + \sum_{i \notin \mathcal{E}} \mu_{1,i} \max(g_i(\mathbf{x}), 0)$$

$$+ \sum_j \mu_{2,j} \max(b_j(\mathbf{y}), 0),$$

where we suppress subscript k for convenience. Here, $\boldsymbol{\mu}_k = \max[\lambda_k^{qsp}, \frac{1}{2}(\lambda_k^{qsp} + \boldsymbol{\mu}_{k-1})]$. The first merit function, exact for the constrained problem (C.5), sets $\text{Loss}(\mathbf{y}) = \gamma$, whereas the second merit function sets $\text{Loss}(\mathbf{y}) = \max_j f_j(\mathbf{x})$, which makes use of to the minimax problem structure (Brayton et al. 1979). If a step length is found such that either of the penalty functions signals improvement, $\mathbf{y}$ and λ are updated accordingly.

As with all Quasi-Newton optimization methods, convergence may occur only to a local minimum. We entertain several starting values to safeguard the results against this feature.

Appendix D: Bayesian Priors From Unconstrained Minimax

The aim of this section is to explain how priors $\tilde{p}_i$ can be found such that

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \max_i f_i(\mathbf{x}) = \arg\min_{\mathbf{x}} \sum_i \tilde{p}_i f_i(\mathbf{x}).$$

In the form of condition (C.5), the first-order conditions to the unconstrained $(\mathbf{g} \equiv \mathbf{0})$ minimax problem are

$$
D_{\mathbf{y}}L := \frac{\partial L}{\partial \mathbf{y}} = \begin{pmatrix} 1 - \sum_i \lambda_i \\ \frac{\partial \sum_i f_i(\mathbf{x})\lambda_i}{\partial \mathbf{x}} \end{pmatrix} = \mathbf{0}.
$$

Define

$$
D_{\mathbf{y}}^2 L := \frac{\partial L}{\partial \mathbf{y} \partial \mathbf{y}'} = \begin{pmatrix} 0 & 0 \\ 0 & \frac{\partial \sum_i f_i(\mathbf{x})\lambda_i}{\partial \mathbf{x} \partial \mathbf{x}'} \end{pmatrix}.
$$

So for any local minimum of the unconstrained minimax, the first and second-order conditions for a local minimum of the unconstrained Bayesian optimization are also met, provided one takes $\boldsymbol{\lambda}^*$ (from the final step of the QSP for the minimax) as the Bayesian priors. Priors $\tilde{p}_i = \lambda_i^*$ support the minimax allocation, $\mathbf{x}^*$. For the case of intermediate ambiguity aversion, the problem is $\min_{\mathbf{x}} \max_i \bar{f}_i(\mathbf{x})$, where $\bar{f}_i(\mathbf{x}) = (1-e)\sum_i p_i f_i(\mathbf{x}) + e f_i(\mathbf{x})$. Substituting for $\bar{f}_i(\cdot)$, the first order conditions are

$$
\sum_i \bar{\lambda}_i^* = 1 \quad \text{and} \quad \sum_i \left\{ (1-e)p_i + e\bar{\lambda}_i^* \right\} \frac{\partial f_i(\mathbf{x})}{\partial \mathbf{x}} \bigg|_{\mathbf{x}^*} = 0 \quad \forall i.
$$

Priors $\tilde{p}_i := (1-e)p_i + e\bar{\lambda}_i^*$ support the ambiguity averse allocation.

Appendix E: Alternative Weights on Interest Variability, $\lambda_{\Delta r}$

This Appendix provides the results in the paper for different degrees of interest rate smoothing, $\lambda_{\Delta r} = 0.1$ and $\lambda_{\Delta r} = 1$.

TABLE E.1. Optimized simple rules for $\lambda_{\Delta r} = .1$ and $\lambda_{\Delta r} = 1$.

Model	$\lambda_y = 0$			$\lambda_y = 0.5$			$\lambda_y = 1$		
	ρ	α	β	ρ	α	β	ρ	α	β
$\lambda_{\Delta r} = .1$									
CW-F	0.9	1.7	0.6	0.8	1.4	1.0	0.8	1.3	1.3
CW-T	1.0	0.7	0.1	0.8	0.4	1.1	0.8	0.3	1.6
SW	1.0	0.9	0.0	1.0	0.5	2.0	1.0	0.4	3.1
AW	0.4	1.2	1.2	0.4	0.8	2.8	0.5	0.6	3.7
$\lambda_{\Delta r} = 1$									
CW-F	0.9	0.6	0.4	0.8	0.5	0.5	0.8	0.5	0.6
CW-T	1.0	0.2	0.1	0.9	0.1	0.4	0.9	0.1	0.6
SW	1.0	0.3	0.0	1.0	0.2	0.6	1.0	0.1	0.8
AW	0.7	0.4	0.4	0.5	0.5	0.9	0.4	0.5	1.2

Note: Parameters of optimal simple rules for each model. Shown is the case $\lambda_{\Delta r} = 1$ for loss function (3).

TABLE E.2. Robustness of rules optimized for a specific model $\lambda_{\Delta r} = 0.1$.

| | CW-F rule evaluated in | | | | | | CW-T rule evaluated in | | | | | |
| | CW-T | | SW | | AW | | CW-F | | SW | | AW | |
λ_y	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP
0.0	25	*.11*	40	*.12*	∞	∞	81	*.53*	9	*.03*	∞	∞
0.5	14	*.10*	26	*.13*	109	*.75*	88	*.84*	12	*.06*	34	*.29*
1.0	15	*.14*	38	*.20*	94	*.81*	179	*1.71*	17	*.09*	37	*.39*

| | SW rule evaluated in | | | | | | AW rule evaluated in | | | | | |
| | CW-F | | CW-T | | AW | | CW-F | | CW-T | | SW | |
λ_y	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP
0.0	160	*.93*	7	*.03*	∞	∞	89	*.57*	32	*.14*	79	*.21*
0.5	86	*.82*	17	*.13*	130	*.85*	155	*1.31*	12	*.09*	25	*.12*
1.0	123	*1.29*	27	*.24*	95	*.82*	257	*2.22*	15	*.14*	25	*.14*

Notes: %$\mathcal{L}$-percentage losses relative to first-best simple rule for each model (in %); IIP-Implied Inflation (Variability) Premium relative to first-best simple rule for each model (in percentage points) (*in italics*). The notation ∞ indicates that the implemented rule results in instability. Shown is the case $\lambda_{\Delta r} = 0.1$ for loss function (3).

TABLE E.3. Robustness of rules optimized for a specific model $\lambda_{\Delta r} = 1$.

| | CW-F rule evaluated in | | | | | | CW-T rule evaluated in | | | | | |
| | CW-T | | SW | | AW | | CW-F | | SW | | AW | |
λ_y	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP
0.0	37	*.18*	69	*.21*	87	*.44*	265	*1.82*	11	*.04*	∞	∞
0.5	12	*.11*	15	*.11*	69	*.76*	87	*.96*	15	*.11*	26	*.35*
1.0	11	*.13*	22	*.19*	63	*.95*	105	*1.25*	16	*.14*	40	*.67*

| | SW rule evaluated in | | | | | | AW rule evaluated in | | | | | |
| | CW-F | | CW-T | | AW | | CW-F | | CW-T | | SW | |
λ_y	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP	%$\mathcal{L}$	IIP
0.0	748	*3.72*	9	*.04*	∞	∞	81	*.71*	22	*.11*	82	*.24*
0.5	63	*.73*	20	*.18*	4299	*10.15*	171	*1.63*	17	*.15*	59	*.37*
1.0	94	*1.17*	26	*.29*	310	*2.79*	ME	ME	ME	ME	67	*.49*

Notes: %$\mathcal{L}$-percentage losses relative to first-best simple rule for each model (in %); IIP-Implied Inflation (Variability) Premium relative to first-best simple rule for each model (in percentage points) (*in italics*). The notation ∞ indicates that the implemented rule results in instability; the notation ME indicates that the implemented rule results in multiple equilibria. Shown is the case $\lambda_{\Delta r} = 1$ for loss function (3).

TABLE E.4. Flat Bayesian priors versus model-specific simple rules different weights $\lambda_{\Delta r}$.

	Optimal rule			Inflation premium				Percentage losses			
λ_y	ρ	α	β	CW-F	CW-T	SW	AW	CW-F	CW-T	SW	AW
$\lambda_{\Delta r} = .1$											
0.0	0.7	1.3	0.7	.09	.10	.14	.13	12	22	49	30
0.5	0.7	1.5	1.7	.11	.07	.09	.21	9	9	19	23
1.0	0.6	1.6	2.4	.12	.11	.13	.24	8	11	24	21
$\lambda_{\Delta r} = 1$											
0.0	0.8	0.5	0.3	.07	.11	.18	.12	7	23	60	21
0.5	0.8	0.5	0.5	.08	.08	.16	.28	6	9	23	20
1.0	0.7	0.6	0.8	.10	.12	.25	.36	6	10	30	19

Notes: Optimal policy rule parameters and IIP relative to the simple rules optimized for each model (in percentage points). Shown are the cases $\lambda_{\Delta r} = 0.1$ and $\lambda_{\Delta r} = 1$ for loss function (3). For comparison with the literature, we also report the percentage increase in losses relative to the simple rules optimized for each model.

TABLE E.5. Minimax policy relative to model-specific rules different weights $\lambda_{\Delta r}$.

	Optimal rule			Inflation premium				Loss premium	
λ_y	ρ	α	β	CW-F	CW-T	SW	AW	$\Delta\mathcal{L}^{worst}$	$\Delta\mathcal{L}^{expect}$
$\lambda_{\Delta r} = .1$									
0.0	0.8	1.7	0.7	.01	.13	.13	.54	9	16
0.5	0.8	1.4	1.0	.00	.10	.13	.75	8	14
1.0	0.8	1.3	1.3	.00	.14	.20	.81	8	12
$\lambda_{\Delta r} = 1$									
0.0	0.9	0.6	0.4	.00	.18	.21	.44	6	11
0.5	0.8	0.6	0.5	.00	.11	.11	.73	6	8
1.0	0.8	0.6	0.6	.01	.14	.22	.70	5	4

Notes: Optimal policy rule parameters and IIP relative to first-best simple rule for each model (in percentage points). Shown are the cases $\lambda_{\Delta r} = 0.1$ and $\lambda_{\Delta r} = 1$ for loss function (3). $\Delta\mathcal{L}^{worst}$-percentage reduction of worst-case loss relative to worst outcome under flat Bayesian priors; $\Delta\mathcal{L}^{expect}$-percentage increase in expected loss relative to Bayesian policy (flat priors).

TABLE E.6. Minimax implied priors different weights $\lambda_{\Delta r}$.

	$\lambda_{\Delta r} = 0.1$				$\lambda_{\Delta r} = 1$			
λ_y	CW-F	CW-T	SW	AW	CW-F	CW-T	SW	AW
0.0	0.971	0.0	0.0	0.029	1.000	0.0	0.0	0.000
0.5	1.000	0.0	0.0	0.000	0.994	0.0	0.0	0.006
1.0	1.000	0.0	0.0	0.000	0.923	0.0	0.0	0.077

Notes: Bayesian priors backing the minimax solution. Shown are the cases $\lambda_{\Delta r} = 0.1$ and $\lambda_{\Delta r} = 1$ for loss function (3).

TABLE E.7. Ambiguity-averse policy for different weights $\lambda_{\Delta r}$.

	Optimal rule			Inflation premium			
λ_y	ρ	α	β	CW-F	CW-T	SW	AW
Flat priors, $\lambda_{\Delta r} = 0.1$							
0.0	0.8	1.6	0.7	.03	.12	.14	.28
0.5	0.7	1.5	1.3	.02	.09	.11	.41
1.0	0.7	1.6	1.8	.02	.12	.17	.47
Flat priors, $\lambda_{\Delta r} = 1$							
0.0	0.8	0.5	0.4	.01	.15	.20	.24
0.5	0.8	0.6	0.5	.02	.10	.14	.46
1.0	0.8	0.6	0.7	.02	.13	.23	.59
AW-leaning, $\lambda_{\Delta r} = 0.1$							
0.0	0.7	1.7	0.8	.09	.14	.16	.12
0.5	0.5	2.1	1.9	.10	.12	.13	.18
1.0	0.5	2.2	2.7	.10	.17	.18	.21
AW-leaning, $\lambda_{\Delta r} = 1$							
0.0	0.8	0.6	0.4	.05	.15	.21	.12
0.5	0.7	0.7	0.6	.10	.15	.22	.21
1.0	0.6	0.8	0.9	.13	.22	.34	.25
SW-leaning, $\lambda_{\Delta r} = .1$							
0.0	0.8	1.6	0.7	.01	.12	.13	.54
0.5	0.9	1.3	1.1	.01	.09	.09	.75
1.0	0.8	1.2	1.5	.02	.11	.14	.79
SW-leaning, $\lambda_{\Delta r} = 1$							
0.0	0.9	0.5	0.4	.00	.15	.19	.47
0.5	0.9	0.5	0.5	.00	.10	.10	.74
1.0	0.8	0.5	0.7	.01	.12	.18	.71

Notes: Optimal policy rule parameters and IIP relative to first-best simple rule for each model (in percentage points). Three ambiguity-averse solutions are shown: one using flat model priors, one giving extra weight only to the AW, and one giving extra weight only to the SW model. Parameters for the ambiguity-averse solution are $p_m = 0.25$ for the former, and $p_{AW} \approx 1$ and $p_{SW} \approx 1$, respectively, for the latter. The policymaker gives a weight of $e = 0.5$ to the worst-case scenario. Shown are the case $\lambda_{\Delta r} = 0.1$ and $\lambda_{\Delta r} = 1$ for loss function (3).

TABLE E.8. Ambiguity-averse implied priors for different weights $\lambda_{\Delta r}$.

	Flat priors				AW leaning				SW leaning			
λ_y	CW-F	CW-T	SW	AW	CW-F	CW-T	SW	AW	CW-F	CW-T	SW	AW
$\lambda_{\Delta r} = .1$												
0.0	0.625	0.125	0.125	0.125	0.500	0.0	0.0	0.5	0.480	0.0	0.5	0.020
0.5	0.625	0.125	0.125	0.125	0.500	0.0	0.0	0.5	0.500	0.0	0.5	0.000
1.0	0.625	0.125	0.125	0.125	0.500	0.0	0.0	0.5	0.500	0.0	0.5	0.000
$\lambda_{\Delta r} = 1$												
0.0	0.625	0.125	0.125	0.125	0.500	0.0	0.0	0.5	0.500	0.0	0.5	0.000
0.5	0.625	0.125	0.125	0.125	0.500	0.0	0.0	0.5	0.475	0.0	0.5	0.025
1.0	0.625	0.125	0.125	0.125	0.500	0.0	0.0	0.5	0.419	0.0	0.5	0.081

Notes: Bayesian priors backing the ambiguity-averse solution. Two ambiguity-averse solutions are shown: one giving extra weight to the AW and one giving extra weight to the SW model. Parameters for the ambiguity averse solution are $p_{SW} \approx 1$ and $p_{AW} \approx 1$, respectively. We set $e = 0.5$, and $p_m \approx 0$ for all other models. Shown are the case $\lambda_{\Delta r} = 0.1$ and $\lambda_{\Delta r} = 1$ for loss function (3).

References

Adalid, Ramón, Guenter Coenen, Peter McAdam, and Stefano Siviero (2005). "The Performance and Robustness of Interest-Rate Rules in Models of the Euro Area." *International Journal of Central Banking*, 1(1), 95–132.

Adam, Klaus (2004). "On the Relation Between Robust and Bayesian Decision Making." *Journal of Economic Dynamics and Control*, 28(10), 2105–2117.

Angelini, Paolo, Paolo Del Giovane, Stefano Siviero, and Daniele Terlizzese (2004). "Monetary Policy Rules for the Euro Area: What Role for National Information?" Temi di discussione [Economic working papers] No. 457, Bank of Italy.

Becker, R., B. Dwolatzky, E. Karakitsos, and Berc Rustem (1986). "The Simultaneous Use of Rival Models in Policy Optimisation." *The Economic Journal*, 96, 89–101.

Brayton, Flint, A. T. Levin, Ralph W. Tryon, and John C. Williams (1997). "The Evolution of Macro Models at the Federal Reserve Board." *Carnegie-Rochester Conference Series on Public Policy*, 42, 115–167.

Brayton, Robert K., S. W. Director, Gary D. Hachtel, and Luis Manuel Vidigal (1979). "A New Algorithm for Statistical Circuit Design Based on Quasi-Newton Methods and Function Splitting." *IEEE Transactions on Circuits and Systems*, 26(9), 784–794.

Brock, William A., Steven N. Durlauf, and Kenneth D. West (2003). "Policy Evaluation in Uncertain Economic Environments." *Brookings Papers on Economic Activity*, 1, 235–322.

Coenen, Guenter (2007). "Inflation Persistence and Robust Monetary Policy Design." *Journal of Economic Dynamics and Control*, 31(1), 111–140.

Coenen, Guenter, and Volker Wieland (2005). "A Small Estimated Euro Area Model with Rational Expectations and Nominal Rigidities." *European Economic Review*, 49, 1081–1104.

Cote, Denise, John Kuszczak, Jean-Paul Lam, Ying Liu, and Pierre St-Amant (2004). "The Performance and Robustness of Simple Monetary Policy Rules in Models of the Canadian Economy." *Canadian Journal of Economics*, 37(4), 978–998.

Dieppe, Alistair, Keith Kuester, and Peter McAdam (2004). "Optimal Monetary Policy Rules for the Euro Area: An Analysis Using the Area Wide Model." ECB Working Paper Series No. 360. European Central Bank, Frankfurt.

Dieppe, Alistair, Keith Kuester, and Peter McAdam (2005). "Optimal Monetary Policy Rules for the Euro Area: An Analysis Using the Area Wide Model." *Journal of Common Market Studies*, 43(3), 507–5372.

Epstein, Lawrence, and Tan Wang (1994). "Intertemporal Asset Pricing under Knightian Uncertainty." *Econometrica*, 62(2), 283–322.

Erceg, C. J., and A. T. Levin (2003). "Imperfect Credibility and Inflation Persistence." *Journal of Monetary Economics*, 50, 915–944.

Fagan, Gabriel, Jerome Henry, and Ricardo Mestre (2001). "An Area-Wide Model (AWM) for the Euro Area." ECB Working Paper Series No. 42, European Central Bank, Frankfurt.

Fagan, Gabriel, Jerome Henry, and Ricardo Mestre (2005). "An Area-Wide Model for the Euro Area." *Economic Modelling*, 22, 39–59.

Fuhrer, Jeff, and George Moore (1995). "Inflation Persistence." *Quarterly Journal of Economics*, 110(1), 127–159.

Fuhrer, Jeffrey C. (2000). "Habit Formation in Consumption and Its Implications for Monetary-Policy Models." *American Economic Review*, 90(3), 367–390.

Gerdesmeier, Dieter, and Barbara Roffia (2004). "Empirical Estimates of Reaction Functions for the Euro Area." *Swiss Journal of Economics and Statistics*, 140, 37–66.

Giannoni, Marc P. (2002). "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model." *Macroeconomic Dynamics*, 6(1), 111–144.

Giannoni, Marc P. (2007). "Robust Optimal Monetary Policy in a Forward-Looking Model with Parameter and Shock Uncertainty." *Journal of Applied Econometrics*, 22(1), 179–213.

Hansen, Lars P., and Thomas J. Sargent (2002). "Robust Control and Model Uncertainty in Economics." Working paper, Nyu.

Kuester, Keith, and Volker Wieland (2005). "Insurance Policies for Monetary Policy in the Euro Area." European Central Bank Working Paper No. 480.

Leeper, Eric, and Tao Zha (2003). "Modest Policy Interventions." *Journal of Monetary Economics*, 50(8), 1673–1700.

Levin, Andrew, Alexei Onatski, John Williams, and Noah Williams (2006). "Monetary Policy under Uncertainty in Micro-Founded Macroeconometric Models." In *NBER Macroeconomics Annual 2005*, edited by M. Gertler, and Kenneth Rogoff. MIT Press.

Levin, Andrew, Volker Wieland, and John Williams (1999). "Robustness of Simple Monetary Policy Rules under Model Uncertainty." In *Monetary Policy Rules*, edited by J. Taylor. University of Chicago Press.

Levin, Andrew, Volker Wieland, and John Williams (2003). "The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty." *American Economic Review*, 93(3), 622–645.

Levin, Andrew, and John Williams (2003). "Robust Monetary Policy with Competing Reference Models." *Journal of Monetary Economics*, 50(5), 945–975.

McCallum, Bennett T. (1988). "Robustness Properties of a Rule for Monetary Policy." *Carnegie-Rochester Conference Series on Public Policy*, 29(autumn), 173–203.

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

Onatski, Alexei, and Noah Williams (2003). "Modeling Model Uncertainty." *Journal of the European Economic Association*, 1(5), 1087–1122.

Orphanides, Athanasios, and Volker Wieland (1998). "Price Stability and Monetary Policy Effectiveness when Nominal Interest Rates are Bounded at Zero." Finance and Economics Discussion Paper No. 98-35. Board of Governors of the Federal Reserve System.

Rudebusch, Glenn D., and Lars E. O. Svensson (1999). "Policy Rules for Inflation Targeting." In *Monetary Policy Rules*, edited by J. B Taylor. University of Chicago Press.

Sack, Brian (2000). "Does the Fed Act Gradually? A VAR Analysis." *Journal of Monetary Economics*, 1, 229–256.

Sargent, Thomas (1999). *The Conquest of American Inflation*. Princeton University Press.

Sims, Christopher A. (2001). "Pitfalls of a Minimax Approach to Model Uncertainty." *American Economic Review*, 91(2), 51–54.

Smets, Frank, and Raf Wouters (2003). "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area." *Journal of European Economic Association*, 1(5), 1123–1175.

Taylor, John B. (1980). "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy*, 88(1), 1–23.

Taylor, John B. (1993). *Macroeconomic Policy in the World Economy: From Econometric Design to Practical Operation*. Norton.

Taylor, John B. (1999). "Introduction." In *Monetary Policy Rules*, edited by J. Taylor. University of Chicago Press.

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(6-7), 911–949.

Tetlow, Robert, and Peter von zur Muehlen (2004). "Avoiding Nash Inflation: Bayesian and Robust Responses to Model Uncertainty." *Review of Economic Dynamics*, 7(4), 869–899.

von zur Muehlen, Peter (1982). "Activist versus Nonactivist Monetary Policy: Optimal Rules under Extreme Uncertainty." Finance and Economics Discussion Series Paper No. 2001-02. Federal Reserve System.

Wieland, Volker (2000a). "Learning by Doing and the Value of Optimal Experimentation." *Journal of Economic Dynamics and Control*, 24, 501–534.

Wieland, Volker (2000b). "Monetary Policy, Parameter Uncertainty and Optimal Learning."
    Journal of Monetary Economics, 46, 199–228.
Wieland, Volker (2006). "Monetary Policy and Uncertainty about the Natural Unemploy-
    ment Rate: Brainard-Style Conservatism versus Experimental Activism." *Advances in
    Macroeconomics*, 6(1), Article 1.
Woodford, Michael (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*.
    Princeton University Press.
Zakovic, Stanislav, Berc Rustem, and Volker Wieland (2007). "Stochastic Optimisation and
    Worst Case Analysis in Monetary Policy Design." *Computational Economics*, 30, 329–347.