Financial Stability and Optimal Interest Rate Policy*

Andrea Ajello, Thomas Laubach, David López-Salido, and Taisuke Nakata
Federal Reserve Board

We study optimal interest rate policy in a New Keynesian framework in which the model economy can experience financial crises and the probability of a crisis depends on credit conditions. We find that the optimal response of the short-term interest rate to credit conditions is (very) small in the model calibrated to match the historical relationship between credit conditions, output, inflation, and likelihood of financial crises. Given the imprecise estimates of key parameters, we also study optimal policy under parameter uncertainty. We find that Bayesian and robust central banks will respond more aggressively to financial instability when the probability and severity of financial crises are uncertain.

JEL Codes: E43, E52, E58, G01.

*We would like to thank Stephen Cecchetti, Frank Smets, Lars Svensson, Javier Garcia-Cicco, Rafael Repullo, and two referees for their comments and suggestions as well as seminar and conference participants at the Board of Governors, the Bank of Japan, the University of Tokyo, the Federal Reserve Bank of San Francisco Annual Conference "The New Normal for Monetary Policy," the 2015 Midwest Macro Meetings at the Federal Reserve Bank of St. Louis, the CEF 2015 meetings, the Georgetown Center for Economic Research Biennial Conference, the ECB, the European Economic Association Meetings, the Swiss National Bank Annual Research Conference, the Dynare Conference, and the Central Bank of Chile Financial Policy Workshop "Macroprudential Policies: Experiences and Challenges." The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System. Miguel Acosta, Timothy Hills, Kathryn Holston, and Jonathan Yu provided excellent research assistance. Author e-mails: andrea.ajello@frb.gov (corresponding author), thomas.laubach@frb.gov, david.lopez-salido@frb.gov, and taisuke.nakata@frb.gov.

> "Monetary policy faces significant limitations as a tool to promote financial stability: Its effects on financial vulnerabilities, such as excessive leverage and maturity transformation, are not well understood . . . in addition, efforts to promote financial stability through adjustments in interest rates would increase the volatility of inflation and employment."
>
> Janet Yellen (2014), Chair of the Board of Governors of the Federal Reserve System

> "There are likely to be potential gains from a more financial stability-oriented monetary policy; and any such policy, if it is to produce gains, would need to take financial developments into account systematically, in both good and bad times."
>
> Claudio Borio (2016), Head of the Monetary and Economic Department, Bank for International Settlements

1. Introduction

The debate on whether central banks should "lean against the wind" of financial imbalances has recently received renewed interest from academics and policymakers. While the consensus before the crisis advocated that central banks should not raise interest rates to contrast exuberance on financial markets (Bernanke and Gertler 1999), the severity of the Great Recession and the constraints imposed on the stabilizing role of monetary policy have spurred questions about what instruments, if any, central banks could deploy to prevent future financial disruptions. A question of particular interest is whether interest rate policy should respond systematically to changes in financial conditions. As suggested by the quotes above, the debate has yet to reach an encompassing conclusion, and the policy prescriptions that are advocated by the different participants rely inherently on the quantitative evaluation of costs and benefits of "leaning against the wind."

We contribute to this debate by studying optimal interest rate policy in a simple New Keynesian model with a rare crisis event, calibrated to portray salient features of the U.S. economic system. In our model, the economy is subject to two-state shocks and goes through normal times and occasional financial crises; the probability at which a crisis can occur is time varying and depends on aggregate financial conditions, as in Woodford (2012b).[1] In this situation, when confronted with the possibility of a financial crisis, the policymaker faces a new intertemporal tradeoff between stabilizing real activity and inflation in normal times and mitigating the possibility of a future financial crisis. The adjustment to the policy rate that is optimal, compared with a setting without financial stability concerns, depends on four sets of parameters: the costs of suffering a financial crisis (and thus the benefits of avoiding this fate), the marginal effect of the policy rate on both the probability of a crisis and its severity, and the output and inflation losses arising in normal times from a policy response that averts future financial stability risks.

To make the exploration empirically relevant, we calibrate the relationship between the likelihood of financial crises and credit conditions to the U.S. experience, borrowing and adapting recent evidence on the cross-country historical data of Schularick and Taylor (2012). Our theoretical analysis shows that the optimal adjustment in the policy rate that arises from financial stability risks is (very) small, less than 10 basis points, when the model is calibrated to match the (estimated) historical relationships between credit conditions, output, and inflation as well as the likelihood and severity of a financial crisis.[2]

Nevertheless, reflecting the infrequent nature of crises episodes, the evidence linking credit conditions to financial crises and the effectiveness of interest rate policy in preventing or reducing the impact of crises are subject to substantial uncertainty. More precisely, we find that a number of key parameters that control the transmission channels of monetary policy appear to be imprecisely estimated in

[1]Throughout the analysis we assume that the only policy tool available to the central bank is the short-term interest rate. Equivalently, we assume that macroprudential policies, if available, have been employed and have exhausted their role in fostering financial stability.

[2]Svensson (2014) uses the Riksbank DSGE model to perform a similar analysis and argues that the cost of "leaning against the wind" (LAW) interest rate policies in terms of current real activity far exceeded the benefits of financial stabilization in the recent Swedish experience. Clouse (2013) instead finds that policymakers may seek to reduce the variance of output by scaling back the level of accommodation in a stylized two-period model that is similar to ours in which loose interest rate policy today can generate sizable future losses in output.

the data. For this reason, we first consider the sensitivity of the optimal policy to alternative parameter values, and then analyze how the optimal policy is affected if the policymaker is confronted with uncertainty about some of the parameters of the model.

Under alternative plausible assumptions regarding the value of key parameters, the optimal policy can call for larger adjustments to the policy rate than in a situation without financial stability concerns. For example, if we assume that the adjustment in the policy rate is two standard deviations more effective in reducing the crisis probability than in the baseline specification, the optimal adjustment in the policy rate can be as large as 50 basis points. Moreover, if we assume that the effects of a financial crisis on inflation and the output gap are comparable in magnitude to those observed during the Great Depression—as opposed to the Great Recession scenario used as our baseline—the optimal policy will call for a riskless short-term interest rate that can be around 75 basis points higher than what would be optimal in the absence of financial stability concerns.

We then consider how the optimal policy is affected if the policymaker is uncertain about three sets of parameters. First, we look at uncertainty regarding the relationship between the crisis probability and aggregate credit conditions. Second, we consider uncertainty regarding the severity of the crisis. Finally, we look at the effects of uncertainty regarding the extent to which changes in the policy rate affect today's inflation and output.[3]

We frame our optimal-policy problem under uncertainty following both of these approaches and consider two types of policymakers. The first type is a Bayesian central bank that aims to maximize the expected welfare of the economy for a given prior distribution of the parameters of the model. This approach originated from the seminal work of Brainard (1967), who showed that in a static framework in

[3]Policymakers called to make interest rate decisions in the presence of financial stability risk may face different types of uncertainty. As discussed in our empirical analysis, the parameters governing the relationship between the probability of a financial crisis and aggregate credit conditions are estimated with wide confidence intervals, reflecting the infrequent nature of crises in history. Similarly, recent studies have documented a large dispersion in the severity of crisis episodes across countries and time (see, for example, Reinhart and Rogoff 2009, 2014; Jordà, Schularick, and Taylor 2013; and Romer and Romer 2014). Moreover, the structure of the economy and of the monetary policy transmission channels can change over time.

which there is uncertainty about the transmission of monetary policy, it is optimal for the central bank to respond less aggressively to shocks than if the parameter were known with certainty (a result known as Brainard's *attenuation* principle).[4] We follow more recent work by Brock, Durlauf, and West (2003), Svensson and Williams (2007), and Cogley et al. (2011) that incorporate Bayesian uncertainty into a linear quadratic framework and characterize optimal policy.[5]

The second type of policymaker is a robust central bank that aims at protecting against worst-case scenarios. To do so, the central bank minimizes the maximum loss over a set of parameters, including those with only a low probability of being realized. Thus, an optimal policy is robust in the sense that it performs best in the worst-case configuration around the (single) reference model, providing a form of insurance against the least favorable scenarios. As in the case of the Bayesian approach to model uncertainty, Brainard's principle can be overturned in this context: the robust policymaker will achieve higher welfare by responding more strongly in advance to forestall the development of future unfavorable outcomes (see Tetlow and von zur Muehlen 2001, Giannoni 2002, and Onatski and Stock 2002). That is, in this case, optimal policy might result in a more aggressive response than in the certainty-equivalent case.[6]

[4]A first step in the implementation of a Bayesian approach consists of building a crisp set of alternative elements of the transmission mechanism or, alternatively, how different economic theories disagree over fundamental aspects of the economy. Then, modeling uncertainty requires the specification of a prior distribution over the space of models, and then propagates this uncertainty to the analysis of monetary policy problem by integrating monetary policy and models out from the posterior distribution. This is what is called Bayesian model averaging (e.g., Brock, Durlauf, and West 2003).

[5]This approach typically implies that the optimal policy exhibits some form of attenuation, as in Brainard (1967), compared with the case of no uncertainty, although this result has some exceptions. While Brainard's analysis is conducted in a static framework, in the dynamic models of Giannoni (2002) and Söderstrom (2002), for example, uncertainty about the persistence of inflation implies that it is optimal for the central bank to respond more aggressively to shocks than if the parameter were known with certainty. In our framework these intertemporal dimensions will arise endogenously from the effects of future likely crises on current outcomes.

[6]To our knowledge, none of the existing studies have considered the nonlinearity coming from the presence of financial crises on the (optimal) nominal

As discussed above, we examine three forms of uncertainty faced by the Bayesian and robust policymakers. First, our main finding is that uncertainty about the effectiveness of the interest rate policy in reducing the probability of a crisis leads both the Bayesian and the robust policymakers to increase the policy rate by more than in the absence of uncertainty, so that the attenuation principle of Brainard (1967) fails. In the model with a Bayesian policymaker, the key to this result is related to the non-linear properties of our crisis probability function: in our model the economy's likelihood of facing a financial crisis is increasing and convex in aggregate credit conditions, and a higher sensitivity to aggregate credit conditions can make the probability of a crisis increase more rapidly for a given change in credit conditions. Uncertainty around this sensitivity parameter tends to make the *expected* probability of a crisis higher and more responsive to credit conditions and hence to the central bank's interest rate policy. In this context, given the higher marginal benefit associated with a tighter policy in lowering the expected future crisis probability (by reducing the availability of credit), the policymaker optimally decides to set the nominal rate higher than in the absence thereof. The same policy prescription follows from a robust perspective since the hypothetical evil agent inside the head of the (robust) policymaker can maximize the welfare loss by increasing the sensitivity of the crisis probability to credit conditions.

Second, in the face of uncertainty about the severity of the crisis, measured in terms of output gap and inflation variability, the same result holds: This type of uncertainty leads both the Bayesian and the robust policymakers to set the policy rate higher than otherwise. In the model with a Bayesian policymaker, this result is driven by the non-linearity of his/her quadratic utility function. In the model with the robust policymaker, this result is more general and does not hinge on the specification of a quadratic loss function.

Third, in the face of uncertainty about the response of today's inflation and output to the policy rate—the same uncertainty considered in Brainard (1967)—the attenuation principle holds for both

risk-free interest rate. In the online appendixes we sketch some of the potential implications for robust optimal policy when an additional non-linearity—the effective lower bound on the short-term interest rate—is introduced in the model.

types of policymakers: the presence of uncertainty leads policymakers to adjust the policy rate by less than otherwise, to avoid increasing the aggregate volatility of output and inflation.

Our optimal policy results under uncertainty differ from Bernanke and Gertler (1999) in a number of dimensions. First, while our analysis is centered around the role of credit growth as a source of financial instability and a predictor of financial crises, their argument considers financial imbalances caused by excess volatility in asset prices. Second, Bernanke and Gertler (1999) highlight how uncertainty in the effects of interest rate policy on financial conditions reduces the benefits of LAW policy. In their analysis the sizable costs inflicted to output and inflation dominate the uncertain benefits of LAW. One key element in Bernanke and Gertler's analysis is that the policymaker can also implement effective countercyclical interest rate policy ex post to stabilize output and inflation, after a financial downturn. In our framework, however, uncertainty about the effectiveness of interest rate policy in reducing financial instability is paired with severe limitations on the conduct of stabilization policy after a financial crisis has materialized. These limitations are exemplified in our analysis by the presence of a large (and inevitable) cost paid by the economy in terms of negative output and inflation gaps once the crisis materializes. Such costs can be rationalized, for example, in model economies in which the policy rate hits its effective lower bound that limits the scope of monetary accommodation. As a result of this assumption, the cost of raising interest rates ex ante to prevent financial imbalances can be of second order, when compared with the expected welfare gain of reducing the likelihood of a large economic downturn. Under uncertainty, such gains can appear sizable to a Bayesian and to a robust policymaker and we find that under plausible model calibrations, a moderate degree of LAW policy is in fact optimal.

Our paper is close in spirit to work by Svensson (2016), who finds that costs of leaning-against-the-wind policy in terms of lower output and inflation outweigh financial stability benefits. His approach differs from ours, as he develops a method to compare costs and benefits of such policy to existing quantitative macro models, while our simple framework allows us to document in detail the role of crucial model assumptions and parameters (including, for example, uncertainty) in affecting the optimal degree of leaning against the wind.

Laseen, Pescatori, and Turunen (2015) reach conclusions similar to Svensson's and to our baseline model calibration, adopting a different analytical and empirical strategy that controls for the effect of interest rate policy on time-varying systemic risk in a non-linear New Keynesian DSGE framework. Galí (2014) also offers a theoretical argument against LAW interest rate policy in an overlapping-generations (OLG) New Keynesian model with rational asset bubbles. He shows that the bubble component of asset prices has to grow at the real rate of interest in equilibrium, so that LAW interest rate policy can in fact promote (and not tame) the increase in asset prices.

On the opposite side of our baseline results and of Svensson's arguments, Adrian and Liang (2013) reach conclusions that resonate with Stein (2013)'s claim that interest rate policy has an advantage over macroprudential policy, as it can "get in all the cracks" of the financial system. They highlight that interest rate policy can influence financial vulnerabilities more uniformly than macroprudential policy, although they warn that the efficacy of LAW as a tool depends on the costs of tighter policy on activity and inflation. They show that if tighter interest rate policy can reduce the degree of risk-taking of the financial sector, the cost–benefit calculation can be significantly altered in favor of LAW policy. Gourio, Kashyap, and Sim (2016) reach a similar result by building a New Keynesian model in which the central bank, while not solving for the optimal interest rate policy, would optimally choose to lean against the wind by adopting a Taylor rule that targets financial imbalances, in order to avoid large losses in output.

The rest of the paper is organized as follows. Section 2 describes the model and discusses the parameterization used in our simulation exercise. Section 3 presents the results based on the baseline and some alternative calibrations. Section 4 formulates the problem of both Bayesian and robust policymakers and presents the results on how uncertainty about the parameters affects our previous prescriptions regarding optimal interest rate policy in the presence of financial stability concerns. A final section concludes. Extra material—including modeling, econometric analyses, and an extension of the analysis that accounts for the presence of the zero lower bound constraint—is presented in the online appendixes, available at http://www.ijcb.org.

2. Financial Crises in a Two-Period New Keynesian Model

The stylized framework is a standard New Keynesian sticky-price model augmented with an endogenous financial crisis event. We use a two-period version of the model to build intuition on the main ingredients that shape the tradeoff faced by the central bank in an economy with possible financial instability. The occurrence of financial crises follows a Markov process, with its transition probability governed by the evolution of aggregate financial conditions. Based on recent empirical work discussed below, we assume that periods of rapid credit growth raise the probability of transitioning from the non-crisis to the crisis state. In this sense, this basic setup closely resembles Woodford (2012a) reducing the infinite horizon of that model to a two-period framework to better isolate the role that model assumptions play in shaping the tradeoff between macroeconomic and financial stabilization.

2.1 Economic Structure and Policy Objectives

The following three equations describe the dynamics of the output gap y, inflation π, and credit conditions L:

$$y_1 = E_1^{ps} y_2 - \sigma\left[i_1 - E_1^{ps}\pi_2\right] \tag{1}$$

$$\pi_1 = \kappa y_1 + E_1^{ps}\pi_2 \tag{2}$$

$$L_1 = \rho_L L_0 + \phi_i(i_1 + i^*) + \phi_y y_1 + \phi_\pi(\pi_1 + \pi^*) + \phi_0. \tag{3}$$

From equation (1), the output gap in period 1 (y_1) depends on the expected output gap in period $t = 2$ ($E_1^{ps} y_2$), and on deviations of the period-1 real rate, defined as $[i_1 - E_1^{ps}\pi_2]$, from its long-run equilibrium level (the relation between the private sector's expectations operator E^{ps} and rational expectations will be discussed below). From equation (2), inflation in period $t = 1$ depends on the current output gap and expected future inflation; while from equation (3), financial conditions in period $t = 1$ depend on their value in period $t = 0$, on the output gap, and on the nominal interest rate and inflation. In particular, $(\pi + \pi^*)$ denotes the rate of inflation (defined as inflation gap plus policymaker's target, π^*); and $(i + i^*)$ is the riskless short-term nominal interest rate (the policy rate, defined as the gap, i plus the long-run equilibrium rate, i^*). L is a proxy for aggregate credit conditions in the model. We choose L to be the five-year cumulative growth rate of real bank loans, expressed in decimal percentages (e.g., 0.2 corresponds to a 20 percent cumulative credit growth over the past five years). We chose the growth rate of real bank loans as a measure of financial stability, based on the findings of Schularick and Taylor (2012), who compare a battery of macro and financial variables as possible early predictors of financial crises in a series of panel logit regressions over fourteen developed countries for over 140 years. Their work, described in more detail in section 2.3.1, suggests that five lagged annual growth rates of real bank loans show the highest predictive power in forecasting the occurrence of financial crises. The also find that other variables, such as measures of real activity, inflation, aggregate loans over GDP, or stock price gains, have little explanatory power when added to their baseline regression that includes lagged real bank loan growth, suggesting that financial crises are in fact "credit booms gone bust."[7]

To keep the analysis focused, we abstract from any direct effect of credit conditions on the output gap and inflation.[8] Instead, credit conditions only affect the probability γ_1 that controls the likelihood of the transition to a crisis state in period $t = 2$. Credit conditions, L_1, affect γ_1 according to the logistic function:

$$\gamma_1 = \frac{\exp(h_0 + h_1 L_1)}{1 + \exp(h_0 + h_1 L_1)}, \tag{4}$$

where h_0 pins down the intercept probability when $L_1 = 0$, and h_1 is the sensitivity of the crisis probability to credit conditions.

Let $\pi_{2,c}$ and $y_{2,c}$ denote inflation and the output gap in the crisis state, while $\pi_{2,nc}$ and $y_{2,nc}$ denote their non-crisis-state values.

[7]In section 2.3.1 we show that we can easily map Schularick and Taylor (2012)'s evidence for the United States to the structure of our model. We believe this is also an important element that favors the adoption of the growth rate of real bank loans as a preferred proxy for financial stability.

[8]Online appendix E.2 discusses an extension of our model in which credit conditions have a positive effect on the output gap.

Then inflation and the output gap outcomes in period $t = 2$ will take values

$$(y_2, \pi_2) = \begin{cases} (y_{2,nc}, \pi_{2,nc}), & \text{with probability } = 1 - \gamma_1 \\ (y_{2,c}, \pi_{2,c}), & \text{with probability } = \gamma_1 \end{cases}$$

with $\pi_{2,c} < \pi_{2,nc} = 0$ and $y_{2,c} < y_{2,nc} = 0$.

Throughout the analysis we assume that the private sector treats γ_1 as fixed and negligible in size and not as a function of L_1, implying that in this regard expectations are *optimistic* and hence not rational. We assume that private agents perceive the probability of the crisis to be different from γ_1 and to be constant and potentially negligible, i.e., a *tail event*. Formally, we assume the following rule regarding private-sector expectations:

$$E_1^{ps} y_2 = (1 - \epsilon) y_{2,nc} + \epsilon y_{2,c} \tag{5}$$

$$E_1^{ps} \pi_2 = (1 - \epsilon) \pi_{2,nc} + \epsilon \pi_{2,c}, \tag{6}$$

where ϵ is arbitrarily small and does not depend on aggregate credit conditions.[9]

We find evidence in support of this assumption in data from the Survey of Professional Forecasters (SPF) on expectations of future GDP growth and inflation. Appendix A shows that over the course of 2007 and 2008 the median forecaster in the SPF assigned a probability close to 0 percent to the event that average real GDP and CPI inflation could fall in 2008. Similarly, the median SPF forecaster reported probabilities below 2 percent when asked to forecast the likelihood of negative growth for average real GDP in 2009, at least until the collapse of Lehman Brothers in 2008:Q3. Only at that point—between 2008:Q3 and 2008:Q4—as more information on the severity of the financial crisis became available, did the median forecasted probability of negative growth and the median forecasted probability of CPI deflation in 2009 increase from 2 percent to 55 percent and from 0 percent to 10 percent respectively (see figure A1 in online appendix A). We interpret these findings as evidence that

[9]The knife-edge assumption of a small and constant perceived crisis probability ϵ can be relaxed without altering our results. In particular, the preceived probability ϵ can be allowed to vary with credit conditions, as long as the parameter that governs the sensitivity to L_1 is small.

expectations of financial market participants on the likelihood of a financial crisis and a prolonged downturn adjusted with a lag to the unfolding of the events over the course of the Great Recession, rather than responding preemptively—for example, to the accumulation of financial imbalances over the course of the economic expansion of the 2000s.

For comparison, in section 3.3 we also present the model solution under rational expectations. In this case the private sector understands that the likelihood of a financial crisis depends on the evolution of credit conditions, as in equation (4), so that

$$E_1^{ps} y_2 = E_1^{re} y_2 = (1 - \gamma_1) y_{2,nc} + \gamma_1 y_{2,c} \tag{7}$$

$$E_1^{ps} \pi_2 = E_1^{re} \pi_2 = (1 - \gamma_1) \pi_{2,nc} + \gamma_1 \pi_{2,c}. \tag{8}$$

Under rational expectations, increasing credit growth increases the likelihood of a financial and reduces the private sector's expectation of future output and inflation, relative to the case of optimistic expectations. In turn, lower expectations lead to lower realizations of output and inflation today (from equations (1) and (2)). In this framework, the phase of buildup of financial instability is characterized by negative output and inflation gaps that the central bank might be tempted to fight by means of accommodative interest rate policy.

2.2 The Policymaker's Problem

Let W denote the policymaker's loss function. The policy problem consists of choosing in period $t = 1$ the policy rate given initial credit conditions, L_0, the only endogenous state variable of the model. Formally, the problem of the central bank at time $t = 1$ is given by

$$W_1 = \min_{i_1} \; u(y_1, \pi_1) + \beta E_1[W_2] \tag{9}$$

subject to the previous private-sector equilibrium conditions (1) to (3) and where

$$u(y_1, \pi_1) = \frac{1}{2}(\lambda y_1^2 + \pi_1^2) \tag{10}$$

and $W_{2,c}$ and $W_{2,nc}$ denote the welfare losses in the crisis and non-crisis states, respectively. $W_{2,c}$ is related to inflation and the output gap in the crisis state by

$$W_{2,c} = \frac{u(y_{2,c}, \pi_{2,c})}{1 - \beta\mu}, \tag{11}$$

where μ is a parameter calibrated to capture the effects of the duration of financial crises on output and inflation, expressed in utility terms. This scaling-up is aimed at ensuring that the costs of financial crises are appropriately captured in our two-period framework. The expected welfare loss at time $t = 2$ is then given by

$$E_1[W_2] = (1 - \gamma_1)W_{2,nc} + \gamma_1 W_{2,c}. \tag{12}$$

We normalize the welfare loss in the non-crisis state to zero, $W_{2,nc} = 0$. One potential shortcoming of our two-period framework is that it may not take full account of the effects of the policy rate setting on forward-looking measures of social welfare (that discount output and inflation gaps that occur many periods into the future) as well as the possibly long-lasting effects of the policy rate on financial stability and on the crisis probability in the long run. Our two-period framework effectively maps into an infinite-horizon model in which the central bank sets the nominal interest rate to minimize the sum of current and future expected welfare losses, knowing that (i) its decision will only affect current output and inflation gaps and current financial conditions; (ii) if a financial crisis does *not* materialize, the economy will be perfectly stabilized starting from period 2 onward (output gap and inflation gaps are assumed to be equal to zero in every period); and (iii) if a financial crisis *does* start at time 2, output gap and inflation gap are assumed to be large and negative for a number of periods (pinned down by the parameter μ) and that the economy will go back to zero output gap and zero inflation once the crisis has ended.

2.3 Parameter Values

Table 1 shows the baseline parameter values. The values for the parameters pertaining to the standard New Keynesian model are chosen to be consistent with many studies in the literature, such as Woodford (2003). The annual inflation target, π^*, is assumed to be 2 percent, and hence our choice of the long-run equilibrium policy rate, i^*, of 4 percent implies an equilibrium real short-term rate of 2 percent in a model without financial instability. The weight $\lambda = \frac{1}{16}$

Table 1. Baseline Parameter Values

Param.	Description	Value	Note
Standard Parameters			
β	Discount Factor	0.995	Standard
σ	Interest Rate Sensitivity of Output	1.0	Standard
κ	Slope of the Phillips Curve	0.024	Standard
λ	Weight on Output Stabilization	1/16	Equal Weights on y and the Annualized π
i^*	Long-Run Natural Rate of Interest	0.01	4% (Annualized)
π^*	Long-Run Inflation Target	0.005	2% (Annualized)
Parameters for the Equation Governing the Crisis Probability			
h_0	Constant Term	–3.396	
h_1	Coefficient on L	1.88	
Parameters for the Equation Governing the Financial Conditions			
ρ_L	Coefficient on the Lagged L	19/20	
ϕ_0	Intercept	$(1-\rho_L)*0.2$	
ϕ_y	Coefficient on Output Gap	0.18	See Appendix B
ϕ_π	Coefficient on Inflation Gap	–0.57	(0.43 – 1); See Appendix B
Parameters Related to the Second Period			
$y_{2,nc}$	Output Gap in the Non-crisis State	0	
$\pi_{2,nc}$	Inflation Gap in the Non-crisis State	0	
$W_{2,nc}$	Loss in the Non-Crisis State	0	
$y_{2,c}$	Output Gap in the Crisis State	–0.1	"Great Recession"
$\pi_{2,c}$	Inflation Gap in the Crisis State	–0.02/4	"Great Recession"
μ	Persistence of the Crisis State	7/8	
$W_{2,c}$	Loss in the Crisis State	$\frac{u(y_{2,c},\pi_{2,c})}{1-\beta\mu}$	
Auxiliary Parameters			
ϵ	Perceived Crisis Probability	0.05/100	Arbitrarily Small

in the central bank's period loss function implies equal concern for annualized inflation gaps and output gaps.[10] We do not attempt to derive this objective from a representative household's utility, but are instead interested in the question of how a policymaker who wants to minimize fluctuations in the output gap and inflation from their targets (reminiscent of the Federal Reserve's traditional dual mandate) would want to alter the macroeconomic stabilization in response to financial stability risks.

In the remainder of this section we will discuss the calibration of the probability of a financial crisis, γ_1, and the evolution of the credit conditions index, L. These are the parameters that influence our results most strongly and that may be considered more controversial in the debate about the appropriate response of interest rate policy to financial stability concerns. Finally, we will also discuss the choice of parameters that affect the severity of the crisis, a key determinant of the welfare losses associated with a crisis state.[11]

2.3.1 A Simple Model of Crisis Probability and Credit Conditions: The U.S. Experience

The ability to predict events such as currency, fiscal, and financial crises by means of econometric models is hindered by the rarity of such episodes in the history of both advanced and emerging economies. Schularick and Taylor (2012) make a thorough attempt to understand the role of bank lending in the buildup to financial

[10]The weight of 1/16 comes from 1/4 squared. Since the model is formulated in quarterly frequency (and an annualized inflation rate is four times a quarterly inflation rate), we need this adjustment factor to make the weight on the volatility of *annualized* inflation gaps and the weight on the volatility of output gaps equal. In online appendix E.4 we consider an alternative value for λ that is consistent with the one obtained under a second-order approximation of welfare, as in Woodford (2003).

[11]There is evidence that credit cycles evolve over longer time horizons than business cycles (Borio 2012; Aikman, Haldane, and Nelson 2015). It is also plausible to assume that tighter interest rate policy may positively affect financial stability only if sustained over time. For these reasons we also provide an annual calibration of our model in online appendix E.1. We find that maintaining a tighter monetary policy stance for one year in normal times improves financial stability only modestly, while inducing higher welfare losses measured in terms of output and inflation gaps than in the quarterly calibration. We also find that when financial crisis are less likely, the central bank might find it optimal to lean *with* the wind, as suggested by Svensson (2016).

crises, using discrete choice models on a panel of fourteen countries over 138 years (1870–2008). The paper characterizes empirical regularities that are common across crisis episodes for different countries and over time, trying to identify early predictors of financial crises. We use their data and analysis to inform the parameterization of our model.[12]

Schularick and Taylor (2012) assume and test that the probability of entering a financial crisis can be a logistic function of macro and financial predictors. Their baseline logit specification finds that the five annual growth rates of bank loans from $t-4$ to t are jointly statistically significant predictors of episodes of financial instability that start in period $t+1$. Other variables, such as measures of real activity, inflation, or stock price gains, have little explanatory power when added to their baseline regression that includes lagged real bank loan growth, suggesting that financial crises are in fact "credit booms gone bust."[13]

Using the data set of Schularick and Taylor (2012), we estimate a slightly simplified version of their model, in which the probability of a financial crisis occurring in country i and year t is $\gamma_{i,t} = \exp(X_{i,t})/(1+\exp(X_{i,t}))$, and $X_{i,t}$ is assumed to be related linearly to the financial condition variable, L_t:

$$X_{i,t} = h_0 + h_i + h_1 L_t, \tag{13}$$

where h_0 is an intercept, h_i denotes a fixed effect for country i, and h_1 is the sensitivity of the crisis probability to the regressor L_t, as in model equation (4).[14]

[12]Schularick and Taylor (2012) also study how the role of monetary policy in sustaining aggregate demand, credit, and money growth has changed after the Great Depression.

[13]Among related studies, Laeven and Valencia (2013) collect a comprehensive database on systemic banking crises and propose a methodology to date banking crises based on policy indexes. Gourinchas and Obstfeld (2012) provide a similar study including developing countries and currency crisis episodes over the years 1973–2010. They find the share of aggregate credit over GDP to be a statistically significant predictor of financial and currency crises. Krishnamurthy and Vissing-Jorgensen (2012) note that short-term lending constitutes the most volatile component of credit over GDP and find that it plays a significant role in event-study logit regressions.

[14]For identification purposes the coefficient h_i for the United States is set to 0.

Let B_t denote the level of nominal bank loans to domestic households and non-financial corporations (henceforth the "non-financial sector") in year t, and P_t the price level. We define our predictor of a financial crisis occurring at time $t+1$ as the five-year cumulative growth rate of real bank loans from time $t-4$ to t:

$$L_t^a = \Sigma_{s=0}^4 \Delta \log \frac{B_{i,t-s}}{P_{i,t-s}}. \tag{14}$$

We verify that the variable L_t^a is a statistically significant predictor of financial crises for Schularick and Taylor's panel of countries, in the spirit of the five years of real bank loans growth in their original specification, and estimate the values of h_0 and h_1 reported in table 1. Our estimates suggest that an increase of 10 percentage points in the five-year real bank loan growth from 20 percent to 30 percent raises the annual probability of a financial crisis by less than 1 percentage point, from 4.9 percent to 5.6 percent. For robustness, in section 3.2 we consider alternative parameterizations in which the crisis probability is more responsive to the changes in credit conditions and economic outlook and (indirectly) to changes in the policy rate.

In order to adapt Schularick and Taylor's logit estimates, based on annual data, to the quarterly frequency of our model, we redefine our predictor L_t^a in equation (14) as the twenty-quarter sum of real bank loans growth:

$$L_t^q := \sum_{s=0}^{19} \Delta \log \frac{B_{t-s}}{P_{t-s}}. \tag{15}$$

We approximate equation (15) by the recursive twenty-quarter sum,

$$L_t^q \approx \Delta \log \frac{B_t}{P_t} + \frac{19}{20} L_{t-1}^q, \tag{16}$$

in order to limit the number of state variable of our model and help reduce the computational burden to find its solution.[15] To calibrate

[15]Figure B1 in online appendix B displays the differences between the financial condition indicators in equations (14), (15), and (16) over the period 1960:Q1–2008:Q4 at annual (left) and quarterly (right) frequencies.

our credit conditions equation (3) in the model, we first observe that the time-t component of the recursive sum in equation (16) is the difference between the nominal growth rate of bank loans and quarterly inflation:

$$\Delta \log \frac{B_t}{P_t} = \Delta \log B_t - \pi_t. \tag{17}$$

We can therefore estimate a reduced-form equation governing the evolution of quarterly nominal credit growth, ΔB_t, on U.S. data for the post-war period. We assume that the quarterly growth rate of nominal bank loans depends on a constant, c, and can vary with the monetary policy instrument, i_t, and with the output gap, y_t, and inflation π_t:

$$\Delta \log B_t = c + \phi_i i_t + \phi_y y_t + \phi_\pi \pi_t + \varepsilon_t^B. \tag{18}$$

Estimating this reduced-form equation for growth of bank lending does not allow us to separately identify how shifts in the demand and supply of credit translate into loan growth. Moreover, the direction of causality between the left- and right-hand-side variables can be questioned. To ameliorate a potential simultaneity bias, we use lagged values of i_t and y_t as instruments for their current values. We find that the coefficient on the policy rate is statistically insignificant and we calibrate it to zero, while the output gap and inflation enter the equation with positive and statistically significant coefficients (see appendix B for more details).[16]

Combining equations (16), (17), and (18), we obtain a simple dynamic equation describing the evolution of our credit conditions variable, L_t:

$$L_t \approx \rho_L L_{t-1} + \phi_0 + \phi_y y_t + (\phi_\pi - 1)\pi_t, \tag{19}$$

which we adapt to our two-period model notation as

$$L_1 \approx \rho_L L_0 + \phi_0 + \phi_y y_1 + (\phi_\pi - 1)\pi_1. \tag{20}$$

[16] As a robustness check, we also study how optimal policy changes when interest rate policy has a direct negative effect on credit growth. We find that when interest rate policy is more effective in reducing the crisis probability, the policymakers would find it optimal to adjust the policy rate upward by a larger margin compared with the baseline calibration. See online appendix E.3 for more details.

As indicated in table 1, a positive output gap of 1 percent is associated with 0.18 percent higher real bank loans growth, while a 1 percent increase in inflation lowers real bank loans growth by 0.57 percent (see parameter estimates of equation (20) in appendix B). Even though the central bank cannot directly affect the crisis predictor L_1 by changing the nominal interest rate i_1, the effects of monetary policy on output and inflation will also influence the growth rate of bank loans in the model and therefore the probability of a crisis. In particular, tighter monetary policy will lower the output gap and inflation and indirectly reduce financial instability. On the other hand, tighter monetary policy can lower inflation and increase financial instability, as in Svensson (2014). Since the Phillips curve is calibrated to be fairly flat to be in line with U.S. empirical estimates over recent decades, the response of inflation to the output gap is only modest and so is the response of inflation to monetary policy. As a result, tighter monetary policy in the model reduces credit growth and financial instability.

2.3.2 The Severity of the Crisis in the Baseline Calibration

Inflation and the output gap in the crisis state are chosen to roughly capture the severity of the Great Recession. In particular, we follow Denes, Eggertsson, and Gilbukh (2013) and assume that a financial crisis leads to a 10 percent decline in the output gap ($y_{2,c}$) and a 2 percent decline in inflation ($\pi_{2,c}$). We assume the expected duration of the crisis to be eight quarters. The continuation loss in the crisis state, $W_{2,c}$, is determined by the crisis-state inflation and the output gap, as well as by the expected crisis duration. In section 3.2, we offer sensitivity analyses under two alternative parameterizations, one in which the crisis episode is more severe (similar in scope to the Great Depression) and one in which the depth of the crisis is increasing in the degree of financial instability.

3. Optimal Policy and Financial Instability

In this section we describe the tradeoff faced by the policymaker and describe the optimal policy results under our baseline calibration with optimistic expectations, and compare the results with the rational expectations case. We also perform some sensitivity analyses by varying key parameters that affect the monetary policy transmission in the model.

3.1 A Key Intertemporal Tradeoff

We begin by illustrating the nature of the tradeoff the central bank faces in choosing the policy rate in the model with optimistic private-sector expectations. For that purpose, the top two panels of figure 1 show how the policy rate affects the output gap and inflation today. The middle panels show how the policy rate affects today's loss (as a function of output gap and inflation today) and the continuation loss. The bottom-left panel shows how the policy rate affects the overall loss function, which is the sum of today's loss and the continuation loss. Finally, the bottom-right panel shows how the policy rate affects the probability that a financial crisis can occur tomorrow. In this figure, L_0 is set to 0.2, roughly corresponding to the average value of the crisis predictor in U.S. data over the past five decades.

The top panels of figure 1 show that as the central bank increases the policy rate, inflation and the output gap decline, from equations (1) and (2). In the absence of any changes in the policy rate from its natural rate, inflation and the output gap are slightly below 2 percent and zero, respectively, because households and firms attach a small probability to large declines in inflation and output in the next period, should a crisis occur. Since the policy rate today reduces inflation and the output gap linearly and the policymaker's loss today is a quadratic function of these two variables, an increase in the policy rate increases today's loss quadratically (middle-left panel). On the other hand, the continuation loss decreases with the policy rate, as shown in the middle-right panel. This is because an increase in the policy rate, together with the associated declines in inflation and the output gap, worsens credit conditions at time $t = 1$, L_1, which in turn lowers the crisis probability, γ_1. The optimal policy rate balances the losses from lower economic activity today against the expected benefits from a reduced crisis probability next period. According to the bottom-left panel, under our baseline parameters, the overall loss is minimized when the nominal policy rate (and the real rate; see red dashed line in figure 2) are about 3 basis points above their long-run natural levels of 4 percent and 2 percent.

Figure 1. A Key Tradeoff Faced by the Central Bank

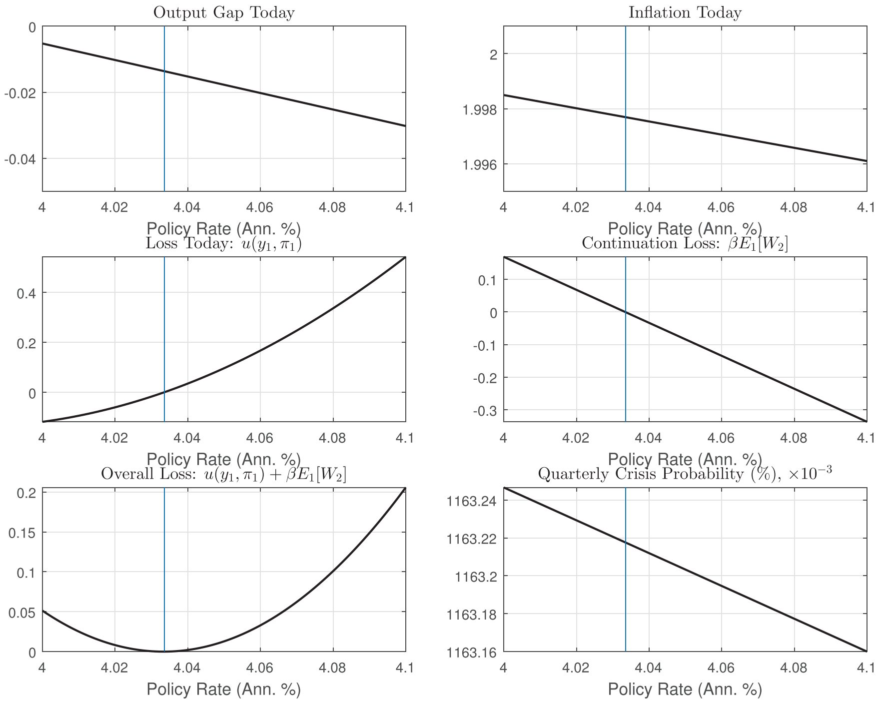

Data Source: Authors' calculations.
Notes: In this figure L_0 is set to 0.2, which is roughly the average value of this variable in the United States over the past five decades. In the bottom-left panel, the blue vertical line shows the optimal policy rate—the policy rate that minimizes the overall loss. The welfare losses are expressed as the one-time consumption transfer at time 1 that would make the household as well off as the household in a hypothetical economy with efficient levels of consumption and labor supply, expressed as a percentage of the steady-state consumption, as described in Nakata and Schmidt (2014). Welfare losses are normalized to be zero at the optimum.

(Figures are shown in color in the online version at the IJCB website.) This is the point at which the marginal cost of increasing the policy rate on today's loss equals the marginal benefit of increasing the policy rate.[17] In non-crisis times, the policymaker is willing to

[17]Under a standard Taylor rule, the policy rate is 4 basis points below the natural rate. At that rate, inflation and output gap are closer to their steady-state level, but the crisis probability is higher. See figure E6 in online appendix E.

Figure 2. Credit Growth and Optimal Policy

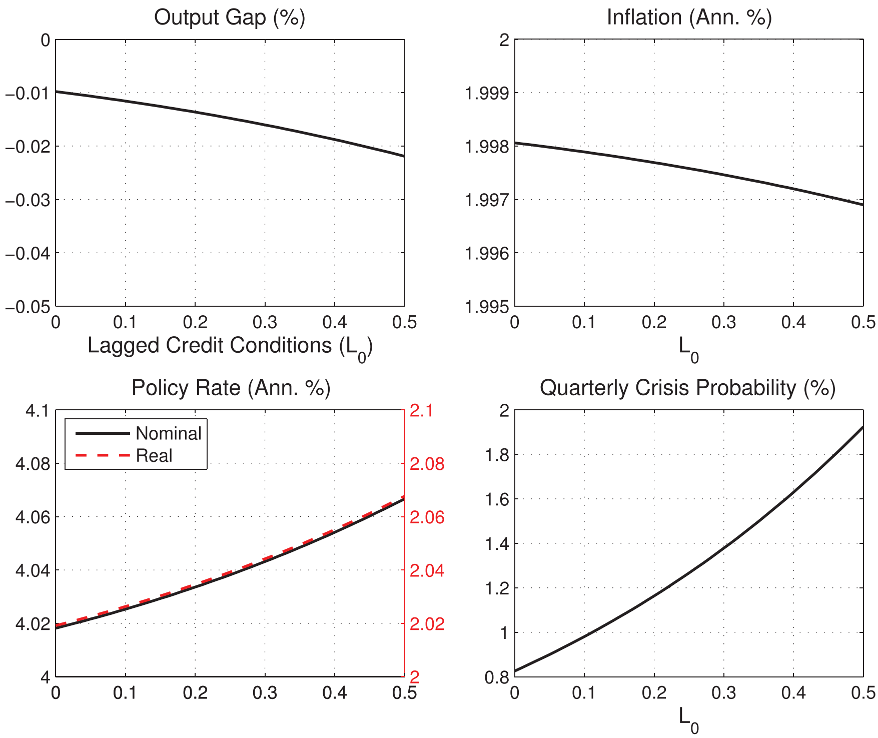

Data Source: Authors' calculations.
Note: This figure shows the optimal policy as a function of the initial level of the credit condition variable, L_0. See also the notes for figure 1.

optimally keep the policy rate slightly higher than its long-run natural rate, inducing a negative output gap and inflation lower than 2 percent, to reduce the probability of a financial crisis driven by exuberant credit conditions.

Our logit specification of the crisis probability equation implies that the effect of marginal changes in the policy rate on the crisis probability, and hence the continuation loss, depends on the lagged credit condition, L_0. To assess the effect of increasing concerns about financial stability on the optimal policy rate in the current period, we therefore vary in figure 2 the level of L_0 along the horizontal axis. Because an increase in the policy rate reduces the crisis probability by more when credit growth is already high, the optimal policy rate increases with lagged credit conditions. When $L_0 = 0$—roughly the minimum of this variable observed in the United States over the past five decades—the optimal increase in the policy rate is about 2 basis points. When $L_0 = 0.5$, the peak observed in the United States in post-war data (see figure B1 in online appendix B), the optimal increase in the policy rate is about 6 basis points.[18] Thus, even under conditions similar to those prevailing immediately prior to the onset of the financial crisis, the optimal adjustment to the short-term interest rate in response to potential financial stability risks would have been very small. The primary reason for this result is that the marginal effect of interest rate changes on the crisis probability, shown in the lower right panel of figure 1, is minuscule under our baseline model calibration. The marginal benefits of higher policy rates are outweighed by their marginal costs in terms of economic outcomes endured in times of no crisis, in line with results in Svensson (2016).

3.2 Alternative Scenarios

While the key parameters governing the crisis probability in equation (4) and the law of motion for credit conditions in equation (20) are based on empirical evidence, they are estimated with substantial uncertainty. In this section, we therefore examine the sensitivity of the result that financial stability considerations have little effect on optimal policy with respect to a range of alternative assumptions. In particular, we now analyze how the optimal policy rate and economic outcomes are affected by alternative assumptions regarding (i) the effectiveness of the policy rate in reducing the crisis probability, (ii) the severity of the crisis, and (iii) the alternative costs of increasing the policy rate on today's loss.[19] Table 2 reports the changes in the baseline parameters of the model that we adopt in

[18] This feature of optimal policy—the policy rate depending on the initial credit condition—would also arise even when the marginal crisis probability is constant if the severity of the crisis increases with the credit condition.

[19] In the online appendix we present an additional sensitivity analysis with respect to the parameter λ that controls the weight the central banker assigns to output stabilization. See figure E5 in online appendix E.

Table 2. Parameter Values for Sensitivity Analyses

Param.	Description	Value	Note
Tightening More Effective in Reducing the Crisis Probability			
h_1	Sensitivity of γ to L	3.0	+2 Std. Dev. from the Baseline
ϕ_y	Sensitivity of L to y	0.258	+2 Std. Dev. From the Baseline
A More Severe Crisis ("Great Depression")			
$y_{2,c}$	Output Gap in the Crisis State	–0.3	30% Drop in Output Gap
$\pi_{2,c}$	Inflation Gap in the Crisis State	–0.1/4	10% Drop in Annual Inflation
$W_{2,c}$	Loss in the Crisis State	$\dfrac{u(y_{2,c}\pi_{2,c})}{1-\beta\mu}$	
Tightening Less Costly for Today's Inflation and Output Gap			
σ	Sensitivity of y to i	1/2	Half of the Baseline Value
κ	Sensitivity of π to y	0.012	Half of the Baseline Value
Depth of Crisis Increasing in Credit Growth			
$\omega_{y,0}$	Drop in y if $L_1 = 0$	–0.03	Mild Recession
$\omega_{y,L}$	Sensitivity of y_c to L_1	–0.2	Great Recession for $L_1 = 0.35$
$\omega_{\pi,0}$	Drop in if $L_1 = 0$	–0.5/400	Mild Recession
$\omega_{\pi,L}$	Sensitivity of π_c to L_1	–0.043	Great Recession for $L_1 = 0.35$

our three sensitivity analyses. In section 4 we will consider how optimal policy is affected when the policymaker explicitly accounts for parameter uncertainty.

The columns of figure 3 show the optimal policy rate and the implied outcomes in terms of the output gap and inflation as functions of initial credit conditions, L_0, under three model parameterizations that differ from the baseline. The left

Figure 3. Credit Growth and Optimal Policy under Alternative Scenarios

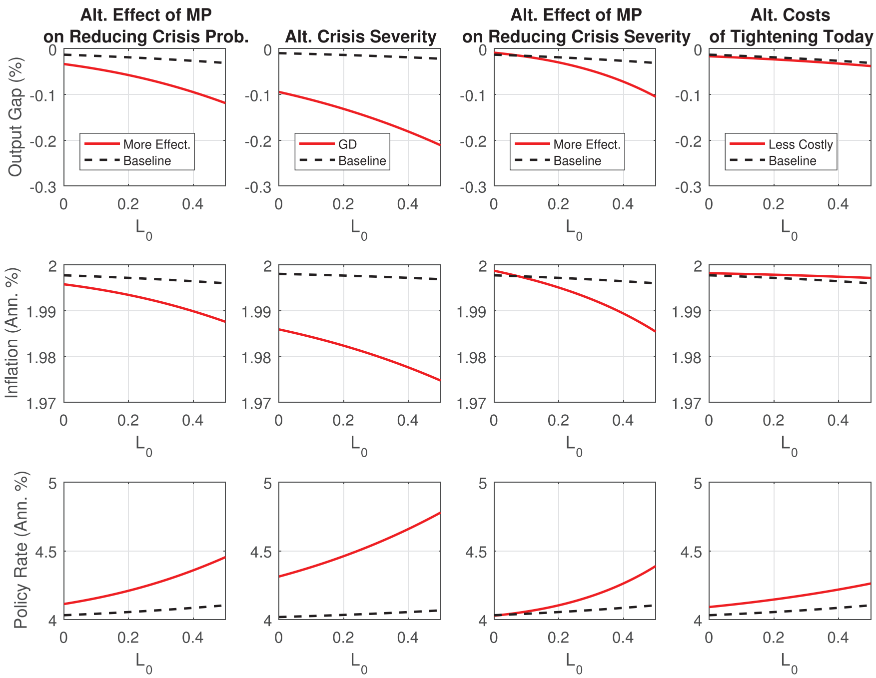

Data Source: Authors' calculations.
Note: This figure shows the optimal policy as a function of the initial level of the credit condition variable, L_0, under alternative calibrations of the model.

column of figure 3 corresponds to a model in which monetary policy tightening is more effective in reducing the crisis probability. As shown in top panel of table 2, we modify the sensitivity of the likelihood of a crisis to credit conditions, h_1, and the sensitivity of credit conditions to the output gap, ϕ_y, to be two standard deviations higher than the point estimates used in our baseline calibration. These higher sensitivities imply that an increase in the policy rate leads to a larger reduction in the crisis probability, and thus the optimal policy rate is higher for any value of L_0. With $L_0 = 0.2$, the optimal policy rate is about 25 basis points higher than the long-run natural rate of 4 percent. With $L_0 = 0.5$, the optimal policy rate is about 45 basis points higher than 4 percent, as seen in the bottom panel of the left column of figure 3. This additional incentive to tighten policy leads to lower inflation and output gap in the non-crisis state compared with our baseline, as shown in the top-left panels of the figure, as well as to a model without financial stability considerations.

The second column of figure 3 shows the output gap, inflation, and the policy rate under optimal policy when the severity of the crisis is of a magnitude roughly similar to that of the Great Depression. As shown in the middle panel of table 2, we assume that the output gap drops by 30 percent and inflation by 10 percent on an annual basis. A more severe crisis means that the benefit of raising the policy rate in reducing the continuation loss is larger, and thus the optimal policy rate is also higher for any values of L_0. With $L_0 = 0.2$, the optimal policy rate adjustment is about 30 basis points above the long-run natural rate of 4 percent. With $L_0 = 0.5$, the optimal policy rate adjustment is about 75 basis points over 4 percent, as seen in the bottom panel of the middle column of figure 3.

In a similar spirit, the third column of figure 3 shows the output gap, inflation, and policy rate under optimal policy when the depth of the crisis depends on the extent of the credit boom that preceded it, as described by the indicator L_1. This assumption echoes findings in Jordà, Schularick, and Taylor (2013) and Mian, Sufi, and Verner (2015) that suggest that excess credit preceding a recession can negatively affect the size of the downturn in real activity. We model the dependence of the size of the crisis on credit by assuming that y_c and π_c depend linearly on L_1:

$$y_c = \omega_{y,0} + \omega_{y,L} L_1 \tag{21}$$

$$\pi_c = \omega_{\pi,0} + \omega_{\pi,L} L_1. \tag{22}$$

We pick the parameters governing the linear relationship ($\omega_{y,0}$, $\omega_{y,L}$, $\omega_{\pi,0}$, $\omega_{\pi,0} < 0$) by assuming that when a crisis starts and the average lagged five-year credit growth, L_1, is 0 percent, then the output and inflation gaps fall, respectively, to –3 percent and –0.5 percent below target (a mild recession), while when average five-year credit growth is 35 percent (the value registered in the United States at the verge of the Great Recession), the output and inflation gaps fall, respectively, to –10 percent and –2 percent below target, as in the baseline calibration. In this framework, higher credit growth forecasts a higher likelihood of a crisis as well as more pronounced drops in output and inflation during a downturn. The higher the credit indicator, the more the central bank will want to avoid incurring in a severe crisis. In line with this intuition, the graphs show that the optimal policy rate adjustment with $L_0 = 0.2$ is about 15 basis points over the long-run natural rate of 4 percent, compared with 3 basis points in our baseline. With $L_0 = 0.5$, the optimal policy rate adjustment is more than 40 basis points above 4 percent.

Finally, the fourth column of figure 3 shows the output gap, inflation, and policy rate under optimal policy when today's inflation and output gap are less affected by the change in the policy rate than under the baseline. As listed in the lower panel of table 2, we assume that the sensitivity of the output gap to the policy rate and the sensitivity of inflation to the output gap are halved with respect to the baseline calibration. Less-responsive inflation and output gap mean that the effect of raising the policy rate on today's loss is small, and thus the optimal policy rate is higher for any values of L_0. With $L_0 = 0.2$, the optimal policy rate adjustment is about 10 basis points over the long-run natural rate of 4 percent. With $L_0 = 0.5$, the optimal policy rate adjustment is more than 20 basis points above 4 percent.

3.3 The Rational Expectations Case

Under rational expectations, in period 1 the private sector will attach probability γ_1 to a financial crisis hitting the economy in period 2, instead of the small and constant probability ϵ adopted under optimistic expectations.

Under this assumption, the private sector understands the link between credit growth and financial instability and forecasts future output and inflation accounting for the true probability of a financial crisis, γ_1, defined in equation (4).

When expectations of γ_1 are modeled as rational, times of plentiful credit conditions are associated with *reductions* of output and inflation, because the increased crisis probability reduces expected future inflation and output gaps, leading to lower inflation and a lower output gap today in the absence of any adjustment in the policy rate.

Figure 4. Credit Growth and Optimal Policy under Rational Expectations

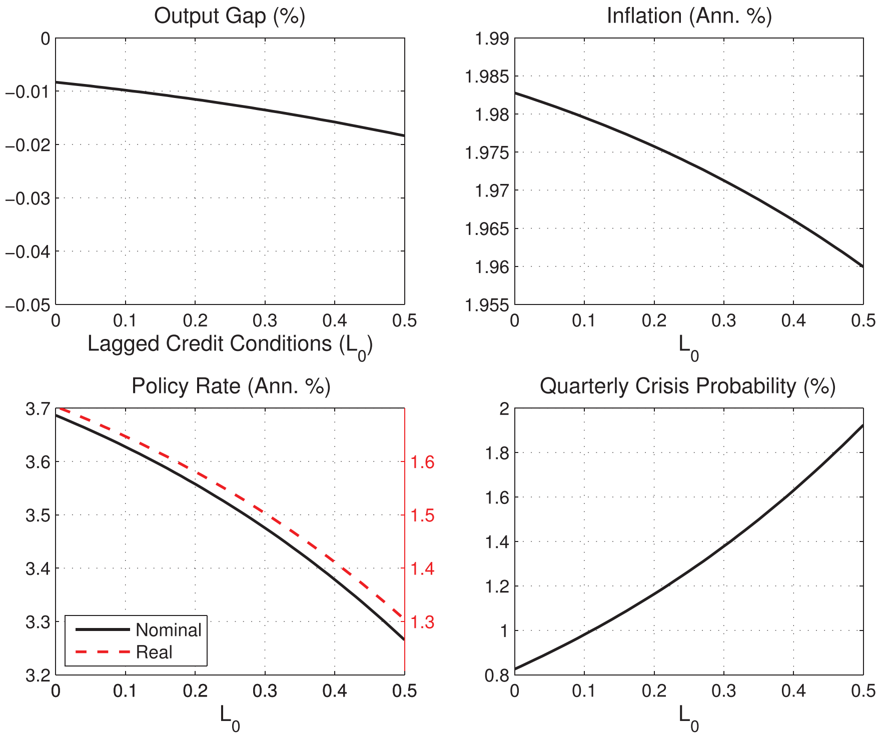

Data Source: Authors' calculations.
Note: This figure shows the optimal policy as a function of the initial level of the credit condition variable, L_0, under rational expectations.

Figure 4 shows optimal interest rate policy and outcomes as a function of initial conditions for credit growth, L_0, for the model with rational expectations. Precautionary-savings motives reduce aggregate output in period 1, in anticipation of a crisis hitting in period 2. Tighter monetary policy in this framework can still marginally reduce the likelihood of a future crisis, but comes at a cost of further worsening output and inflation outcomes in period 1. Optimal policy is in fact more accommodative than under optimistic expectations for any level of L_0, calling for nominal interest rates that are below the long-run natural rate of 4 percent and decreasing with the degree of financial instability (see Woodford 2012a).

This result seems inconsistent with much empirical evidence suggesting that times of buoyant financial conditions tend to be associated with private agents' expectations that these conditions will continue going forward (Shiller 2005, 2006). It is also at odds with the survey evidence discussed in section 2.1 and presented in appendix A as well as with experimental evidence on the existence of positive feedback in expectation formation discussed in Hommes (2011). Accordingly, in the remainder of the paper, we concentrate on studying optimal interest rate policy in several variants of our baseline model with optimistic expectations in which the relevance of the precautionary-savings channel is mitigated.

4. Optimal Policy under Parameter Uncertainty

We now consider how optimal policy is affected when the policymaker explicitly accounts for parameter uncertainty. We assume that the policymaker is uncertain about the value of specific parameters that affect the monetary policy transmission channel. We assume uncertainty around (i) the effectiveness of the policy rate in reducing the crisis probability, (ii) the severity of the crisis, and (iii) the alternative costs of increasing the policy rate on today's loss. We solve the model under two different assumptions on the policymaker's attitude towards uncertainty. We compute optimal interest rate policy both under the assumption of a Bayesian policymaker, as in Brainard (1967), and under the assumption of a robust policymaker, as in Hansen and Sargent (2008).

4.1 Sources of Uncertainty and Alternative Policymakers

Table 3 displays the prior distributions that we use to characterize uncertainty about the parameters. The first type of uncertainty is about two parameters related to the effectiveness of the policy rate in reducing the crisis probability: h_1 and ϕ_y. In our analysis below, we consider uncertainty about these two parameters separately. In the "no-uncertainty" case, h_1 takes the value of $h_{1,base}$ with probability one. When there is uncertainty and the policymaker is Bayesian, h_1 follows a discrete uniform distribution that takes the values of

Table 3. Calibration of Uncertainty

Parameter	Value	Probability
Uncertain Elasticity of Crisis Probability to Credit Conditions		
$h_{1,min}$	0.74	1/3
$h_{1,base}$	1.88	1/3
$h_{1,max}$	3.02	1/3
Uncertain Elasticity of Credit Conditions to Output		
$\phi_{y,min}$	0.102	1/3
$\phi_{y,base}$	0.18	1/3
$\phi_{y,max}$	0.258	1/3
Uncertain Severity of the Crisis		
$\pi_{2,c,min}$	−0.03/4	1/3
$\pi_{2,c,base}$	−0.02/4	1/3
$\pi_{2,c,max}$	−0.01/4	1/3
$y_{2,c,min}$	−0.15	1/3
$y_{2,c,base}$	−0.1	1/3
$y_{2,c,max}$	−0.05	1/3
Uncertain Effects of the Interest Rate on Today's π and y		
σ_{min}	0.5	1/3
σ_{base}	1	1/3
σ_{max}	1.5	1/3
κ_{min}	0.012	1/3
κ_{base}	0.024	1/3
κ_{max}	0.036	1/3

$h_{1,min}$, $h_{1,base}$, and $h_{1,max}$, each with probability 1/3.[20] Notice that the expected value of h_1 is $h_{1,base}$.[21] When the policymaker is a robust decision maker, he or she considers the value of h_1 in the closed interval $[h_{1,min}, h_{1,max}]$. Uncertainty about ϕ_y follows a similar structure. Specific parameter values are listed in the top panel of table 3.

[20] This is done for computational tractability.

[21] That is, these distributions imply mean-preserving spreads on these parameters.

The second type of parameter uncertainty is related to the severity of the crisis in terms of inflation and output outcomes in period $t = 2$: $\pi_{2,c}$ and $y_{2,c}$. Uncertainty regarding them is jointly analyzed and structured in the same way (see the middle panel of table 3). Finally, we consider the effects of uncertainty about two parameters that directly control the effects of changes in the policy rate on today's inflation and output: σ and κ, respectively. Uncertainty regarding them is analyzed jointly and structured in the same manner as above (see the bottom panel of table 3).

4.1.1 A Bayesian Policymaker

The Bayesian policymaker problem at time 1 is given by

$$W_1 = \min_{i_1} \int E_1\left[u(y_1, \pi_1) + \beta W_2 \mid \theta\right] dp(\theta) \tag{23}$$

subject to the private-sector equilibrium conditions described in the previous section and assuming that the private-sector agents perceive the probability of the crisis as constant and negligible; but now the policymaker takes expectations of future welfare losses with respect to the joint distribution of future states and the uncertain subset of parameters θ. This formulation of the problem follows that in the classic work of Brainard (1967) as recently restated by Brock, Durlauf, and West (2003) and Cogley et al. (2011).

4.1.2 A Robust Policymaker

The problem faced by a policymaker following a robust strategy is given by

$$W_1 = \min_{i_1} \left[\max_{\theta \in [\theta_{min}, \theta_{max}]} u(y_1, \pi_1) + \beta E_1[W_2] \right] \tag{24}$$

subject to the same set of private-sector equilibrium constraints and private agents' expectations. Following the literature, we will refer to the hypothetical agent who maximizes the welfare loss as the hypothetical evil agent who resides inside the head of the robust policymaker. The vector of parameters θ is a subset of the model parameters that are subject to uncertainty, and θ_{min} and θ_{max} are the lower and upper bounds considered by the hypothetical evil agent when he or she maximizes the welfare loss, respectively. This *min-max* formulation is standard in the literature on robustness (Hansen and Sargent 2008). While the robustness literature typically focuses on uncertainty arising from the distribution of exogenous shocks, uncertainty in our model comes from parameter values. Thus, our analysis closely follows those of Giannoni (2002) and Barlevy (2009), who also consider the problem of the robust decision maker under parameter uncertainty.[22]

4.2 Uncertainty about the Crisis Probability

Figure 5 illustrates how the presence of uncertainty around the estimate of the sensitivity of the crisis probability to credit conditions, h_1, affects the intertemporal tradeoff faced by the Bayesian policymaker. For each panel, red solid and black dashed lines refer to the cases with and without uncertainty, respectively.

An increase in uncertainty regarding the effectiveness of interest rate policy in reducing the crisis probability leads the Bayesian policymaker to adjust the policy rate by a larger amount, which can be seen in the bottom-left panel of figure 5 for the case with initial credit conditions $L_0 = 0.2$.

The presence of uncertainty about the parameter h_1 does not alter the period $t = 1$ loss function since the crisis probability does not affect how the policy rate influences today's inflation and output outcomes. This can be seen in the middle-left panel and top two panels of figure 5. However, the presence of uncertainty does affect the expected continuation loss for period $t = 2$. As shown in the middle-right panel, the slope of the expected welfare loss function is steeper with uncertainty than without it. This means that the marginal gain of policy tightening is larger with uncertainty than without it. With the marginal costs of policy tightening unchanged in $t = 1$, this higher marginal gain of policy tightening translates into an optimal policy rate that is higher than in the absence of

[22]Hansen and Sargent (2014) also consider the problem of the robust policymaker under parameter uncertainty. In their work, a parameter is a random variable and the hypothetical evil agent is allowed to twist the probability distribution of uncertain parameters. In our paper as well as in Giannoni (2002) and Barlevy (2009), a parameter is a scalar and the hypothetical evil agent is only allowed to choose an alternative value for the parameter.

Figure 5. The Tradeoff Facing the Bayesian Policymaker

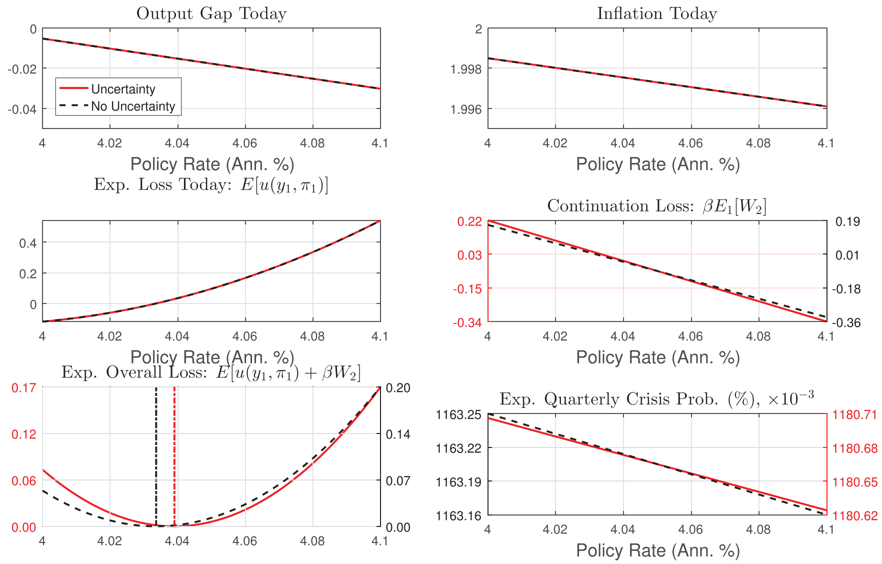

Data Source: Authors' calculations.
Notes: In this figure, L_0 is set to 0.2, which is roughly the average value of this variable in the United States over the past five decades. In the bottom-left panel, vertical black dashed and red solid lines are for the optimal policy rates without and with uncertainty. The welfare losses are expressed as the one-time consumption transfer at time 1 that would make the household as well off as the household in a hypothetical economy with efficient levels of consumption and labor supply, expressed as a percentage of the steady-state consumption, as described in Nakata and Schmidt (2014). Welfare losses are normalized to be zero at the optimum.

uncertainty even if just by a few decimals of a basis point, as seen in the lower-left panel.

As shown in the middle-right panel of figure 5, the slope of the expected continuation loss is steeper under uncertainty because the *expected* crisis probability under uncertainty is steeper than that of the (expected) crisis probability without uncertainty.[23] Figure 6 shows that when h_1 increases, both the slope and the level of the

[23] Note that, since the non-crisis value is zero (i.e., $W_{2,nc} = 0$), the expected continuation loss is a constant times crisis probability (i.e., $\beta E_1[W_2]=\beta\left[(1-\gamma_1) W_{2,nc} + \gamma_1 W_{2,c}\right] = \beta W_{2,c}\gamma_1$).

Figure 6. The Effect of a Mean-Preserving Spread on h_1 for the Crisis Probability Function: $\gamma_1 = \frac{\exp(h_0+h_1L_1)}{1+\exp(h_0+h_1L_1)}$

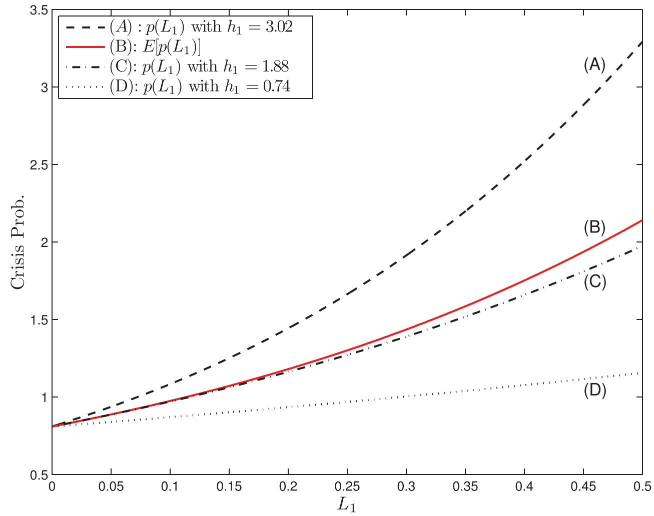

Data Source: Authors' calculations.
Note: This figure shows the effect of a mean-preserving spread on h_1 for the crisis probability function γ_1.

crisis probability function increase; this is captured in the steeper slope of line (A) with respect to line (C) (the baseline) in figure 6. When h_1 decreases, the slope, as well as the level, of the crisis probability function decreases, which is captured in the flatter slope of line (D) with respect to line (C) in figure 6. The convexity of the logit function implies that the increase in the slope of the crisis probability due to an increase in h_1 is larger than the decrease in the slope of the crisis probability due to a decrease in h_1 of the same magnitude. As a result, the slope of the *expected* crisis probability is steeper than that of the crisis probability function, captured by the fact that the slope of the red solid line (B) is steeper than that of the baseline calibrated function (C). That is, a mean-preserving spread in h_1 increases the slope of the (expected) crisis probability function.

As demonstrated in figure 7, this result does not depend on the level of credit conditions in the economy. The optimal adjustment

Figure 7. Optimal Policy under Uncertainty: Bayesian Policymaker (uncertain effects of policy on the crisis probability)

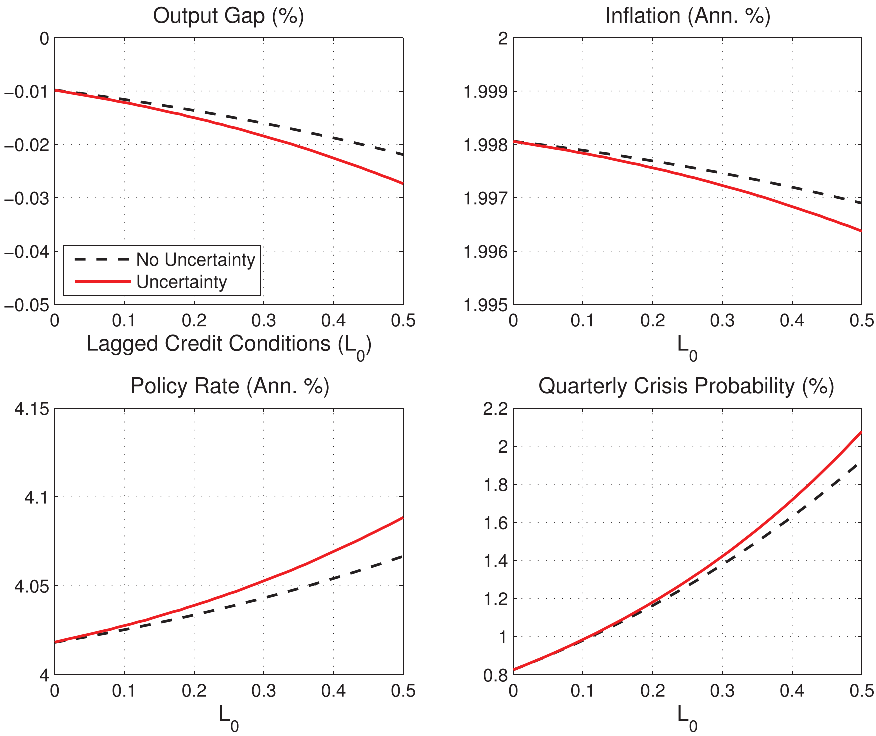

Data Source: Authors' calculations.
Note: This figure shows the optimal policy as a function of the initial level of the credit condition variable, L_0, with and without uncertainty, for a Bayesian policymaker.

of the policy rate is about 10–20 percent larger in the presence of uncertainty than in its absence and it is increasing in initial credit conditions, L_0.

Uncertainty about the effects of policy on the probability of a crisis also leads the robust policymaker to choose a higher policy rate, which is shown in figure 8. The policymaker following robust policies chooses the policy rate to minimize the welfare loss under the worst-case scenario. In the present context, the parameter value

Figure 8. Optimal Policy under Uncertainty: Robust Policymaker (uncertain effects of policy on the crisis probability)

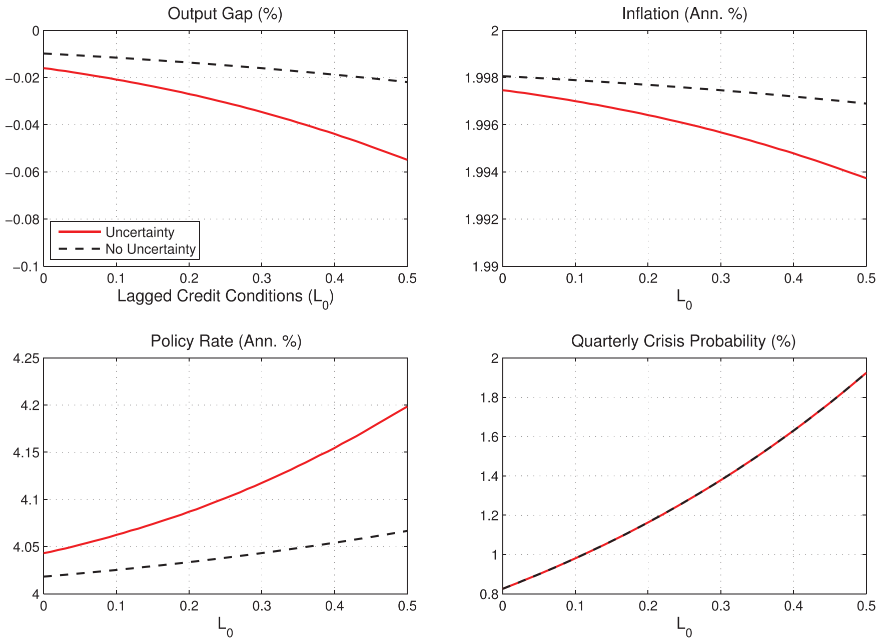

Data Source: Authors' calculations.
Note: This figure shows the optimal policy as a function of the initial level of the credit condition variable, L_0, with and without uncertainty, for a robust policymaker.

that leads to the maximum welfare loss is the highest h_1, as this implies higher crisis probability for *any given choice of* i_1. This is illustrated in figure 9, which shows the payoff function of the hypothetical evil agent when the robust policymaker chooses the optimal policy rate under no uncertainty of 4.03 percent. By choosing the maximum possible h_1, the hypothetical evil agent can cause the largest damage to the robust policymaker. Thus, the robust policymaker chooses the policy rate in order to minimize the welfare loss, *anticipating* that the hypothetical evil agent would choose the highest possible h_1. A higher h_1 means that an increase in the policy rate leads to a larger decline in the continuation value.

Figure 9. The Objective Function of the Hypothetical Evil Agent inside the Head of the Robust Policymaker

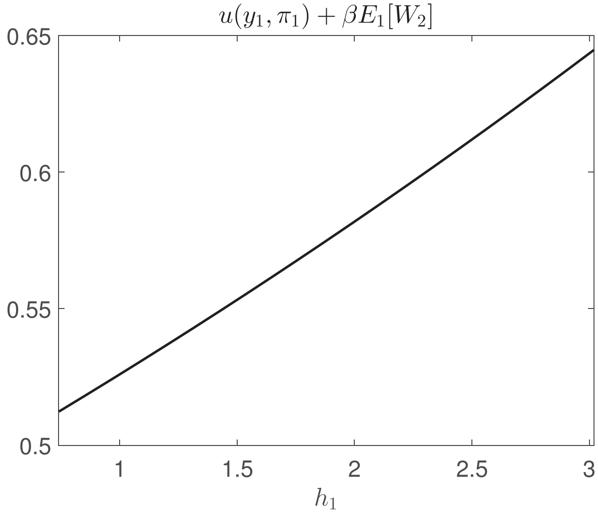

Data Source: Authors' calculations.

Thus, the robust policymaker adjusts the policy rate by more under uncertainty. In our calibration, the presence of uncertainty leads the robust policymaker to adjust the policy rate by 100–200 percent more. If, for example, initial credit conditions are particularly buoyant, with $L_0 = 0.5$, the robust policymaker would want to set the policy rate in the non-crisis state just below 4.2 percent, compared with 4.06 percent in the absence of uncertainty.

4.3 Uncertainty about Credit Conditions

Figure 10 shows how the uncertainty regarding the elasticity of credit conditions to output affects optimal policy. The left and right columns are for the Bayesian policymaker and the robust policymaker, respectively.

We verify that the presence of uncertainty leads the Bayesian policymaker to choose a higher policy rate; however, the difference between optimal policy with and without uncertainty is negligible,

Figure 10. Optimal Policy under Uncertainty: Uncertain Elasticity of Credit Conditions to Output

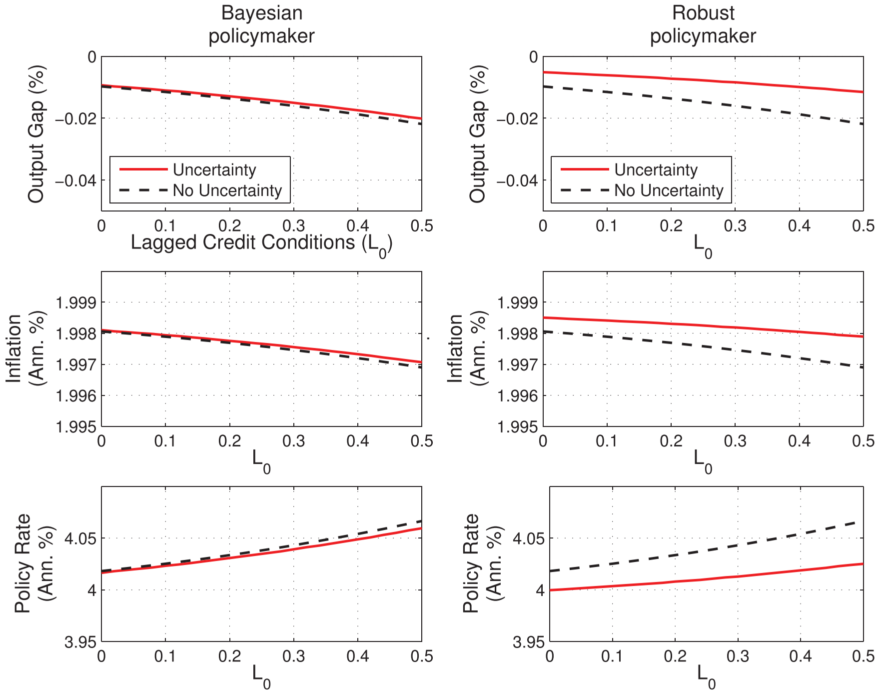

Data Source: Authors' calculations.
Note: This figure shows the optimal policy as a function of the initial level of the credit condition variable, L_0, with and without uncertainty, for a Bayesian (left) and robust (right) policymaker.

as shown by the black dashed and red solid lines in the left column. We find that uncertainty regarding the elasticity of credit conditions to output induces uncertainty about credit conditions today. This also makes the crisis probability uncertain. The convexity of the logit function implies that a mean-preserving spread in L_1 increases the level and slope of the (expected) crisis probability, which in turn increases the marginal benefit of policy tightening. However, in our calibration, this effect is very small.

The right column shows that this type of uncertainty leads the robust policymaker to choose a lower policy rate instead of a higher one. In this context, the parameter value that leads to the maximum welfare loss is the lowest possible value for the parameter ϕ_y, as it implies a higher crisis probability *for any choice of* i_1. Thus, the robust policymaker sets the policy rate in order to minimize the welfare loss, *expecting* the hypothetical evil agent to choose the lowest possible ϕ_y. Notice that a low value of the parameter ϕ_y means that a policy tightening has, via aggregate demand, a weaker effect on credit conditions and the crisis probability. Facing a lower marginal benefit of policy tightening and a lower unchanged marginal cost, the policymaker chooses a lower policy rate than in the absence of uncertainty.

4.4 Uncertain Severity of the Crisis

Figure 11 shows how uncertainty regarding both inflation and output levels induced by a crisis, $(\pi_{2,c}, y_{2,c})$, affects the optimal policy under a Bayesian and robust policymaker, respectively. The figure shows that, regardless of the type, the policymaker chooses a higher policy rate in the presence of uncertainty than in the absence of it.

Why does uncertainty about the severity of the crisis lead the Bayesian policymaker to choose a higher policy rate? Uncertainty regarding the severity of the crisis does not affect today's output gap, inflation, and loss. However, it does affect the (expected) continuation loss. In particular, the slope of the (expected) continuation value is steeper with uncertainty than without it. This is because the loss associated with the crisis state tomorrow is quadratic. As a result, an increase in the loss due to a decline in inflation is larger than a decline in the loss due to an increase in inflation of the same magnitude. Similarly, an increase in the loss due to a decline in output gap is larger than a decline in the loss due to an increase in output gap of the same magnitude. Thus, the presence of uncertainty regarding $\pi_{2,c}$ and $y_{2,c}$ increases the expected loss associated with the crisis state. The marginal benefit of policy tightening increases when the expected crisis loss increases (i.e., $\beta E_1[W_2] = \beta\gamma_1 W_{2,c}$). Accordingly, the marginal benefit of policy tightening is higher with uncertainty than without it. With the marginal cost of policy tightening unchanged, the higher marginal benefit of a reduced expected loss means that the optimal policy rate will be higher.

Figure 11. Optimal Policy under Uncertainty: Uncertain Severity of the Crisis

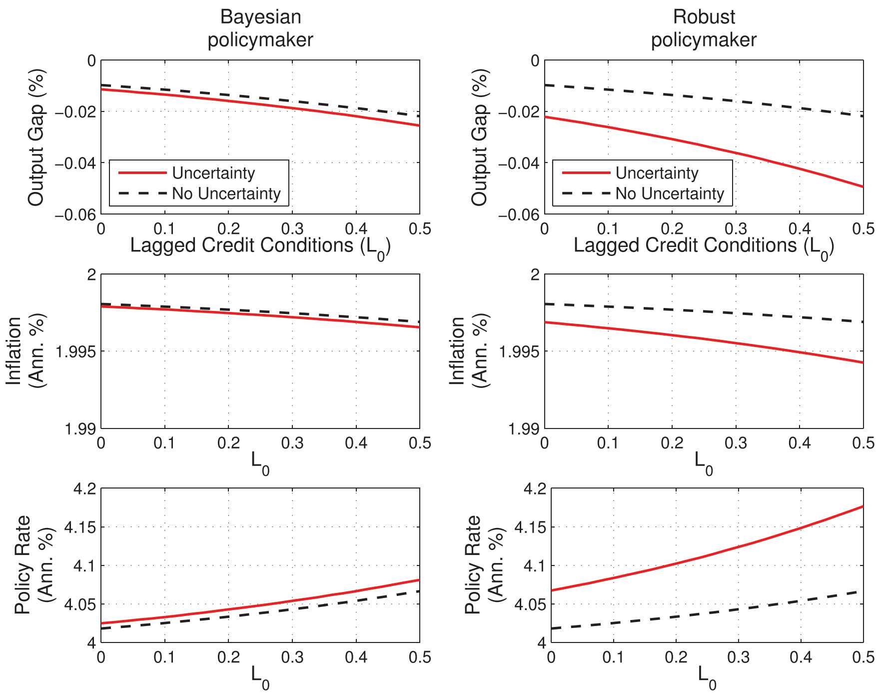

Data Source: Authors' calculations.
Note: This figure shows the optimal policy as a function of the initial level of the credit condition variable, L_0, with and without uncertainty, for a Bayesian (left) and robust (right) policymaker.

Similarly, the robust policymaker chooses a higher policy rate in the presence of this uncertainty. The hypothetical evil agent can reduce the welfare by choosing the largest possible declines in inflation and output gaps in the crisis state. This means that, for the robust policymaker, the marginal change in the continuation value associated with an adjustment of the policy rate is larger under uncertainty. Accordingly, the presence of uncertainty leads the robust policymaker to adjust the policy rate by more under uncertainty to avoid the unpleasant crisis scenario, as shown in the right-hand-side panels of figure 11.

Figure 12. Optimal Policy under Uncertainty: Uncertain Effects of Monetary Policy on Today's Inflation and Output

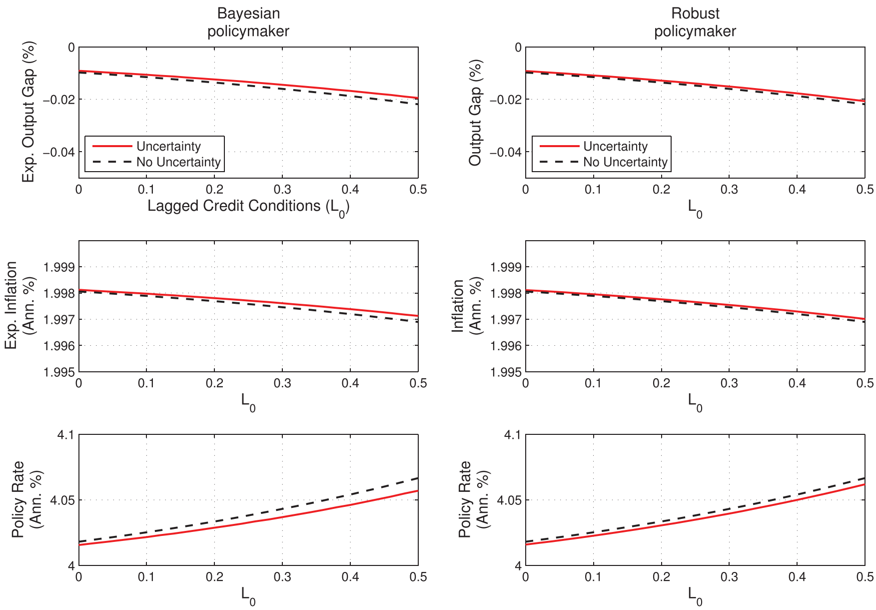

Data Source: Authors' calculations.
Note: This figure shows the optimal policy as a function of the initial level of the credit condition variable, L_0, with and without uncertainty, for a Bayesian (left) and robust (right) policymaker.

4.5 Uncertain Effects of Policy on Today's Inflation and Output

Figure 12 shows the effect on the optimal policy of uncertainty over the parameters (σ, κ), that is the effects of interest rates on today's inflation and output. The two columns correspond to the Bayesian policymaker and the robust policymaker, respectively. They show that both types of agents choose a lower policy rate in the presence of uncertainty than in the absence of it. This is the same type of uncertainty considered in Brainard (1967), and our result is consistent with his conclusion.

Figure 13. The Tradeoff Facing the Bayesian Policymaker: Uncertain Effects of Policy on Today's Allocations

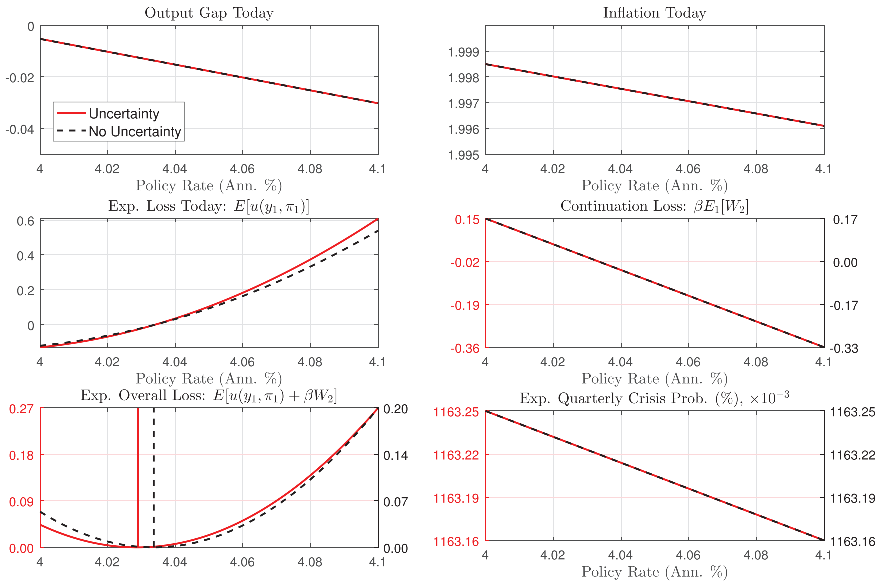

Data Source: Authors' calculations.
Notes: In this figure, L_0 is set to 0.2, which is roughly the average value of this variable in the United States over the past five decades. In the bottom-left panel, vertical black dashed and red solid lines are for the optimal policy rates without and with uncertainty. The welfare losses are expressed as the one-time consumption transfer at time 1 that would make the household as well off as the household in a hypothetical economy with efficient levels of consumption and labor supply, expressed as a percentage of the steady-state consumption, as described in Nakata and Schmidt (2014). Welfare losses are normalized to be zero at the optimum.

Why does uncertainty lead the Bayesian policymaker to choose a lower policy rate? As shown in the top two panels of figure 13, uncertainty about the parameters σ and κ does not change the expected inflation and output gap today. This is because today's inflation and output depend linearly on the policy rate. However, this uncertainty does affect the expected loss today. This is because the central bank's welfare loss today is quadratic in inflation and output. As shown in the middle-left panel, the expected loss is larger with uncertainty than without it, and so is the marginal cost of policy tightening. While the presence of uncertainty has some effects on the (expected) continuation loss, they are negligible and the marginal benefits of policy tightening are essentially unchanged under uncertainty. Accordingly, the central bank will optimally set the policy rate lower in the presence of uncertainty than in the absence of it.

The robust policymaker also chooses a lower policy rate under uncertainty. In our calibration, the hypothetical evil agent inside the head of the central banker chooses the smallest possible σ and the largest possible κ. While a smaller σ increases welfare through higher (or less negative) output gap and inflation, it decreases welfare through higher L_1 and crisis probability. The hypothetical evil agent chooses the smallest possible σ because the second force dominates the first. The evil agent chooses the highest possible κ because a higher κ is associated with a lower (more negative) inflation and a higher L_1, both of which reduce welfare. Anticipating that the hypothetical evil agent would choose a smaller σ, the robust policymaker has an incentive to adjust the policy rate by more because an increase in the policy rate has a smaller consequence on today's output. Anticipating that the evil agent would choose a higher κ, the robust policymaker has an incentive to adjust the policy rate by less because an increase in the policy rate has a larger consequence on today's output. In our calibration, the second effect dominates the first, leading the robust policymaker to choose a lower policy rate under uncertainty, as shown in the right column of figure 12.

The effect of uncertainty over the parameters σ and κ that control the effects of policy on today's inflation and output is consistent with Brainard's *attenuation principle*: the policymaker optimally sets the policy rate lower than in the absence of uncertainty. However, our analysis shows that this principle does not generalize to other types of uncertainty. There are several reasons why this difference arises. On the one hand, in Brainard's work uncertainty increases the marginal cost of monetary policy tightening today with negligible changes in the marginal expected loss tomorrow, so that the policymaker chooses a lower optimal policy rate in equilibrium. On the other hand, we find that uncertainty that increases the future marginal benefit of monetary policy interventions (either because the policymaker is unsure about how credit conditions affect the probability of a crisis or is uncertain about the size of the output and inflation drops in the crisis state) tends to amplify the preemptive response of the policymaker. In these cases, uncertainty calls for a higher optimal rate due to the non-linearity of the expected future loss derived either from the logit function or the quadratic nature of the per-period loss.

5. Conclusions

We have analyzed how the central bank should respond in normal times to financial imbalances in a stylized model of financial crises. For the version of the model that is calibrated to match the historical correlation of credit booms and financial crises in advanced economies, we find that the optimal increase in the policy rate due to financial imbalances is negligible. We also take an additional step to identify circumstances that would lead the central bank to adjust the policy rate more aggressively. We show that if (i) the severity of the crisis is comparable to that of the Great Depression, or (ii) the crisis probability is twice more responsive to financial conditions in the economy, then the optimal adjustment to the policy rate can be as large as, or can even exceed, 50 basis points. Finally, we demonstrated that parameter uncertainty can induce a Bayesian and a robust policymaker to respond more aggressively to financial crises by setting the policy rate higher than in absence of uncertainty. This happens if the source of uncertainty can increase the expected marginal benefit of policy interventions aimed at reducing the likelihood of a crisis and its expected welfare loss.

To conclude, we want to highlight two important limitations of our analysis and suggest avenues for future research that could address them. First, while the two-period model setup proves useful in understanding and describing the key tradeoffs faced by policymakers, this modeling choice can also impose some serious constraints on the quantitative implications of the model. For example, relaxing the two-period assumption might allow to properly account for the effects of the policy rate setting on future allocations and welfare, as well as for the possibly long-lasting effects of interest rate policy on financial conditions and on the probability of a crisis occurring in the long run that are now left unmodeled.

Second, we took a semi-structural approach in which the relationship between the financial condition variable, L_t, and output, inflation, and the short-term nominal interest rate, y_t, π_t, and i_t, cannot be directly subsumed from standard microfounded DSGE models with financial frictions. In addition, the non-linear relationship between the crisis probability and financial conditions, while empirically motivated by Schularick and Taylor's work, is also not derived in the context of a general equilibrium model from the interaction of agents' optimizing behavior. This limitation is not unique to our paper; other recent papers examining the implication of financial stability for interest rate policy also assume that at least one of these the relationships is non-structural (Svensson 2014, 2016; Gourio, Kashyap, and Sim 2016). It would be interesting and important to examine whether the main findings of our analysis would extend to an environment in which the macrofinancial linkages are derived structurally.

References

Adrian, T., and N. Liang. 2013. "Financial Stability Monitoring." Finance and Economic Discussion Series No. 2013-21.

Aikman, D., A. G. Haldane, and B. D. Nelson. 2015. "Curbing the Credit Cycle." *Economic Journal* 125 (585): 1072–1109.

Barlevy, G. 2009. "Policymaking under Uncertainty: Gradualism and Robustness." *Economic Perspectives* (Federal Reserve Bank of Chicago) 2Q.

Bernanke, B. S., and M. Gertler. 1999. "Monetary Policy and Asset Price Volatility." *Economic Review* (Federal Reserve Bank of Kansas City)(Q IV): 17–51.

Borio, C. 2012. "The Financial Cycle and Macroeconomics: What Have We Learnt?" Working Paper No. 395, Bank for International Settlements.

———. 2016. "Towards a Financial Stability-Oriented Monetary Policy Framework?" Presentation at the "Central Banking in Times of Change" conference sponsored by the Central Bank of the Republic of Austria, Vienna, September 13–14.

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

Brock, W. A., S. N. Durlauf, and K. D. West. 2003. "Policy Evaluation in Uncertain Economic Environments." *Brookings Papers on Economic Activity* 34 (1): 235–322.

Clouse, J. A. 2013. "Monetary Policy and Financial Stability Risk: An Example." Finance and Economics Discussion Series No. 2013-41, Board of Governors of the Federal Reserve System.

Cogley, T., B. De Paoli, C. Matthes, K. O. Nikolov, and A. Yates. 2011. "A Bayesian Approach to Optimal Monetary Policy with Parameter and Model Uncertainty." *Journal of Economic Dynamics and Control* 35 (12): 2186–2212.

Denes, M., G. Eggertsson, and S. Gilbukh. 2013. "Deficits, Public Debt Dynamics and Tax and Spending Multipliers." *Economic Journal* 123 (566): 133–63.

Galí, J. 2014. "Monetary Policy and Rational Asset Price Bubbles." *American Economic Review* 104 (3): 721–52.

Giannoni, M. P. 2002. "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model." *Macroeconomic Dynamics* 6 (1): 111–44.

Gourinchas, P.-O., and M. Obstfeld. 2012. "Stories of the Twentieth Century for the Twenty-First." *American Economic Journal: Macroeconomics* 4 (1): 226–65.

Gourio, F., A. Kashyap, and J. Sim. 2016. "The Tradeoffs in Leaning Against the Wind." Working Paper.

Hansen, L. P., and T. J. Sargent. 2008. *Robustness.* Princeton, NJ: Princeton University Press.

———. 2014. "Four Types of Ignorance." Mimeo.

Hommes, C. 2011. "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab." *Journal of Economic Dynamics and Control* 35 (1): 1–24.

Jordà, Ò., M. Schularick, and A. M. Taylor. 2013. "When Credit Bites Back." *Journal of Money Credit and Banking* 45 (s2): 3–28.

Krishnamurthy, A., and A. Vissing-Jorgensen. 2012. "Short-term Debt and Financial Crises: What We Can Learn from U.S. Treasury Supply."

Laeven, L., and F. Valencia. 2013. "Systemic Banking Crises Database." *IMF Economic Review* 61 (2): 225–70.

Laseen, S., A. Pescatori, and J. Turunen. 2015. “Systemic Risk: A New Trade-off for Monetary Policy?” IMF Working Paper No. 15/142.

Mian, A. R., A. Sufi, and E. Verner. 2015. “Household Debt and Business Cycles Worldwide.” Working Paper No. 21581, National Bureau of Economic Research.

Nakata, T., and S. Schmidt. 2014. “Conservatism and Liquidity Traps.” Finance and Economics Discussion Series No. 2014-105, Board of Governors of the Federal Reserve System.

Onatski, A., and J. H. Stock. 2002. “Robust Monetary Policy under Model Uncertainty in a Small Model of the U.S. Economy.” *Macroeconomic Dynamics* 6 (1, February): 85–110.

Reinhart, C. M., and K. S. Rogoff. 2009. “The Aftermath of Financial Crises.” *American Economic Review* 99 (2): 466–72.

———. 2014. “Recovery from Financial Crises: Evidence from 100 Episodes.” *American Economic Review* 104 (5): 50–55.

Romer, C., and D. Romer. 2014. “New Evidence on the Impact of Financial Crises in Advanced Countries.” Mimeo, UC Berkeley.

Schularick, M., and A. M. Taylor. 2012. “Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008.” *American Economic Review* 102 (2): 1029–61.

Shiller, R. J. 2005. *Irrational Exuberance.* Princeton, NJ: Princeton University Press.

———. 2006. “Long-Term Perspectives on the Current Boom in Home Prices.” *Economists’ Voice* 3 (4).

Söderstrom, U. 2002. “Monetary Policy with Uncertain Parameters.” *Scandinavian Journal of Economics* 104 (1, March): 125–45.

Stein, J. 2013. “Overheating in Credit Markets: Origins, Measurement, and Policy Responses.” Speech delivered at the “Restoring Household Financial Stability after the Great Recession: Why Household Balance Sheets Matter” research symposium sponsored by the Federal Reserve Bank of St. Louis, St. Louis, Missouri, February 7.

Svensson, L. E. O. 2014. “Inflation Targeting and ‘Leaning against the Wind’.” *International Journal of Central Banking* 10 (2): 103–14.

———. 2016. "Cost-Benefit Analysis of Leaning Against the Wind: Are Costs Larger Also with Less Effective Macroprudential Policy?" IMF Working Paper No. 16/3.

Svensson, L. E. O., and N. Williams. 2007. "Bayesian and Adaptive Optimal Policy Under Model Uncertainty." Mimeo.

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

Woodford, M. 2003. *Interest and Prices*. Princeton, NJ: Princeton University Press.

———. 2012a. "Inflation Targeting and Financial Stability." *Economic Review* (Sveriges Riksbank) 1: 7–32.

———. 2012b. "Methods of Policy Accommodation at the Interest-Rate Lower Bound." In *The Changing Policy Landscape*. Proceedings of the Economic Policy Symposium sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyoming, August 30–September 1.

Yellen, J. L. 2014. "Monetary Policy and Financial Stability." Speech given at the 2014 Michel Camdessus Central Banking Lecture, International Monetary Fund, Washington, D.C., July 2.