ISSN 1936-5330

The Federal Reserve Bank of Kansas City
Economic Research Department

Optimal Inflation for the U.S. Economy

Roberto M. Billi

April 2007; Revised June 2008

RWP 07-03

Research Working Papers

Optimal Inflation for the U.S. Economy[1]

Roberto M. Billi[2]

First version: April 2007
This version: June 2008

RWP 07-03

Abstract: This paper characterizes the optimal inflation rate in a standard macro model, which accounts for an occasionally-binding zero lower bound on nominal interest rates and model uncertainty. Estimates of the optimal rate of inflation, as measured by the PCE price index, range 0.7 to 1.4 percent per year. Even under extreme model uncertainty, the optimal inflation rate is not as high as previous studies suggest.

Keywords: commitment, liquidity trap, long-run tradeoffs, monetary policy, nonlinear, robust control, stationary distribution, worst-case scenarios

JEL classification: C63, E31, E52

---

[1] This paper is a substantial revision of the third chapter of the author's doctoral dissertation at Goethe University Frankfurt, see Billi (2005). For many helpful comments and discussions, the author thanks Klaus Adam, Larry Ball, Marco Bassetto, Günter Beck, Ben Bernanke, Michael Binder, Brent Bundick, Larry Christiano, Richard Dennis, Steven Durlauf, Martin Eichenbaum, Jon Faust, Ben Friedman, Dale Henderson, Peter Ireland, George Kahn, Jinill Kim, Andy Levin, Ben McCallum, Benoit Mojon, Athanasios Orphanides, Dave Reifschneider, Stephanie Schmitt-Grohé, Pu Shen, Frank Smets, Jón Steinsson, Lars Svensson, Eric Swanson, Bob Tetlow, Hiroshi Ugai, Volker Wieland, John Williams, Noah Williams, Alex Wolman, Mike Woodford, seminar participants at the Federal Reserve Board, the Federal Reserve Banks of Atlanta, Boston, Chicago, Cleveland, Kansas City, and San Francisco, as well as participants at the ASSA Annual Meetings, the SCE Annual Meetings and the Joint Bank of Canada and European Central Bank Conference on "Defining price stability: Theoretical options and practical experience". The views expressed herein are solely those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System.

[2] Federal Reserve Bank of Kansas City, One Memorial Drive, Kansas City, MO 64198, United States, E-mail: Roberto.Billi@kc.frb.org

1 Introduction

Economists, policymakers, and the public agree that high inflation is costly for the economy because it ultimately reduces standards of living.[1] For a number of reasons, however, inflation can also be too low. In practice, policymakers aim for low, positive rates of inflation rather than zero inflation. However, rigorous estimates of the 'optimal inflation rate' which maximizes the economic well-being of the public are not available.

1.1 The Case for Low, Positive Rates of Inflation

One reason for keeping inflation above zero stems from the fact that nominal interest rates cannot fall below zero under normal circumstances.[2] Nominal interest rates are low when inflation is expected to remain low. Since monetary policy counteracts slowing economic activity by lowering short-term interest rates, very low inflation limits the extent policymakers can respond to an economic slowdown. Vickrey (1954), Phelps (1972), Okun (1981), Summers (1991), Fischer (1996), Fuhrer and Madigan (1997), Krugman (1998), McCallum (2000), Orphanides and Wieland (2000), Svensson (2003), Bernanke, Reinhart and Sack (2004), Kato and Nishiyama (2005), Wolman (2005), and others, argue in favor of low, positive rates of inflation as to leave room to ease policy in response to economic weakness.[3]

There are a number of other reasons, however, for which inflation can be too low:

Measurement error in inflation. Available measures of inflation tend to be biased upward. Recent estimates place the measurement bias for the personal consumption expenditure (PCE) price index, the Federal Reserve's preferred measure of inflation, at about 0.5 percentage point per year.[4]

Downward wage rigidity. Nominal wages may be downwardly rigid if firms are unable to make nominal wage cuts because workers are unwilling to accept them. Tobin (1972) conjectures that a little inflation may make it easier for firms to lower real wages and maintain employment in response to declining demand.

Debt–Deflation. A negative inflation rate—deflation—may be a more serious problem than inflation because deflation causes a decreases in the value of collateral used to secure a loan, or another form of debt. Fisher (1933) argues that falling prices could create a vicious cycle of financial distress for banks and other lenders which would lead to more downward pressure on prices.[5] Thus, a little inflation may be desirable to insure against collateral deflation.

The possibility of hitting the zero lower bound on nominal interest rates, along with measurement error in inflation, are the key reasons policymakers aim for rates of inflation above zero. By addressing

[1] See Fischer (1996) for a discussion of the costs of inflation.

[2] In theory, achieving negative nominal interest rates is feasible by levying a tax on money holdings or giving up free convertibility of financial assets into cash—Buiter and Panigirtzoglou (2003), and Goodfriend (2000) discuss this idea.

[3] However, Friedman (1969) argues that the opportunity cost of holding money is zero when nominal interest rates are zero. Nominal interest rate are equal to real interest rates plus expected inflation (the Fisher identity). Since real interest rates are usually positive, inflation is expected to be negative when nominal interest rates are zero.

[4] Former Federal Reserve Governor Gramlich (2003) discusses the measurement bias in inflation.

[5] See Bernanke (2004) for a recent study of the debt-deflation problem in the Great Depression.

the zero bound problem through a low and positive inflation objective, policymakers may simultaneously insure against downward wage rigidity and debt-deflation. Central banks tend not to emphasize downward wage rigidity in their monetary policy frameworks since its relevance is not clear.[6] Although policymakers recognize the risks associated with collateral deflation, most macro models do not incorporate a debt-deflation channel.

1.2 Contribution

This paper estimates the optimal inflation rate in a small New-Keynesian sticky-price model often used for monetary policy analysis. In addition, the model accounts for an occasionally-binding zero lower bound on nominal interest rates and model uncertainty. Unlike previous studies, which impose an arbitrary inflation objective and analyze its effects, policymakers aim for the rate of inflation which maximizes the economic well-being of the public.

Besides providing estimates of the optimal inflation rate which account for the zero bound and model uncertainty, this paper unifies two separate strands of economics literature. First, studies of optimal policy and the zero bound in the New-Keynesian model, such as the studies by Eggertsson and Woodford (2003), Jung, Teranishi and Watanabe (2005), Adam and Billi (2006, 2007), do not estimate the optimal inflation rate under model uncertainty.[7] Second, as Sims (2001) points out, studies of model uncertainty and monetary policy, such as Giannoni (2002), Onatski and Stock (2002), Onatski and Williams (2003), Giordani and Söderlind (2004), Walsh (2004), Woodford (2005), Cateau (2007), Dennis (2007), Hansen and Sargent (2008), do not explain the long-run tradeoffs between the levels and variability of inflation and output.

This paper also makes a technical contribution. The numerical procedure in this paper is more efficient and more general than previous numerical procedures. The problem in this paper is more difficult to solve, compared for example to the studies by Billi (2005) and Adam and Billi (2006, 2007), because the number of state variables is unusually high for a problem with an occasionally-binding constraint on policy. The problem has five continuos state variables. Despite the use of a state-of-the-art computing environment, algorithms from previous studies are not suitable for solving the problem in this paper.[8]

Inflation is costly in the New-Keynesian framework because it distorts relative prices and leads consumers and firms to make suboptimal decisions. In absence of other factors of influence such as the

[6] Akerlof and Dickens (2007) argue whether downward wage rigidity is relevant for monetary policy.

[7] The model in this paper has some important differences relative to specifications used before. Adam and Billi (2006) solve a model that differs from the one in this paper by not including lagged inflation in the Phillips Curve. Adam and Billi (2007) analyze the case with lagged inflation but they drop the shock to the Phillips Curve. Yet lagged inflation and the shock to the Phillips Curve are both important determinants of the optimal inflation rate (Tables 4 and 6).

[8] The algorithms of Billi (2005) and Adam and Billi (2006, 2007) are based on a value function approach, which evaluates the Jacobian of the value function. The algorithm in this paper is a Euler-equation method, which bypasses the calculation of the Jacobian and is less prone to memory limitations. Another difference is that the previous algorithms employ high-order polynomials and obtain a 'smooth' approximation of the value function. Yet functions with 'kinks' are more accurately approximated by a low-order polynomial. This paper employs a piecewise-linear to approximate the response function. Indeed, any algorithm designed for solving models with smooth response functions, as for example the method of Krueger and Kubler (2004), is not designed for solving models with occasionally-binding constraints.

	Baseline	Scenario of Extreme Model Uncertainty
Optimal inflation rate:		
With no measurement bias	0.2	0.9
With 0.5% measurement bias (PCE Price Index)	0.7	1.4
Percent of time federal funds rate is zero	3.7	7.0
Number of consecutive quarters funds rate is zero	1.8	2.4
Standard deviation of:		
Output Gap	1.2	1.6
Inflation	1.9	2.9
Federal funds rate	2.4	3.2

Table 1: Optimal Inflation in the small New-Keynesian Model (Annual Percent)

zero bound, zero inflation appears optimal as it offsets these distortions. In addition to concern about inflation, policymakers in the New-Keynesian model care also about output stability. Thus, they face a dual mandate similar to the Federal Reserve. Concern about inflation and output stability and the zero bound on nominal interest rates implies that literal price stability, or zero inflation, is not the optimal goal for monetary policy.

Table 1 summarizes the results from the small New-Keynesian model. To obtain estimates of the optimal inflation rate which maximizes the economic well-being of the public, the private sector in the model perceives policy as perfectly credible. The baseline estimate is constructed as to buffer the economy from the zero bound given adverse shocks comparable in size to shocks that have hit the U.S. economy in recent decades. Assuming the model provides a correct description of the economy, the optimal rate of inflation as measured by the PCE price index (after accounting for 0.5 percentage point per year measurement bias) is 0.7 percent per year. At that rate of inflation, the federal funds rate hits the zero bound just under 4 percent of the time and stays there for only about two consecutive quarters. The standard deviation of the output gap is 1.2 percent and the standard deviation of inflation is just below 2 percent per year.

With greater uncertainty surrounding the parameters of the model, uncertainty about the actual response of the economy to shocks increases. This uncertainty about the structure of the economy leads to uncertainty about the effects of monetary policy on the economy. Thus, higher inflation is required to buffer the economy from the zero bound when there is model uncertainty.

Table 1 shows that under the scenario of extreme model uncertainty the federal funds rate reaches the zero bound 7 percent of the time and stays there for only about two consecutive quarters. In this context, extreme model uncertainty is the greatest uncertainty surrounding the parameters of the model for which inflation expectations remain anchored. Macroeconomic performance may deteriorate significantly under model uncertainty. With extreme model uncertainty, the standard deviation of output and inflation is as much as 50 percent higher than experienced in recent decades. Yet the

	Inflation Objective (PCE Price Index)				
	0	1	2	3	4
Percent of time federal funds rate is zero	14	9	5	1	< 1
Number of consecutive quarters funds rate is zero	6	5	4	3	2
Standard deviation of:					
Output Gap (CBO estimate)	3.6	3.2	3.0	2.9	2.9
Inflation	2.0	1.9	1.9	1.9	1.9
Federal funds rate	2.3	2.4	2.5	2.5	2.5

Table 2: Alternative Inflation Objectives in the FRB/US Model (Annual Percent)

optimal inflation rate rises only to 1.4 percent per year.

The results of Table 1 suggest that previous studies overstate the consequences of the zero lower bound on nominal interest rates. The range of estimates for the optimal inflation rate in the small New-Keynesian model, 0.7 to 1.4 percent per year for the PCE price index, falls below the 2 percent threshold identified in the FRB/US model and other macro models.

1.3 Previous Studies of the Zero Bound and Inflation

Previous work simulates a variety of macro models under alternative scenarios for the central bank's inflation objective in order to examine the severity of the zero bound problem. In one such study, Reifschneider and Williams (2000) simulate the FRB/US model, which is a large-scale structural model the Federal Reserve Board uses for forecasting and policy analysis.[9]

Table 2 illustrates the results from the FRB/US model. With an inflation objective of 4 percent per year, the federal funds rate reaches the zero bound less than 1 percent of the time and stays there for only two consecutive quarters. As the inflation objective is lowered and the funds rate is on average closer to zero, however, policy becomes more constrained. For a zero-inflation objective, the funds rate hits the zero bound 14 percent of the time and remains at zero for six consecutive quarters.

Table 2 shows that an increased incidence of hitting the zero bound is associated with worse macroeconomic performance. In particular, output stabilization is problematic since monetary policy is less effective at stabilizing the economy in a very low-inflation regime. If the inflation objective is lowered from 4 to 0 percent, the standard deviation of the output gap rises more than 20 percent.[10]

Estimates of the effects of the zero bound in the FRB/US model, and also other macro models, suggest that policymakers should be cautious in pursuing rates of inflation much below 2 percent per year. Such studies estimate the tradeoff between the inflation objective, the frequency of hitting the zero bound and macroeconomic performance; however, they do not provide a method to determine

[9] As most macro models, the FRB/US model does not incorporate downward wage rigidity or a debt-deflation channel.

[10] Other studies using different models and different data, but the same approach, reached very similar conclusions. Coenen, Orphanides, and Wieland (2004) were among the first to quantify the impact of the zero bound in a structural model of the U.S. economy. The model they use has about a dozen equations and thereby less sectoral detail than the FRB/US model, which comprises a few hundred equations. Despite differences in size and complexity, the Coenen, Orphanides, and Wieland model shares the same basic features of the FRB/US model.

what point along the tradeoff represents the optimal inflation rate.[11] As the inflation objective rises, the variability of both output *and inflation* falls (Table 2). As a result, higher inflation objectives appear unambiguously better.

The remainder of this paper is structured as follows. Section 2 introduces the model, then Section 3 illustrates the solution strategy and the equilibrium definition under the assumption of no model uncertainty. In Section 4, the model is calibrated to recent U.S. data. Section 5 discusses the determinants of optimal inflation for the case of no model uncertainty. Then, Section 6 quantifies optimal inflation for worst-case scenarios of model uncertainty. Section 7 briefly concludes. The Appendix describes the numerical algorithm, the derivation of the permanent consumption loss, and the calculation of the detection error probability for calibrating scenarios of model uncertainty.

2 Model

The setting adopts the well-known sticky-price version of the small New-Keynesian model, which is discussed in-depth by Clarida, Galí and Gertler (1999), Woodford (2003a), Galí (2008), and others.[12] The model consists of a representative consumer and firms in monopolistic competition facing restrictions on the frequency of price adjustments à la Calvo (1983). The policymaker commits to the objective of maximizing welfare for the consumer.[13] Thus, the optimal policy problem is

$$\max_{\{\pi_t, x_t, i_t\}} -E_0 \sum_{t=0}^{\infty} \beta^t \left[(\pi_t - \gamma \pi_{t-1})^2 + \lambda x_t^2 \right] \tag{1}$$

s.t.

$$\pi_t - \gamma \pi_{t-1} = \beta E_t (\pi_{t+1} - \gamma \pi_t) + \kappa x_t + u_t \tag{2}$$

$$x_t = E_t x_{t+1} - \varphi (i_t - E_t \pi_{t+1} - r_t^n) \tag{3}$$

$$u_t = \rho_u u_{t-1} + \sigma_{\varepsilon u} \varepsilon_{ut} \tag{4}$$

$$r_t^n = (1 - \rho_r) r_{ss} + \rho_r r_{t-1}^n + \sigma_{\varepsilon r} \varepsilon_{rt} \tag{5}$$

$$i_t \geq 0 \tag{6}$$

where E_t denotes the expectations operator conditional on all information available at time t. Expectations are rational with no uncertainty surrounding the true model of the economy. (The assumption of no model uncertainty is relaxed in Section 6.) π_t denotes the inflation rate, and x_t is the output gap or

[11] The results from the small New-Keynesian model suggest lower output and inflation variability than in the FRB/US model. However, the model-based measure of the output gap is not directly comparable to the CBO estimate of the output gap. Moreover, Reifschneider and Williams show the incorporating interest rate inertia in the policy rule would lower the variability of output and inflation in the FRB/US model.

[12] To save space, the complete derivation of the small New-Keynesian model is not shown here.

[13] The setting is a social planning problem, whereby the policymaker can implement time-zero optimal policy by selecting the equilibrium paths of inflation, the output gap, and the nominal interest rate: $\{\pi_t, x_t, i_t\}_{t=0}^{\infty}$.

the deviation of output from its flexible-price equilibrium.[14] And i_t is the nominal interest rate under the control of monetary policy.[15]

The small New-Keynesian model is developed from explicit micro-foundations. As a result, the policymaker's objective can be derived by taking a second-order Taylor series approximation to the expected life-time utility of the consumer. The resulting welfare-theoretic objective (1) is quadratic in deviations of output from the socially efficient level and deviations of the unanticipated component of inflation from zero.[16] The subjective discount factor is denoted by $\beta \in (0,1)$. The weight assigned to the goal of output stability, relative to price stability,

$$\lambda \equiv \frac{\kappa}{\theta} > 0 \tag{7}$$

is a function of the structure of the model economy. $\theta > 1$ represents the price elasticity of demand substitution among differentiated goods produced by firms operating in monopolistic competition.

Equation (2) is a log-linear approximation to the aggregate-supply relation, which describes the optimal price-setting behavior of firms under staggered price setting. The slope parameter

$$\kappa \equiv \frac{(1-\alpha)(1-\alpha\beta)}{\alpha} \frac{\varphi^{-1}+\omega}{1+\omega\theta} > 0 \tag{8}$$

depends on the structure of the model economy. $\omega > 0$ denotes the elasticity of a firm's real marginal cost with respect to its own output level. Each period, a share $\alpha \in (0,1)$ of randomly picked firms cannot adjust their prices and the remaining $(1-\alpha)$ firms get to choose prices optimally. Prices that are not optimized are indexed to the most recent aggregate price index, and $\gamma \in [0,1)$ denotes the degree of indexation. The shifter of the aggregate-supply curve, u_t, is interpreted as a 'mark-up' shock or the variation over time in the degree of monopolistic competition between firms.

Equation (3) is a log-linear approximation to the intertemporal Euler equation describing the representative consumer's private expenditure decisions. $\varphi > 0$ denotes the intertemporal elasticity of substitution or the real-rate elasticity of output. Shifting the Euler equation is the 'natural' real-rate of interest shock, r_t^n.[17]

Equations (4) and (5) describe the evolution of the exogenous mark-up shock (u_t) and the real-rate shock (r_t^n). The shocks follow AR(1) stochastic processes with autoregressive coefficients denoted by $\rho_j \in (-1,1)$ for $j = u, r$. The steady state real interest rate is $r_{ss} \equiv 1/\beta - 1$, such that $r_{ss} \in (0, +\infty)$. The innovations ($\sigma_{\varepsilon j}\varepsilon_{jt}$ for $j = u, r$) are independent across time and cross-sectionally, and normally distributed with mean zero and standard deviations denoted by $\sigma_{\varepsilon j} \geq 0$ for $j = u, r$.

Equation (6) represents the zero lower bound on the nominal interest rate. Mainly for reasons of

[14] Output is assumed to be efficient at its deterministic steady state level thanks to an output subsidy that neutralizes the distortions from monopolistic competition.

[15] By abstracting from money-demand distortions associated with positive nominal interest rates, the model can be interpreted as the 'cashless limit' of a model with money holdings.

[16] The unanticipated component of price changes ($\pi_t - \gamma\pi_{t-1}$), not the anticipated component from indexation ($\gamma\pi_{t-1}$), matters for social welfare.

[17] The real-rate shock summarizes all shocks that under flexible prices generate variation in the real interest rate; it captures the combined effects of preference shocks, productivity shocks, and exogenous changes in government expenditure.

analytical tractability, however, the economics literature often forgets the zero bound. By assuming that the policymaker can achieve negative nominal interest rates, studies based on the small New-Keynesian model typically solve the simpler problem (1)-(5) with standard linear-quadratic methods.[18] Solving the nonlinear problem (1)-(6) is far more difficult.

3 Solving the Model

The infinite-horizon Lagrangian of the time-zero optimal policy problem (1)-(6) in recursive form is

$$\begin{aligned}\max_{\{\pi_t,x_t,i_t\}}\min_{\{m_{1t},m_{2t}\}}\mathcal{L} \equiv E_0\sum_{t=0}^{\infty}\beta^t\Big\{&-(\pi_t-\gamma\pi_{t-1})^2-\lambda x_t^2 \\ &+m_{1t}\left[(1+\beta\gamma)\pi_t-\gamma\pi_{t-1}-\kappa x_t-u_t\right]-m_{1t-1}\pi_t \\ &+m_{2t}\left[-x_t-\varphi\left(i_t-r_t^n\right)\right]+m_{2t-1}\beta^{-1}\left(x_t+\varphi\pi_t\right)\Big\} \\ &\text{s.t.} \\ \text{Equations (4)-(6) for all } t\end{aligned} \tag{9}$$

where m_{1t} and m_{2t} denote the Lagrange multiplier for the aggregate-supply relation (2) and the intertemporal Euler equation (3).

The Kuhn-Tucker conditions of the Lagrangian (9) are equations (2), (3) and the marginal conditions

$$\partial\mathcal{L}/\partial\pi_t = -2\left(\pi_t-\gamma\pi_{t-1}\right)+(1+\beta\gamma)\,m_{1t}-m_{1t-1}+\beta^{-1}\varphi m_{2t-1}=0 \tag{10}$$

$$\partial\mathcal{L}/\partial x_t = -2\lambda x_t-\kappa m_{1t}-m_{2t}+\beta^{-1}m_{2t-1}=0 \tag{11}$$

$$\partial\mathcal{L}/\partial i_t\cdot i_t = -\varphi m_{2t}\cdot i_t=0,\quad m_{2t}\geq 0,\quad i_t\geq 0 \tag{12}$$

where the last equation imposes that either the Lagrange multiplier on the Euler equation ($m_{2t}\geq 0$) or the nominal interest rate ($i_t\geq 0$) must be zero, at each state and for all periods. When the nominal interest rate reaches zero ($i_t=0$), the Euler equation is binding ($m_{2t}>0$).

Equations (2), (3) and (10)-(12) form a nonlinear system of five equations with five unknowns, which must be satisfied by optimal policy in equilibrium. Solving the system delivers a five-dimensional nonlinear equilibrium response function

$$y\left(s_t\right)\equiv\left(\pi_t,x_t,i_t\geq 0,m_{1t},m_{2t}\geq 0\right)\subset R^5$$

over a five-dimensional state space

$$s_t\equiv\left(u_t,r_t^n,\pi_{t-1},m_{1t-1},m_{2t-1}\geq 0\right)\subset R^5$$

Besides the three natural state variables (u_t,r_t^n,π_{t-1}), there are two endogenous co-state variables

[18] See for example Woodford (2003a) or Galí (2008), and references therein.

given by the lagged values of the Lagrange multipliers ($m_{1t-1}, m_{2t-1} \geq 0$). The co-state variables represent 'promises' kept from past policy commitments, which lead to deviations from purely forward-looking policy whenever their value differs from zero.[19] The state in period $t+1$ depends on the state and equilibrium response in period t and the shock innovations that are unknown in period t,

$$s_{t+1} = g(s_t, y(s_t), \varepsilon_{t+1}) \tag{13}$$

Associated with the equilibrium response function, the expectations function is

$$E_t y_{t+1}(s_t) = \int y(g(s_t, y(s_t), \varepsilon_{t+1})) f(\varepsilon_{jt+1}) d(\varepsilon_{t+1}) \tag{14}$$

where $f(\cdot)$ is the probability density function of the shock innovations, $\varepsilon_t \equiv (\varepsilon_{ut}, \varepsilon_{rt}) \in R^2$.[20] The following definition of a stochastic rational expectations equilibrium is proposed.

Definition 1 (SREE) *Assume $\sigma_{\varepsilon j} \geq 0$ for $j = u, r$. A 'stochastic rational expectations equilibrium' of the optimal policy problem (1)-(6) is a nonlinear policy response function, $y(s_t)$, over the state of the model economy, s_t, with law of motion (13), such that the nonlinear system of equilibrium conditions (2), (3) and (10)-(12) is satisfied.*

Importantly, the nonlinear system in Definition 1 does not have a closed-form solution. A numerical procedure must be used to find a fixed-point in the space of nonlinear response functions. Since the number of state variables is unusually high for a model with an occasionally-binding constraint on policy, the algorithm must be highly efficient. Appendix A.1 illustrates the numerical procedure.

4 Calibration

The model is calibrated to the U.S. economy and the time period is one quarter. Table 3 summarizes the baseline parameter values, which are expressed in quarters unless otherwise noted.

The values for the main structural parameters, $(\varphi, \alpha, \theta, \omega)$ and the resulting (κ, λ), are from Tables 5.1 and 6.1 of Woodford (2003a). The degree of inflation indexation is $\gamma = 0.9$, which is consistent with the estimates of Giannoni and Woodford (2005) and Milani (2007) under the assumption of rational expectations.[21]

[19] Woodford (2003b) argues that it is desirable for policy to respond to the co-state variables: The dependence of current policy actions on past commitments allows the policymaker to steer private-sector expectations of future policy which makes policy more effective.

[20] However, when agents have 'perfect foresight' ($\sigma_{\varepsilon j} \to 0$ for $j = u, r$), the state in period $t+1$ is completely described by the state and equilibrium response in period t, $s_{t+1} = g(s_t, y(s_t))$. Since agents can anticipate future variables with certainty, the expectations function (14) is not integrated over the probability density function of the shock innovations, $E_t y_{t+1}(s_t) = y(g(s_t, y(s_t)))$.

[21] Christiano, Eichenbaum and Evans (2005) assume full inflation indexation ($\gamma = 1$) in a model that does not impose the zero lower bound on nominal interest rates. However, it is easily verified that under full indexation the optimal policy problem is not well defined because the nonlinear system in Definition 1 does not have a determinate steady state for inflation and the nominal interest rate. If there is full indexation, the change in inflation ($\pi_t - \pi_{t-1}$) matters for the welfare objective (1) and thus inflation is nonstationary.

Parameter Definition	Assigned Value
Subjective discount factor	$\beta = 0.9913$
Real-rate elasticity of output	$\varphi = 6.25$
Share of firms keeping prices fixed	$\alpha = 0.66$
Price elasticity of demand	$\theta = 7.66$
Elasticity of firms' marginal cost	$\omega = 0.47$
Slope of the Phillips curve	$\kappa = 0.024$
Weight on output in the loss function	$\lambda = 0.003$
Degree of inflation indexation	$\gamma = 0.9$
Steady state real interest rate	$r_{ss} = 3.5\%$ per year
s.d. real-rate shock innovation	$\sigma_{\varepsilon r} = 0.24\%$
s.d. mark-up shock innovation	$\sigma_{\varepsilon u} = 0.30\%$
AR(1)-coefficient of real-rate shock	$\rho_r = 0.8$
AR(1)-coefficient of mark-up shock	$\rho_u = 0$

Table 3: Baseline Calibration (Quarterly Model)

The parameters describing the two shock processes, $(r_{ss}, \sigma_{\varepsilon r}, \rho_r)$ and $(\sigma_{\varepsilon u}, \rho_u)$, are estimated over the period 1983:1-2002:4, with the same approach of Rotemberg and Woodford (1997) and Adam and Billi (2006). The expectations of inflation and the output gap are constructed from the predictions of an unconstrained VAR in inflation, the output gap, and the federal funds rate. The estimated expectations are then plugged into the intertemporal equilibrium conditions, (2) and (3), along with the actual data. The historical shock processes, u_t and r_t^n, are identified with the equation residuals. Fitting AR(1) processes to the identified historical shocks justifies the estimates for the shock processes reported in Table 3.[22] The quarterly subjective discount factor is $\beta = (1 + r_{ss})^{-1} \approx 0.9913$, as implied by the estimate for the steady state real interest rate $r_{ss} = 3.5$ percent per year.

5 Stochastic Rational Expectations Equilibrium

This section presents results for the stochastic rational expectations equilibrium under the baseline calibration in Table 3. For readability, all results are presented as annualized percentage values.

5.1 Stationary Distribution

Figure 1 shows the stationary distribution, presented in terms of probability density, for the inflation rate (top panel), the output gap (middle panel) and the nominal interest rate (bottom panel).[23] In the various panels, the dashed-vertical lines indicate the unconditional mean of the endogenous variables.

The left-hand panel ('Without Lower Bound') displays the stationary distribution for the standard

[22] Adam and Billi (2006) estimate, over the same period 1983:1−2002:4, a model with no inflation indexation ($\gamma = 0$). However, the standard deviation of the mark-up shock is almost double the size with indexation ($\gamma = 0.9$).

[23] The distribution is computed by assembling 10^5 stochastic simulations at a specific time period. The simulations are initialized to the deterministic steady state of the model. By tracking the time-evolution of the mean, standard deviation, skewness, and kurtosis, it is ascertained that the distribution did reach its stationary configuration.

linear-quadratic solution of the model, which fails to consider the zero lower bound on the nominal interest rate. As expected, the linear-quadratic solution delivers a stationary distribution that is symmetric and normally distributed around the unconditional mean.[24]

Instead, the right-hand panel ('With Lower Bound') shows the stationary distribution for the nonlinear solution of the model that does take the zero lower bound into consideration. With the lower bound, the stationary distribution remains almost normal and symmetric around the unconditional mean for inflation and the output gap; however, the stationary distribution is truncated at zero and positively skewed for the nominal interest rate.[25]

[Figure 1 about here]

Figure 2 shows the stationary distribution for an alternative scenario of far greater uncertainty about the future state of the model. The standard deviation of the shock innovations is 50 percent larger than the baseline calibration, while other parameters take the baseline values. Although the stationary distribution remains almost normal and symmetric around the unconditional mean for inflation and the output gap, the mean inflation rate and the mean nominal interest rate are higher for the model with the zero lower bound.[26]

[Figure 2 about here]

All else equal, higher inflation allows the policymaker to support a higher nominal interest rate and protects the economy against frequent episodes of zero nominal interest rates. Inflation, however, is costly to the economy. The next section investigates how optimal policy resolves the long-run tradeoff between the benefits of a more flexible stabilization policy and the cost of higher inflation.

5.2 Optimal Inflation and Welfare Cost from the Zero Bound

Table 4 reports statistics describing the stationary distribution for inflation and the nominal interest rate. The table compares results for the baseline level of uncertainty to alternative scenarios for the shock processes.

Without the zero lower bound in the model (top panel), the inflation rate and the nominal interest rate are constant at their steady state values, and the economy is half the time in deflation, for any level of uncertainty. With the zero lower bound (bottom panel), the mean inflation rate is 0.2 percent per year for the baseline, 0.3 percent per year if real-rate shocks are 50 percent larger, and 0.5 percent per year if mark-up shocks are 50 percent larger. The mean inflation rate rises further to 0.7 percent per year for the scenario of both type of shocks larger.

[24] In a linearized model the endogenous variables inherit the properties of the exogenous shock processes. Since mark-up shocks and real-rate shocks are normally distributed, the inflation rate, the output gap, and the nominal rate are also normally distributed.

[25] For the baseline, the coefficient of skewness for the nominal interest rate is 0.4 and the kurtosis is not significantly different from that of a normal distribution.

[26] For the scenario of much larger shocks, the coefficient of skewness for the nominal interest rate rises to 0.6 but the kurtosis remains not significantly different from that of a normal distribution.

	$1 \cdot \sigma_\varepsilon$	$1.5 \cdot \sigma_{\varepsilon u}$	$1.5 \cdot \sigma_{\varepsilon r}$	$1.5 \cdot \sigma_\varepsilon$
Without Lower Bound:				
$E(\pi)$	0	0	0	0
s.d.(π)	2.1	3.1	2.1	3.1
Freq($\pi < 0$)	50	50	50	50
$E(i)$	3.5	3.5	3.5	3.5
s.d.(i)	2.6	3.5	3.2	4.0
Freq($i \leq 0$)	9.2	16.0	13.7	18.8
With Lower Bound:				
$E(\pi)$	0.2	0.5	0.3	0.7
s.d.(π)	1.9	2.8	1.9	2.8
Freq($\pi < 0$)	47	45	44	41
$E(i)$	3.6	3.9	3.8	4.1
s.d.(i)	2.4	3.0	2.8	3.3
Freq($i = 0$)	3.7	8.1	6.0	8.9

Table 4: Optimal Inflation and Nominal Interest Rate (Annual Percent)

Table 4 also shows that the variability of inflation and the nominal interest rate are both lower in the model with the zero lower bound relative to the one without. The positive mean bias for inflation and the nominal interest rate stems from a long-run policy tension between the level and the variability of inflation. Higher inflation provides the monetary policymaker more room for the conduct of stabilization policy since without higher inflation the frequency of zero nominal interest rates would be higher. But higher inflation is costly to the economy and it is optimal to occasionally hit the zero bound. With the zero bound, the frequency of zero nominal interest rates is 3.7 percent for the baseline, and rises to 8.9 percent if both type of shocks are larger.

The long-run tradeoff between the benefits of a more flexible stabilization policy and the cost of higher inflation can be interpreted in terms of an 'insurance' policy. The level of inflation can be seen as the amount of insurance taken by the policymaker to avoid hitting the zero bound. At the same time, the variability of inflation is the cost of insurance. If there is greater uncertainty about the future, it is optimal for the policymaker to take more insurance but the cost of insurance is higher. So, it is not optimal to fully insure against the occurrence of zero nominal interest rates.

Table 5 reports the welfare cost under alternative scenarios for the shock processes. The welfare cost is measured in terms of the permanent reduction in consumption for the representative agent, which is derived via a transformation of the unconditional loss in the objective function (1).[27] Appendix A.2 explains the computation of the permanent consumption loss. Table 5 shows the permanent consumption loss for the model with the zero lower bound, μ, as well as the loss for the model without the lower bound, μ_{LQ}. The welfare cost from the lower bound is measured by the additional loss for the model with the lower bound relative to the one without, $\Delta(\mu) \equiv \mu - \mu_{LQ} \leq 0$.

[27] The unconditional loss is computed as the average discounted loss across $5 \cdot 10^3$ stochastic simulations, each 10^3 periods long after discarding several pre-simulated periods in order to ascertain that the distribution did reach its stationary configuration prior to the computation of the loss.

		$1 \cdot \sigma_\varepsilon$	$1.5 \cdot \sigma_{\varepsilon u}$	$1 \cdot \sigma_{\varepsilon r}$	$1.5 \cdot \sigma_\varepsilon$
+	Consumption Loss With Lower Bound: μ	-0.28	-0.63	-0.29	-0.64
−	Consumption Loss Without Lower Bound: μ_{LQ}	-0.28	-0.62	-0.28	-0.62
=	Additional Loss With Lower Bound: $\Delta(\mu)$	-0.00	-0.01	-0.01	-0.02

Table 5: Welfare Cost from the Zero Lower Bound (Annual Percent)

The welfare cost from the zero lower bound increases if either type of shocks is larger. However, the welfare cost is not very large under optimal policy.[28] Even allowing for real-rate shocks and mark-up shocks 50 percent larger than the baseline, the welfare cost from the zero lower bound amounts to a permanent reduction in consumption of less than 0.02 percent per year.

5.3 Robustness of Results to Extreme Calibrations

Table 6 shows the results for a wide range of changes to each structural parameter of the model. As is customary in the calibration literature, only one parameter is changed for each scenario.

Parameter changes can affect the results via two main channels in the model. The parameter change modifies the equilibrium conditions describing the behavior of the economy and alters the 'slope' (8) of the aggregate-supply relation. The parameter change also modifies the welfare 'weight' (7) assigned to output stability relative to price stability.

The real-rate elasticity of output ($\varphi > 0$) determines the leverage of nominal-interest-rate policy on consumption. Higher elasticity implies stronger leverage and thereby lowers the frequency of zero nominal interest rates. With higher elasticity, the slope of the aggregate-supply relation falls which makes it optimal to tolerate more variability of inflation. Yet the weight on output stability also falls which implies that lower variability of inflation is desirable instead. Depending on which of these two opposing effects dominates, the variability of inflation can turn out either higher or lower. Thus, it is not obvious whether the optimal inflation rate should rise or fall. With a very high real-rate elasticity of output—φ equal to 10—the optimal inflation rate rises to 0.3 percent per year.

The optimal inflation rate goes down for both very low and very high degrees of price stickiness ($0 < \alpha < 1$). If prices are more flexible, the variability of inflation is lower and the zero lower bound is reached less often, so there is less incentive for policy to support high inflation. Conversely, if prices are more sticky, the variability of inflation is higher and the zero bound is reached more often, so high inflation is more costly. The results show that for both very low and very high degrees of price stickiness—α equal to 0.1 and 0.9—the optimal inflation rate falls to 0.1 percent per year.

The price elasticity of demand substitution among differentiated goods produced by firms ($\theta > 1$) determines the degree of monopolistic competition. The less competition among firms, the stronger the incentive for policy to support high inflation, which restrains the frequency of zero nominal interest rates. With very low competition—θ equal to 3—the optimal inflation rate rises to 0.5 percent per year.

[28] The welfare cost of 'suboptimally' high inflation is expected to be greater than the cost under optimal policy. Lucas (2000) surveys research on the welfare cost of suboptimal inflation in a monetary economy.

Alternative Calibrations	$E(\pi)$	s.d.(π)	Fr($\pi < 0$)	$E(i)$	s.d.(i)	Fr($i = 0$)	$\Delta(\mu)$
Baseline	0.2	1.9	47	3.6	2.4	3.7	−0.00
Very low elasticity of output ($\varphi = 1$)	0.0	1.5	50	3.5	3.1	11.6	−0.01
Very high elasticity of output ($\varphi = 10$)	0.3	2.0	45	3.8	2.3	2.9	−0.00
Very flexible prices ($\alpha = 0.1$)	0.1	0.3	43	3.6	1.6	0.0	−0.00
More flexible prices ($\alpha = 0.3$)	0.2	0.8	43	3.6	1.7	0.1	−0.00
More sticky prices ($\alpha = 0.5$)	0.2	1.4	47	3.6	2.1	1.7	−0.00
Very sticky prices ($\alpha = 0.9$)	0.1	2.6	49	3.5	2.9	9.7	−0.04
Very low competition ($\theta = 3$)	0.5	2.0	41	4.0	2.3	1.5	−0.00
Very high competition ($\theta = 15$)	0.0	1.9	50	3.5	2.6	5.1	−0.02
Very inelastic marginal cost ($\omega = 0.1$)	0.2	1.9	47	3.6	2.4	3.6	−0.00
Very elastic marginal cost ($\omega = 10$)	0.2	2.0	47	3.6	2.4	3.7	−0.06
No inflation indexation ($\gamma = 0$)	0.0	0.9	50	3.5	1.6	1.3	−0.00
Less inflation indexation ($\gamma = 0.85$)	0.1	1.7	48	3.6	2.2	3.0	−0.00
More inflation indexation ($\gamma = 0.95$)	0.4	2.4	46	3.8	2.8	5.3	−0.01
Almost full indexation ($\gamma = 0.99$)	0.7	3.0	44	4.2	3.3	6.8	−0.01
Very low steady state rate ($r_{ss} = 2\%$)	0.6	1.8	38	2.5	2.1	9.2	−0.01
Very high steady state rate ($r_{ss} = 5\%$)	0.1	2.0	49	5.0	2.5	0.4	−0.00

Table 6: Robustness of Results to Extreme Calibrations (Annual Percent)

The elasticity of a firm's real marginal cost with respect to its own output level ($\omega > 0$) has little effect on the results. Less elastic marginal cost produces stronger incentives for policy to avoid zero nominal interest rates. However, such incentives are not strong enough to generate any noticeable change in the optimal inflation rate.

The ability of the policymaker to steer private-sector expectations depends on the degree of inflation indexation ($0 \leq \gamma < 1$). In particular, the expectations channel of monetary policy is less effective when private-sector expectations are less forward-looking. If there is more inflation indexation, and thus private-sector expectations are less forward-looking, the optimal inflation rate and the frequency of zero nominal interest rates are higher. With almost full inflation indexation—γ equal to 0.99—the optimal inflation rate rises to 0.7 percent per year.[29]

Lower steady state real interest rates ($0 < r_{ss} < +\infty$) imply lower steady state nominal interest rates, and thereby the zero lower bound is reached more often. When interest rates are lower, the policymaker has a stronger incentive to support higher inflation. With a very low steady state real interest rate—r_{ss} equal to 2 percent per year—the optimal inflation rate rises to 0.6 percent per year.

The findings overall are robust to a wide range of changes to each structural parameter. Although changing one parameter at a time helps understand the role of a particular dimension of the model in generating the results, there may be uncertainty along multiple dimensions of the model. As a result, changing only one parameter may understate the extent of the model misspecification.

The researcher may assess the extent of the model misspecification by changing all the parameters

[29] For the reasons explained earlier in footnote 21, the scenario of full indexation is not well defined because the inflation rate would be nonstationary.

of the model at once, each in the direction that gives rise to the worst scenario. Yet, it is not clear whether the worst scenario is towards the boundary of the joint parameter space. In fact, as can be seen in Table 6, the worst scenario for the optimal inflation rate occurs for intermediate values of the degrees of price stickiness ($0 < \alpha < 1$) rather than at the boundary.

6 Model Uncertainty

This section discusses the findings for worst-case scenarios of model uncertainty. Following the approach of Hansen and Sargent (2008), the optimal policy problem is generalized to a robust control problem.

6.1 Robust Control Model

In the robust control approach, the model is viewed as an approximation of the true model of the economy which is surrounded by unmeasurable uncertainty. The true model is unknown to the policymaker, but is known to be in a neighborhood around its approximating model. The robust control problem can be viewed as a game played between the policymaker who aims to maximize its objective function, and a fictitious adversary agent (or nature) who at the same time seeks to minimize the objective function of the policymaker.

Thus, the robust control version of the optimal policy problem is

$$\max_{\{\pi_t, x_t, i_t\}} \min_{\{w_{1t}, w_{2t}\}} -\hat{E}_0 \sum_{t=0}^{\infty} \beta^t \left[(\pi_t - \gamma\pi_{t-1})^2 + \lambda x_t^2 - \Theta \left(w_{1t}^2 + w_{2t}^2 \right) \right] \tag{15}$$

s.t.

$$\pi_t - \gamma\pi_{t-1} = \beta \hat{E}_t (\pi_{t+1} - \gamma\pi_t) + \kappa x_t + u_t \tag{16}$$

$$x_t = \hat{E}_t x_{t+1} - \varphi \left(i_t - \hat{E}_t \pi_{t+1} - r_t^n \right) \tag{17}$$

$$u_t = \rho_u u_{t-1} + \sigma_{\varepsilon u} (\varepsilon_{ut} + w_{1t}) \tag{18}$$

$$r_t^n = (1 - \rho_r) r_{ss} + \rho_r r_{t-1}^n + \sigma_{\varepsilon r} (\varepsilon_{rt} + w_{2t}) \tag{19}$$

$$i_t \geq 0 \tag{20}$$

where $\hat{E}_t$ denotes the expectations operator conditional on information available at time t. The accent is added above the expectations operator to indicate that expectations are formed under the worst-case scenario of model uncertainty. As in the standard rational expectations framework, the policymaker and the private sector share the same model of the economy because the policymaker aims to maximize welfare for the representative consumer.

However, the policymaker faces model uncertainty and the adversary agent is able to hamper the policymaker's ability to shape private-sector expectations. The adversary agent seeks to minimize the objective function (15) of the policymaker by choosing the variables w_{1t} and w_{2t} and manipulating the

evolution of the exogenous shock processes (18) and (19). The variables w_{1t} and w_{2t} are interpreted in the literature as 'wort-case shocks' and represent unmeasurable uncertainty surrounding the true model of the economy. By manipulating the (expected) evolution of the economy, the adversary agent at the same time can directly affect the formation of expectations.

The parameter $\Theta \geq 0$ in the objective function (15) determines the degree of misspecification or the distance between the approximating model and the worst-case scenario. Intuitively, when Θ is small a more severe modeling misspecification can arise. When $\Theta \to +\infty$, however, there is no model uncertainty and the robust control problem simplifies to the standard rational expectations case, already analyzed in the previous sections of this paper.

6.2 Solving the Robust Control Model

The infinite-horizon Lagrangian of the time-zero optimal policy problem (15)-(20) is

$$\begin{aligned}
\max_{\{\pi_t, x_t, i_t\}} \min_{\{m_{1t}, m_{2t}, w_{1t}, w_{2t}\}} \hat{\mathcal{L}} \equiv \hat{E}_0 \sum_{t=0}^{\infty} \beta^t \Big\{ & -(\pi_t - \gamma\pi_{t-1})^2 - \lambda x_t^2 + \Theta\left(w_{1t}^2 + w_{2t}^2\right) \\
& + m_{1t}\left[(1+\beta\gamma)\pi_t - \gamma\pi_{t-1} - \kappa x_t - u_t\right] - m_{1t-1}\pi_t \\
& + m_{2t}\left[-x_t - \varphi\left(i_t - r_t^n\right)\right] + m_{2t-1}\beta^{-1}\left(x_t + \varphi\pi_t\right)\Big\} \\
& \text{s.t.} \\
& \text{Equations (18)-(20) for all } t
\end{aligned} \tag{21}$$

where m_{1t} and m_{2t} denote the Lagrange multipliers for the equilibrium conditions (16) and (17).

The Kuhn-Tucker conditions are equations (16), (17) and the marginal conditions

$$\partial\hat{\mathcal{L}}/\partial\pi_t = -2(\pi_t - \gamma\pi_{t-1}) + (1+\beta\gamma)m_{1t} - m_{1t-1} + \beta^{-1}\varphi m_{2t-1} = 0 \tag{22}$$

$$\partial\hat{\mathcal{L}}/\partial x_t = -2\lambda x_t - \kappa m_{1t} - m_{2t} + \beta^{-1}m_{2t-1} = 0 \tag{23}$$

$$\partial\hat{\mathcal{L}}/\partial i_t \cdot i_t = -\varphi m_{2t} \cdot i_t = 0, \quad m_{2t} \geq 0, \quad i_t \geq 0 \tag{24}$$

$$\partial\hat{\mathcal{L}}/\partial w_{1t} = 2\Theta w_{1t} - \sigma_{\varepsilon u} m_{1t} = 0 \tag{25}$$

$$\partial\hat{\mathcal{L}}/\partial w_{2t} = 2\Theta w_{2t} + \sigma_{\varepsilon r}\varphi m_{2t} = 0 \tag{26}$$

where equation (24) imposes that either the Lagrange multiplier on the Euler equation ($m_{2t} \geq 0$) or the nominal interest rate ($i_t \geq 0$) is zero, at all states and for each period. The adversary agent chooses the variables w_{1t} and $w_{2t} \leq 0$ so to constrain, through the marginal conditions (25) and (26), the policymaker's ability to make commitments through the endogenous co-state variables m_{1t} and $m_{2t} \geq 0$.[30] As a result of such constraints on the use of policy commitments, model uncertainty can

[30] Equation (26), $m_{2t} \geq 0$, $\sigma_{\varepsilon r} \geq 0$, $\varphi > 0$ and $\Theta \geq 0$ implies $w_{2t} \leq 0$.

make policy less effective.

Solving the nonlinear system of equilibrium conditions (16), (17) and (22)-(26) delivers a seven-dimensional nonlinear equilibrium response function

$$\hat{y}\left(s_t\right) \equiv \left(\pi_t, x_t, i_t \geq 0, m_{1t}, m_{2t} \geq 0, w_{1t}, w_{2t} \leq 0\right) \subset R^7$$

where the variables w_{1t} and w_{2t} represent the policymaker's uncertainty surrounding the true model of the economy. The equilibrium response function has a five-dimensional state space

$$s_t \equiv \left(u_t, r_t^n, \pi_{t-1}, m_{1t-1}, m_{2t-1} \geq 0\right) \subset R^5$$

The state in period $t+1$ depends on the state and equilibrium response in period t and the shock innovations that are unknown in period t,

$$s_{t+1} = g(s_t, \hat{y}\left(s_t\right), \varepsilon_{t+1}) \tag{27}$$

Associated with the equilibrium response function, the expectations function is

$$\hat{E}_t \hat{y}_{t+1} = \int \hat{y}\left(g(s_t, \hat{y}\left(s_t\right), \varepsilon_{t+1})\right) f\left(\varepsilon_{jt+1}\right) d\left(\varepsilon_{t+1}\right) \tag{28}$$

where $f\left(\cdot\right)$ is the probability density function of the shock innovations, $\varepsilon_t \equiv \left(\varepsilon_{ut}, \varepsilon_{rt}\right) \in R^2$.[31] Since the variables w_{1t} and w_{2t} enter directly the equilibrium expectations function, model uncertainty hampers the policymaker's ability to shape expectations. The following definition of a stochastic robust control equilibrium is advanced.

Definition 2 (SRCE) *Assume $\sigma_{\varepsilon j} \geq 0$ for $j = u, r$ and $\Theta \geq 0$. A 'stochastic robust control equilibrium' of the optimal policy problem (15)-(20) is a nonlinear policy response function, $\hat{y}\left(s_t\right)$, over the state space, s_t, with law of motion (27), such that the nonlinear system of equilibrium conditions (16), (17) and (22)-(26) is satisfied.*

When $\Theta \to +\infty$, however, the robust control problem simplifies to the standard rational expectations case. With no model uncertainty, a stochastic robust control equilibrium (SRCE) coincides in the limit with a standard stochastic rational expectations equilibrium (SREE). This observation justifies the following corollary.

Corollary 3 *Assume $\sigma_{\varepsilon j} \geq 0$ for $j = u, r$ and $\Theta \geq 0$. Given $\hat{y}\left(s_t\right)$ solving a SRCE and $y\left(s_t\right)$ solving a SREE, $\hat{y}\left(s_t\right) \to y\left(s_t\right)$ for $\Theta \to +\infty$.*

[31] As in the standard rational expectations framework, the expectations function is not integrated over the probability density function of shock innovations when agents have perfect foresight for the reasons explained in footnote 20.

Proof. When $\Theta \to +\infty$ equations (25) and (26) imply $w_{1t} = w_{2t} = 0$. Thus, the system of equilibrium conditions (16), (17) and (22)-(26) in Definition 2 (SRCE) coincides in the limit with the system of equilibrium conditions (2), (3) and (10)-(12) in Definition 1 (SREE). ■

6.3 Stochastic Robust Control Equilibrium

This section presents results for the stochastic robust control equilibrium under the baseline calibration in Table 3. For readability, all results are presented as annualized percentage values.

With respect to the standard rational expectations case, there is one additional parameter that must be calibrated, $\Theta \geq 0$, which determines the degree of model uncertainty. Hansen and Sargent (2008) propose a statistical theory of model selection to determine a context-specific value for Θ. They propose choosing a reasonable probability of making a detection error, $p(\Theta)$, about whether observed equilibrium outcomes may have originated from the approximating model with or without a worst-case shock.

When $\Theta \to +\infty$, and there is no model uncertainty, the probability of making a detection error is 50 percent because in the limit there is no difference between the approximating model with or without a worst-case shock. When Θ is smaller, however, the model misspecification is more severe and more easily detected. Appendix A.3 explains the computation of the detection error probability.

6.3.1 Stationary Distribution

Figure 3 shows the stationary distribution for the worst-case equilibrium with an extreme degree of model uncertainty, which corresponds to the lowest detection error probability, 29 percent, for which the algorithm can identify an equilibrium with the zero lower bound in the model.[32] The stationary distribution is almost normal and symmetric around the unconditional mean for inflation and the output gap. However, the mean inflation rate and the mean nominal interest rate are higher for the model with the zero lower bound.[33]

Figure 4 shows the stationary distribution for the approximating equilibrium without a worst-case shock. Since the approximating equilibrium is an 'intermediate' case between the worst-case equilibrium and the standard rational expectations equilibrium, as expected the mean inflation rate and the mean nominal interest rate are not as high as under the worst-case equilibrium of the model with the zero lower bound (Figure 3).

6.3.2 Robustly Optimal Inflation and Welfare Cost from the Zero Bound

Figure 5 depicts average inflation (top panel), the standard deviation of inflation (middle panel) and the frequency of deflation (bottom panel) over the range of feasible detection error probabilities for the

[32] The detection error probability is computed by averaging across 10^4 stochastic simulations. The sample size for the simulations is set to 80 periods, which is equal to number of observations over the estimation period 1983:1−2002:4 used to calibrate the quarterly model in Section 4.

[33] For the worst-case equilibrium with extreme model uncertainty, the coefficient of skewness for the nominal interest rate is 0.6 and the kurtosis is not significantly different from that of a normal distribution.

model with the zero lower bound. The left-hand panel shows the results for the worst-case equilibrium—for completeness the right-hand panel shows the results for the approximating equilibrium. The optimal inflation rate after accounting for the zero lower bound is increasing in a nonlinear fashion with the degree of model uncertainty.

Although the detection error probability ranges 29 to 50 percent for the model with the zero lower bound, a detection error probability even as low as 10 percent—not shown in the figures—is feasible if the zero lower bound is ignored. As a result, models without the zero lower bound will overstate the effectiveness of monetary policy and exaggerate the feasible degree of model uncertainty.

Figure 5 shows that the optimal inflation rate for the worst-case equilibrium rises from 0.2 percent per year with no model uncertainty to 0.5 percent per year when the detection error probability is 30 percent. The optimal inflation rate then jumps to 0.9 percent per year when the detection error probability is 29 percent. Thus, results from the robust control model encompass all the scenarios for the optimal inflation rate under the extreme calibrations (Table 6). The standard deviation of inflation rises from 1.9 per year with no model uncertainty to 2.9 per year with extreme model uncertainty, and the frequency of deflation falls from 47 to 40 percent.

The model provides estimates of the optimal inflation rate based on a hypothetical measure of inflation with no measurement error. As a result, the model-based estimates are independent of any specific measure of inflation (PCE price index, or others). As explained in section 1, however, available measures of inflation tend to be biased upward. Thus, to convert the model-based estimate of inflation into an actual measure of inflation, an estimate of the bias has to be added.

Figure 6 compares inflation correctly measured (left-hand panel) to inflation in terms of the PCE price index (right-hand panel), for which the bias is roughly 0.5 percentage point per year, under the worst-case equilibrium. The optimal rate of inflation as measured by the PCE price index ranges 0.7 to 1.4 percent per year depending on the degree of model uncertainty. At the same time, the frequency of deflation for the PCE price index ranges 33 to 37 percent. The frequency of deflation for the PCE price index is lower, than for a hypothetical measure of inflation with no measurement bias, because the stationary distribution for inflation is shifted by the estimate of the bias.

Figure 7 depicts the average nominal interest rate (top panel), the standard deviation of the nominal interest rate (middle panel) and the frequency of zero nominal interest rates (bottom panel). With more severe model uncertainty, higher nominal interest rates are required to protect the economy against frequent episodes of zero nominal interest rates. The optimal nominal interest rate for the worst-case equilibrium rises from 3.6 percent per year with no model uncertainty to 4.3 percent per year with an extreme degree of model uncertainty. The standard deviation of the nominal interest rate rises from 2.4 to 3.2 percent per year, and the frequency of zero nominal interest rates rises from 3.7 to 7.0 percent.

Figure 8 shows the welfare cost for the worst-case equilibrium (top panel) and the approximating equilibrium (bottom panel). The welfare cost is measured in terms of the permanent reduction in consumption for the representative agent. The welfare cost from the zero lower bound is given by the additional loss for the model with the zero lower bound (line with circles) relative to the one without

(line with squares). More severe model uncertainty produces a larger welfare cost. Yet the welfare cost from the zero lower bound—the distance between the two lines—amounts to a permanent reduction in consumption of less than 0.02 percent per year.

6.3.3 Worst-Case Shocks

Without the zero lower bound, the worst-case shocks are linear in the state because the model is linear-quadratic. When the detection error probability is 29 percent, the worst-case shock to the Phillips Curve is $w_{1t} = 0.34u_t + 0r_t^n + 0\pi_{t-1} + 0.17m_{1t-1} + 0m_{2t-1}$. Yet the worst-case shock to the Euler Equation is irrelevant, $w_{2t} = \underline{0} \cdot s_t$, because the Euler Equation is redundant to the solution of the model. With the zero bound, however, the worst-case shocks do not have a closed-form solution.

Figure 9 characterizes the worst-case shocks to the Phillips Curve (left-hand panel) and the Euler Equation (right-hand panel). The variability of the worst-case shocks (top panel) explains the model uncertainty. The correlation of the worst-case shocks to either the mark-up shock (middle panel) or the real-rate shock (bottom panel) is positive, which means that the worst-case shocks are 'piled onto' the exogenous shock processes. The worst-case shock to the Phillips Curve piles onto the mark-up shock (left-hand, middle panel) but not the real-rate shock (left-hand, bottom panel). The worst-case shock to the Euler Equation piles onto the mark-up shock (right-hand, middle panel) and the real-rate shock (right-hand, bottom panel).

7 Conclusions

When inflation is low, nominal interest rates may approach zero which limits a central bank's ability to stabilize the economy by lowering its policy rate. Researchers who have analyzed this issue have estimated a tradeoff between the central bank's inflation objective and the likelihood of hitting the zero lower bound on nominal interest rates. Previous studies show that the incidence of hitting the zero lower bound falls quickly as the inflation objective rises from 0 to roughly 4 percent per year. While this line of research provides information policymakers can use in formulating an inflation objective, it does not provide them a direct estimate of the optimal inflation rate.

An estimate of the optimal inflation rate, however, can be obtained by simulating a small New-Keynesian model. In the model described in this paper, estimates of the optimal inflation rate which accounts for the zero bound range 0.7 to 1.4 percent per year, for the PCE price index, depending on the degree of model uncertainty. This range of estimates is somewhat lower than that suggested by previous researchers.

In interpreting these results, a number of caveats must be kept in mind. First, the model focuses on the effects of the zero nominal interest rate bound and ignores other potential factors that might lead policymakers to pursue positive rates of inflation. These factors include downward wage rigidity and the potential cost of debt-deflation. Second, the results are derived from a very simple model that abstracts from many real world features. Finally, the estimates of measurement error in the various

price indexes are themselves subject to error, creating some uncertainty around the optimal rate of inflation as measured using available price indexes.

Even with these caveats, the results suggest that the zero nominal interest rate bound may not warrant quite the concern that some economists and policymakers have attributed to it. Still, further research is needed to confirm or refine these results in models that incorporate a more realistic and complete description of the economy.

A Appendix

A.1 Numerical Procedure

Solving the robust control model requires finding a response function which satisfies the high-dimensional nonlinear system of equilibrium conditions in Definition 2.

The state space, $s \subset R^5$, is discretized into a set of N interpolation nodes, $\{s_n | n = 1, ..., N\}$ where $s_n \in s$. The response function, $\hat{y}(s) \subset R^7$, is evaluated at intermediate values of the discretization grid by resorting to multilinear interpolation. Next period's state depends on the current state and response and the shock innovations that are unknown in the current period: $s_{+1} = g(s, \hat{y}(s), \varepsilon_{+1})$. Since the shock innovations, $\varepsilon \in R^2$, are normally distributed, the expectations function, $\hat{E}\hat{y}_{+1}(s)$, is evaluated accurately and efficiently with an M-node Gaussian-Hermite quadrature scheme, $\{\varepsilon_m | m = 1, ..., M\}$ where $\varepsilon_m \in \varepsilon$, as explained in Chapter 7 of Judd (1998).[34] Quadrature-based integration is accurate only if the integrands to be evaluated are smooth. The integrands are smooth because the underlying distributions of inflation and the output gap are smooth, which the top panels (inflation) and middle panels (output gap) of Figures 1 to 4 show.

The fixed-point of the nonlinear system in Definition 2 is found with an iterative update rule

$$\hat{y}^{k+1} \leftarrow \hat{y}^k + \iota^k \left(\hat{y}^{k+1} - \hat{y}^k \right), \text{ from step } k \text{ to } k+1 \tag{29}$$

where $\iota^k \in (0, 1]$ is the step size chosen to guarantee algorithm stability and convergence, as explained in Chapter 4 of Bertsekas (1999).

The **Algorithm** proceeds as follows:

Step 1: Assign the interpolation nodes with an efficient, sparse-grid method. Guess an initial value for the response function, $\hat{y}^0$.

Step 2: Update the state, evaluate the expectations function and apply the iterative update rule (29) to derive a new guess for the response function, $\hat{y}^{+1}$.

[34] However, when solving for the perfect-foresight limit ($\sigma_\varepsilon \to 0$), next period's state depends only on the current state and response, $s_{+1} = g(s, y(s))$, and the expectations function is $\hat{E}\hat{y}_{+1}(s) = \hat{y}(g(s, y(s)))$. The stationary distribution in the limit collapses to a mass point at the deterministic steady state of the model. Since a quadrature scheme is not necessary, the numerical procedure is less involved.

Step 3: Stop if $\max_{n=1,\ldots,N} \left\| \hat{y}^{k+1} - \hat{y}^{k} \right\| < \tau$, where $\tau > 0$ is the convergence tolerance level. Otherwise repeat step 2.

The convergence tolerance level is set to the square root of machine precision, $\tau = 1.49 \cdot 10^{-8}$. The accuracy of the solution is checked by computing the Euler equation residuals at an arbitrary set of R residual interpolation nodes, $\{s_r | r = 1, ..., R\}$ where $s_r \in s$, as explained by Santos (2000). To assure that the approximation error is acceptably small and does not affect the results, the Euler equation residuals are verified at the interpolation nodes and over a finer grid.

To cope with the curse of dimensionality, the procedure employs a sparse grid which assigns more weight to the regions of the state space where the occasionally-binding constraint is active. The support for the exogenous shocks is chosen to cover ± 4 unconditional standard deviations, which is more than sufficient for a mark-up shock or a real-rate shock to drive the nominal interest rate to its zero lower bound. The support for the endogenous state variables is chosen large enough to avoid erroneous extrapolation. To achieve for the baseline an acceptable degree of approximation with a sparse grid requires $N \approx 3.6 \cdot 10^4$ interpolation nodes and $M \approx 45$ quadrature nodes, as opposed to a linearly spaced grid which would require over 10^6 interpolation nodes.

To achieve greater efficiency, the procedure employs an approximation refinement method. The initial guess for the response function is set to the linearized solution around the deterministic steady state on a coarse grid, $N^0 \times M^0$. The nonlinear solution obtained for the coarse grid is interpolated over a finer grid and used as a new guess to resolve. The degree of approximation is progressively increased towards the final set of nodes, $\{N^0 \times M^0 < N^1 \times M^1 < \ldots < N \times M\}$. Experimentation with alternative initial guesses did not lead to differences in the results.

A.2 Permanent Consumption Loss

The expected utility of the representative household, as shown in Chapter 6 of Woodford (2003a), can be validly approximated up to second order by

$$\hat{E}_0 \sum_{t=0}^{\infty} \beta^t U_t = \frac{U_c \overline{Y}}{2} \frac{\alpha\theta\left(1+\omega\theta\right)}{\left(1-\alpha\right)\left(1-\alpha\beta\right)} \hat{L} \tag{30}$$

where $\overline{Y}$ denotes steady state output, $U_c > 0$ is the steady state marginal utility of consumption and

$$\hat{L} \equiv -\hat{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\left(\pi_t - \gamma\pi_{t-1}\right)^2 + \lambda x_t^2 - \Theta\left(w_{1t}^2 + w_{2t}^2\right) \right] \leq 0$$

is the objective function (15) of the policymaker.

At the same time, the utility loss generated by a permanent reduction in consumption of $\mu \leq 0$ can be validly approximated up to second order by

$$\frac{1}{1-\beta}\left(U_c \overline{Y} \mu + \frac{U_{cc}}{2} \left(\overline{Y}\mu\right)^2 \right) = \frac{U_c \overline{Y}}{1-\beta}\left(\mu - \frac{1}{2\varphi}\mu^2 \right) \tag{31}$$

where $U_{cc} < 0$ is the second derivative of utility with respect to consumption evaluated at the steady state and $\varphi \equiv -U_c / \left(\overline{Y} U_{cc}\right) > 0$ is the intertemporal elasticity of substitution of aggregate expenditure.

Equating the right-hand sides of (30) and (31) delivers

$$\frac{1}{2\varphi}\mu^2 - \mu + \frac{1-\beta}{2}\frac{\alpha\theta\left(1+\omega\theta\right)}{\left(1-\alpha\right)\left(1-\alpha\beta\right)}\hat{L} = 0$$

Therefore, the loss in terms of permanent consumption is

$$\mu = \varphi\left(1 - \sqrt{1 - \frac{1-\beta}{\varphi}\frac{\alpha\theta\left(1+\omega\theta\right)}{\left(1-\alpha\right)\left(1-\alpha\beta\right)}\hat{L}}\right) \tag{32}$$

A.3 Detection Error Probability

To compute the detection error probability, as explained in Chapter 9 of Hansen and Sargent (2008), the researcher should simulate the model for a small sample. For a long enough sample the misspecification will be easy to detect. In the limit, for $T \to +\infty$, the probability of a detection error is 50 percent.

The researcher can estimate the log-likelihood ratio that the data was generated by the approximating model without a worst-case shock

$$r_A = \frac{1}{T}\sum_{t=0}^{T-1}\left[\frac{1}{2}w_{At}'w_{At} - w_{At}'\varepsilon_{At}\right]$$

where w_{At} is a vector of control variables set by the adversary agent so to distort the dynamics of the exogenous shock processes (18) and (19). ε_{At} is a vector of innovations normally distributed and independent across time and cross-sectionally. The paths for w_{At} are obtained, however, by performing stochastic simulations with the dynamics of the shocks following the undistorted processes (4) and (5).

The researcher can also estimate the log-likelihood ratio that the data was generated by the approximating model with a worst-case shock

$$r_B = \frac{1}{T}\sum_{t=0}^{T-1}\left[\frac{1}{2}w_{Bt}'w_{Bt} + w_{Bt}'\varepsilon_{Bt}\right]$$

where w_{Bt} is a vector of control variables distorting the dynamics of the shock processes (18) and (19). ε_{Bt} is a vector of innovations assumed normally distributed and independent across time and cross-sectionally. The paths for w_{Bt} are stochastically simulated with the shocks following the distorted processes (18) and (19).

Therefore, by assigning equal prior weights to the approximating model with or without a worst-case shock, the overall detection error probability is

$$p\left(\Theta\right) = \frac{1}{2}\left(p_A + p_B\right)$$

where $p_j = \text{Freq}(r_j \leq 0)$, for $j = A, B$.

References

Adam, Klaus and Roberto M. Billi, "Optimal Monetary Policy under Commitment with a Zero Bound on Nominal Interest Rates," *Journal of Money, Credit, and Banking*, 2006, *38* (7), 1877–1905.

____ and ____ , "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," *Journal of Monetary Economics*, 2007, *54* (3), 728–752.

Akerlof, George A. and William T. Dickens, "Unfinished Business in the Macroeconomics of Low Inflation: A Tribute to George and Bill by Bill and George," *Brookings Papers on Economic Activity*, 2007, *2*, 31–47.

Bernanke, Ben S., *Essays on the Great Depression*, Princeton: Princeton University Press, 2004.

____ , Vincent R. Reinhart, and Brian P. Sack, "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment," *Brookings Papers on Economic Activity*, 2004, *2*, 1–100.

Bertsekas, Dimitri P., *Nonlinear Programming*, 2 ed., Belmont, Massachusetts: Athena Scientific, 1999.

Billi, Roberto M., "The Optimal Inflation Buffer with a Zero Bound on Nominal Interest Rates," *Center for Financial Studies, Frankfurt, Germany, Working Paper 17*, 2005.

Buiter, Willem H. and Nikolaos Panigirtzoglou, "Overcoming the Zero Bound on Nominal Interest Rates with Negative Interest on Currency: Gesell's Solution," *Economic Journal*, 2003, *113*, 723–746.

Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 1983, *12*, 383–398.

Cateau, Gino, "Monetary Policy under Model and Data-Parameter Uncertainty," *Journal of Monetary Economics*, 2007, *54* (7), 2083–2101.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Clarida, Richard, Jordi Galí, and Mark Gertler, "The Science of Monetary Policy: Evidence and Some Theory," *Journal of Economic Literature*, 1999, *37*, 1661–1707.

Coenen, Günter, Athanasios Orphanides, and Volker Wieland, "Price Stability and Monetary Policy Effectiveness When Nominal Interest Rates are Bounded at Zero," *Advances in Macroeconomics*, 2004, *4* (1), Article 1.

Dennis, Richard, "Robust Control with Commitment: A Modification to Hansen-Sargent," *Journal of Economic Dynamics and Control*, 2007. Forthcoming.

Eggertsson, Gauti B. and Michael Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 2003, *1*, 139–233.

Fischer, Stanley, "Why Are Central Banks Pursuing Long-Run Price Stability?," in "Achieving Price Stability," Federal Reserve Bank of Kansas City, 1996, pp. 7–34.

Fisher, Irving, "The Debt-Deflation Theory of Great Depressions," *Econometrica*, 1933, *1* (4), 337–357.

Friedman, Milton, "The Optimum Quantity of Money," in "The Optimum Quantity of Money and Other Essays," Chicago: Aldine Publishing Company, 1969, chapter 1, pp. 1–50.

Fuhrer, Jeffrey C. and Brian F. Madigan, "Monetary Policy When Interest Rates are Bounded at Zero," *Review of Economics and Statistics*, 1997, *79*, 573–585.

Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton: Princeton University Press, 2008.

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

—— **and Michael Woodford**, "Optimal Inflation-Targeting Rules," in Ben S. Bernanke and Michael Woodford, eds., *The Inflation-Targeting Debate*, Chicago: University of Chicago Press, 2005, chapter 3, pp. 93–162.

Giordani, Paolo and Paul Söderlind, "Solution of Macromodels with Hansen-Sargent Robust Policies: Some Extensions," *Journal of Economic Dynamics and Control*, 2004, *28*, 2367–2397.

Goodfriend, Marvin, "Overcoming the Zero Bound on Interest Rate Policy," *Journal of Money, Credit, and Banking*, 2000, *32* (4), 1007–1035.

Gramlich, Edward M., "Maintaining Price Stability," *Remarks before the Economic Club of Toronto*, Canada, October 1, 2003.

Hansen, Lars P. and Thomas J. Sargent, *Robustness*, Princeton: Princeton University Press, 2008.

Judd, Kenneth L., *Numerical Methods in Economics*, Cambridge, Massachusetts: MIT Press, 1998.

Jung, Taehun, Yuki Teranishi, and Tsutomu Watanabe, "Optimal Monetary Policy at the Zero-Interest-Rate-Bound," *Journal of Money, Credit, and Banking*, 2005, *37* (5), 813–835.

Kato, Ryo and Shin-Ichi Nishiyama, "Optimal Monetary Policy When Interest Rates are Bounded at Zero," *Journal of Economic Dynamics and Control*, 2005, *29*, 97–133.

Krueger, Dirk and Felix Kubler, "Computing Equilibrium in OLG Models with Stochastic Production," *Journal of Economic Dynamics and Control*, 2004, *28* (7), 1411–1436.

Krugman, Paul R., "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 1998, *2*, 137–205.

Lucas, Robert E., "Inflation and Welfare," *Econometrica*, 2000, *68* (2), 247–274.

McCallum, Bennett T., "Theoretical Analysis Regarding a Zero Lower Bound on Nominal Interest Rates," *Journal of Money, Credit, and Banking*, 2000, *32* (4), 870–904.

Milani, Fabio, "Expectations, Learning and Macroeconomic Persistence," *Journal of Monetary Economics*, 2007, *54* (7), 2065–2082.

Okun, Arthur M., *Prices and Quantities: A Macroeconomic Analysis*, Washington, D.C.: The Brookings Institution, 1981.

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

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

Orphanides, Athanasios and Volker Wieland, "Efficient Monetary Policy Design Near Price Stability," *Journal of the Japanese and International Economies*, 2000, *14*, 327–365.

Phelps, Edmund S., *Inflation Policy and Unemployment Theory*, New York: W. W. Norton and Company, 1972.

Reifschneider, David and John C. Williams, "Three Lessons for Monetary Policy in a Low-Inflation Era," *Journal of Money, Credit, and Banking*, 2000, *32* (4), 936–966.

Rotemberg, Julio J. and Michael Woodford, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 1997, *12*, 297–346.

Santos, Manuel S., "Accuracy of Numerical Solutions Using the Euler Equation Residuals," *Econometrica*, 2000, *68* (6), 1377–1402.

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

Summers, Lawrence, "Panel Discussion: Price Stability: How Should Long-Term Monetary Policy Be Determined?," *Journal of Money Credit and Banking*, 1991, *23* (3), 625–631.

Svensson, Lars E. O., "Escaping from a Liquidity Trap and Deflation: The Foolproof Way and Others," *Journal of Economic Perspectives*, 2003, *17* (4), 145–166.

Tobin, James, "Inflation and Unemployment," *American Economic Review*, 1972, *62* (1/2), 1–18.

Vickrey, William S., "Stability Through Inflation," in Kenneth K. Kurihara, ed., *Post-Keynesian Economics*, New Brunswick, New Jersey: Rutgers University Press, 1954, chapter 4, pp. 89–122.

Walsh, Carl E., "Robustly Optimal Instrument Rules and Robust Control: An Equivalence Result," *Journal of Money, Credit, and Banking*, 2004, *36* (6), 1105–1113.

Wolman, Alexander L., "Real Implications of the Zero Bound on Nominal Interest Rates," *Journal of Money, Credit, and Banking*, 2005, *37* (2), 273–296.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.

——, "Optimal Interest-Rate Smoothing," *Review of Economic Studies*, 2003, *70* (245), 861–886.

——, "Robustly Optimal Monetary Policy with Near-Rational Expectations," *NBER Working Paper No. 11896*, 2005.

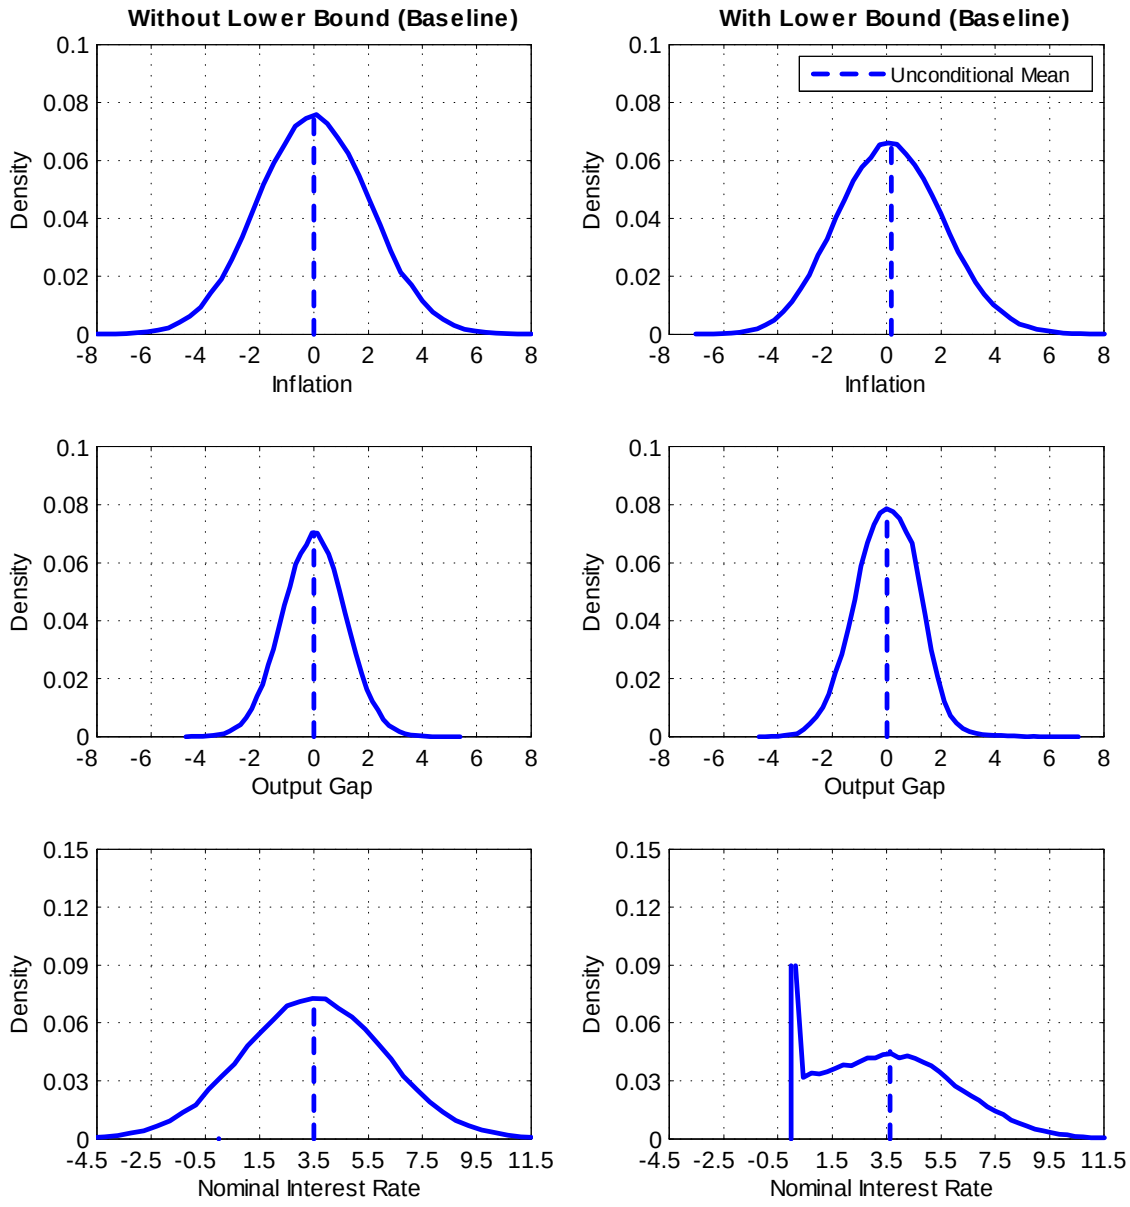

Figure 1: Stationary Distribution of Inflation, Output Gap and Nominal Interest Rate under the Baseline Calibration (Annual Percent)

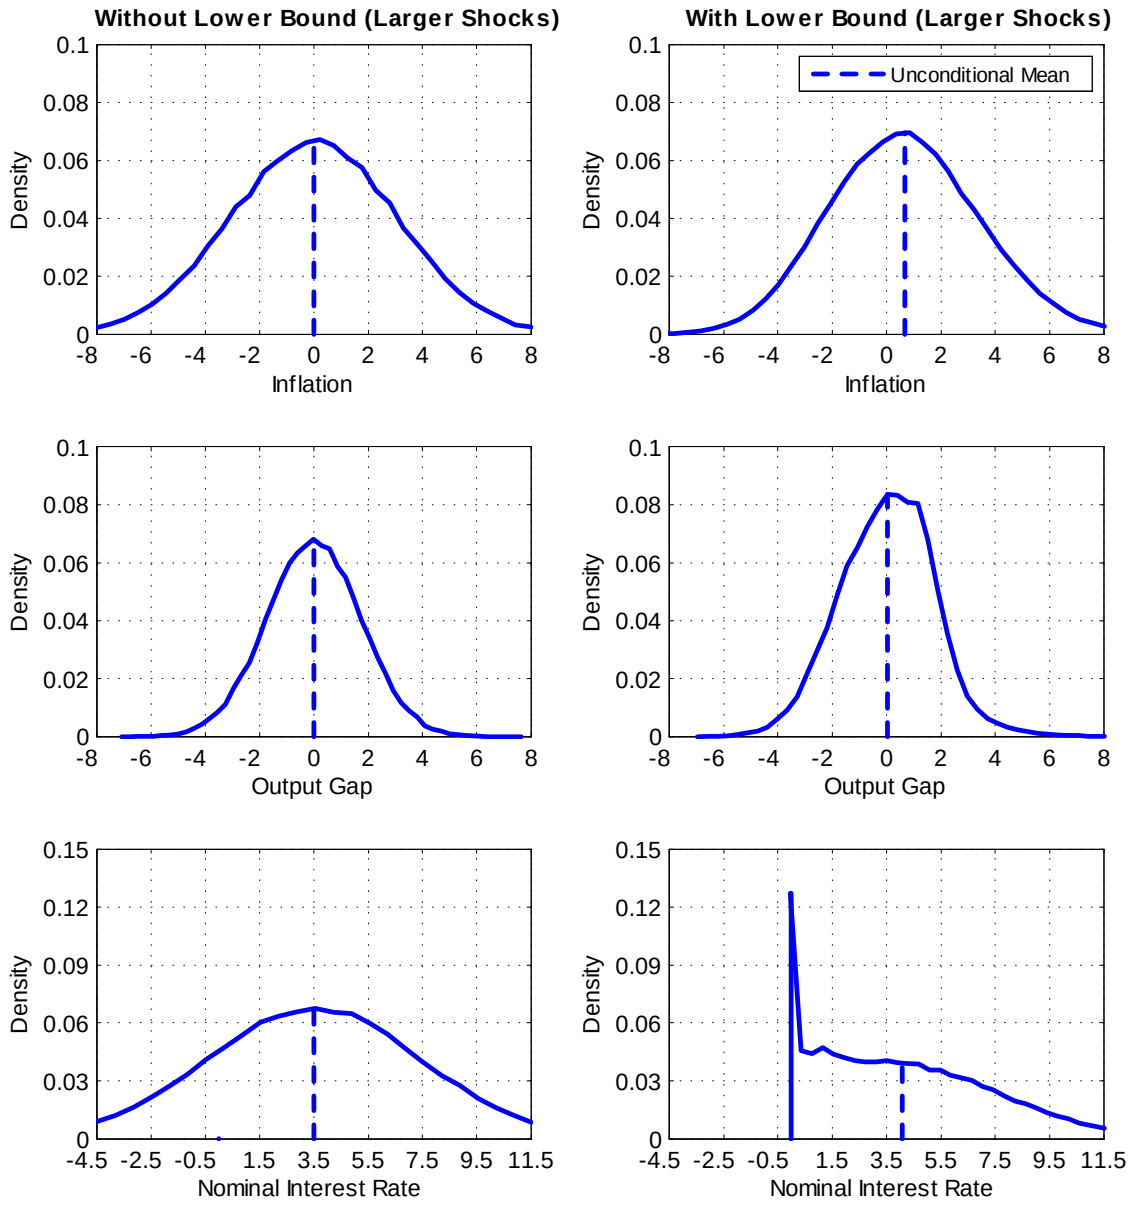

Figure 2: Stationary Distribution for Extremely Uncertain Economy: Standard Deviation of Innovations 50% Larger than Baseline (Annual Percent)

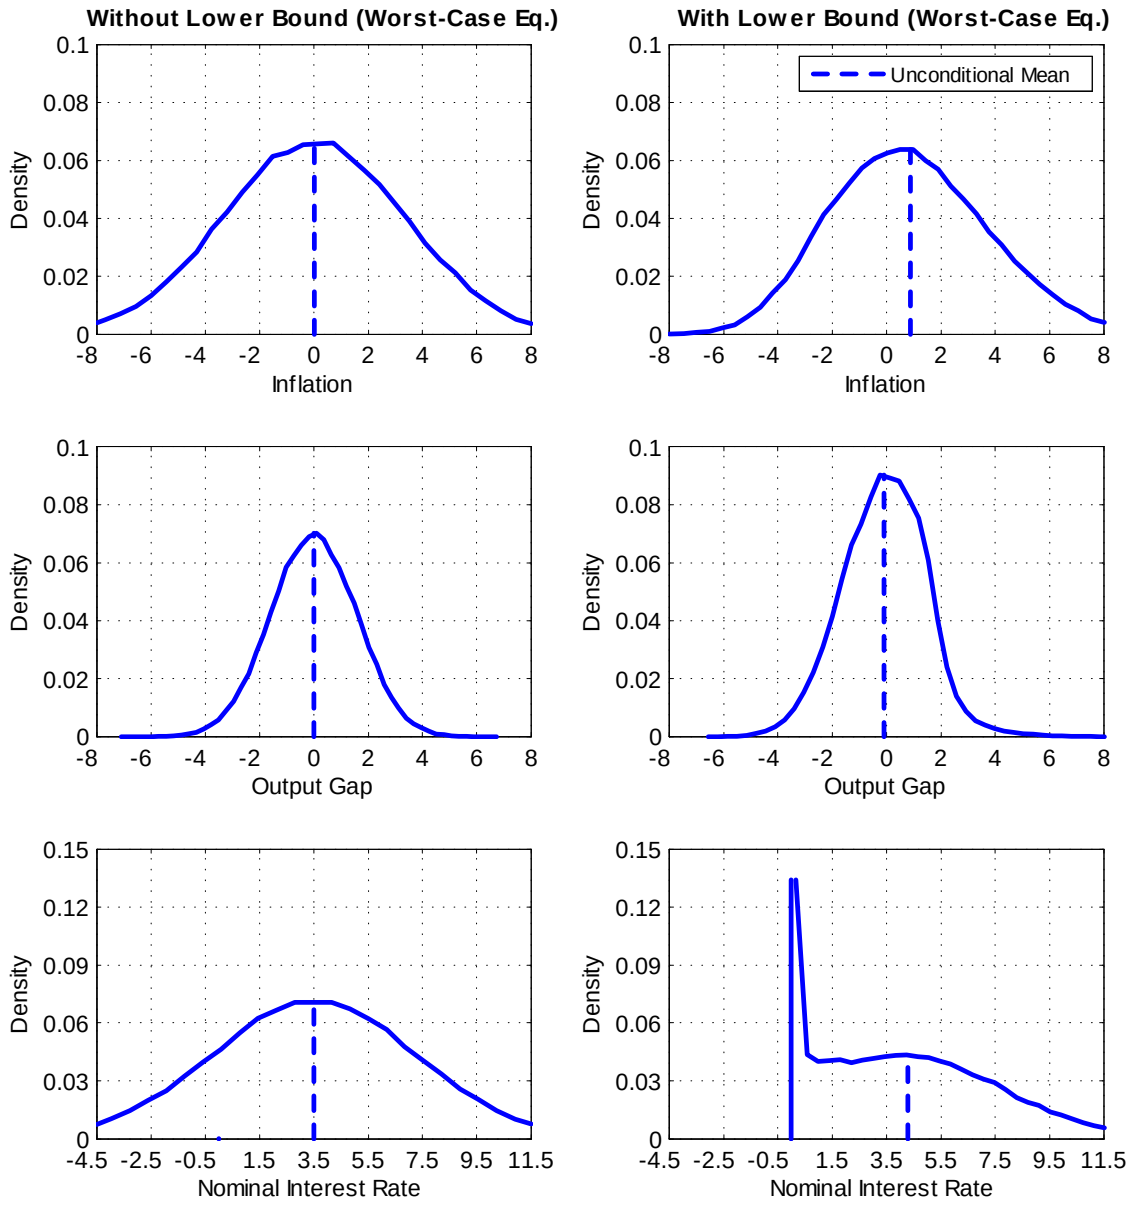

Figure 3: Stationary Distribution for Worst-Case Equilibrium with Extreme Model Uncertainty: Lowest Detection Error Probability with Lower Bound (Annual Percent)

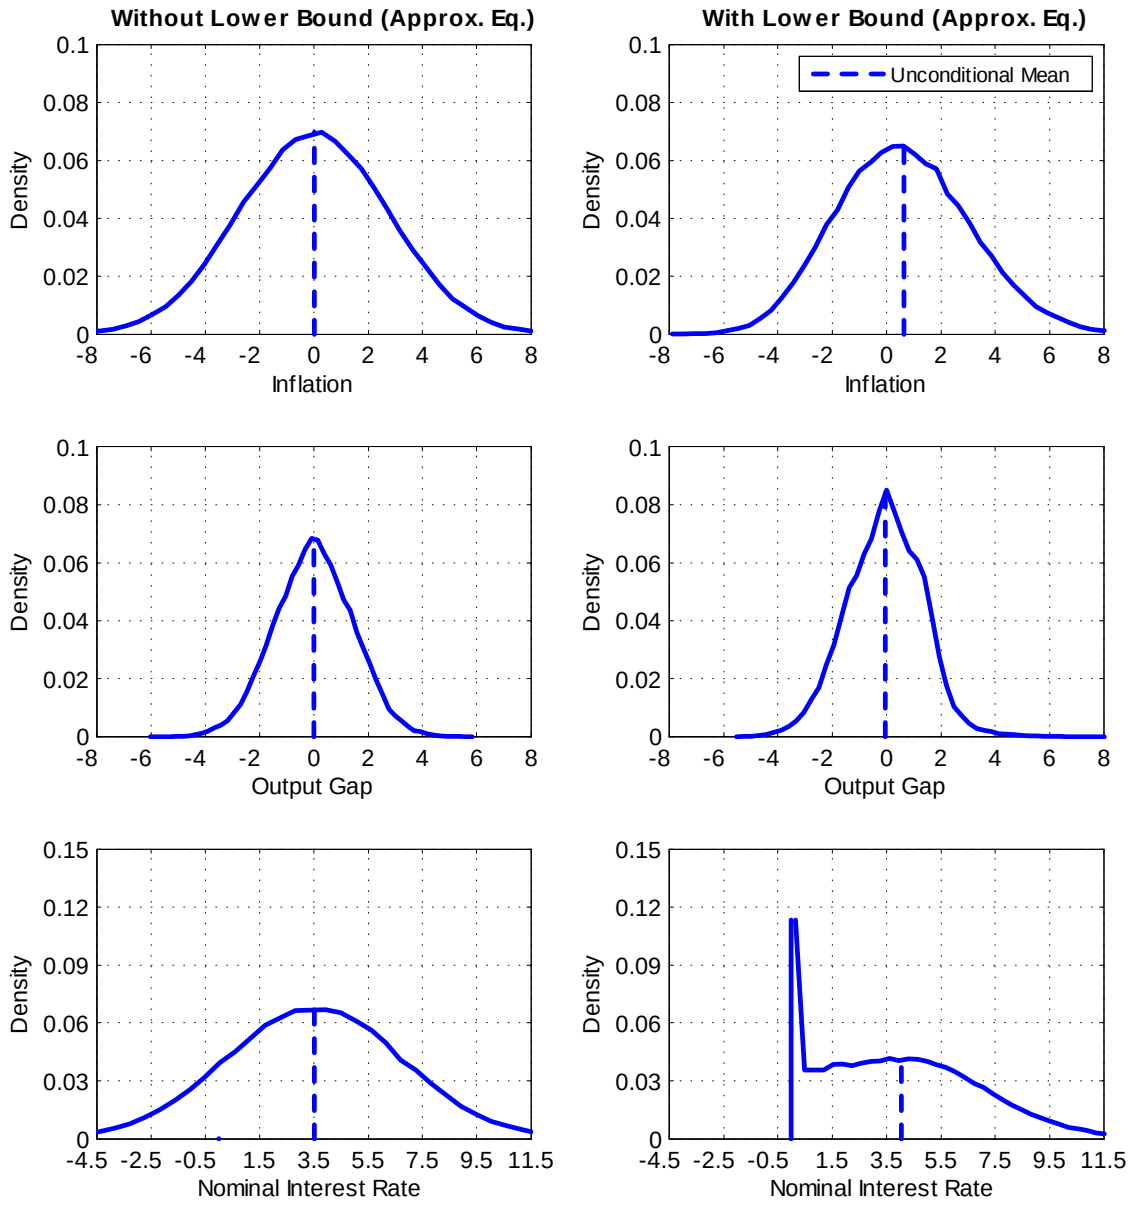

Figure 4: Stationary Distribution for Approximating Equilibrium with Extreme Model Uncertainty: Lowest Detection Error Probability with Lower Bound (Annual Percent)

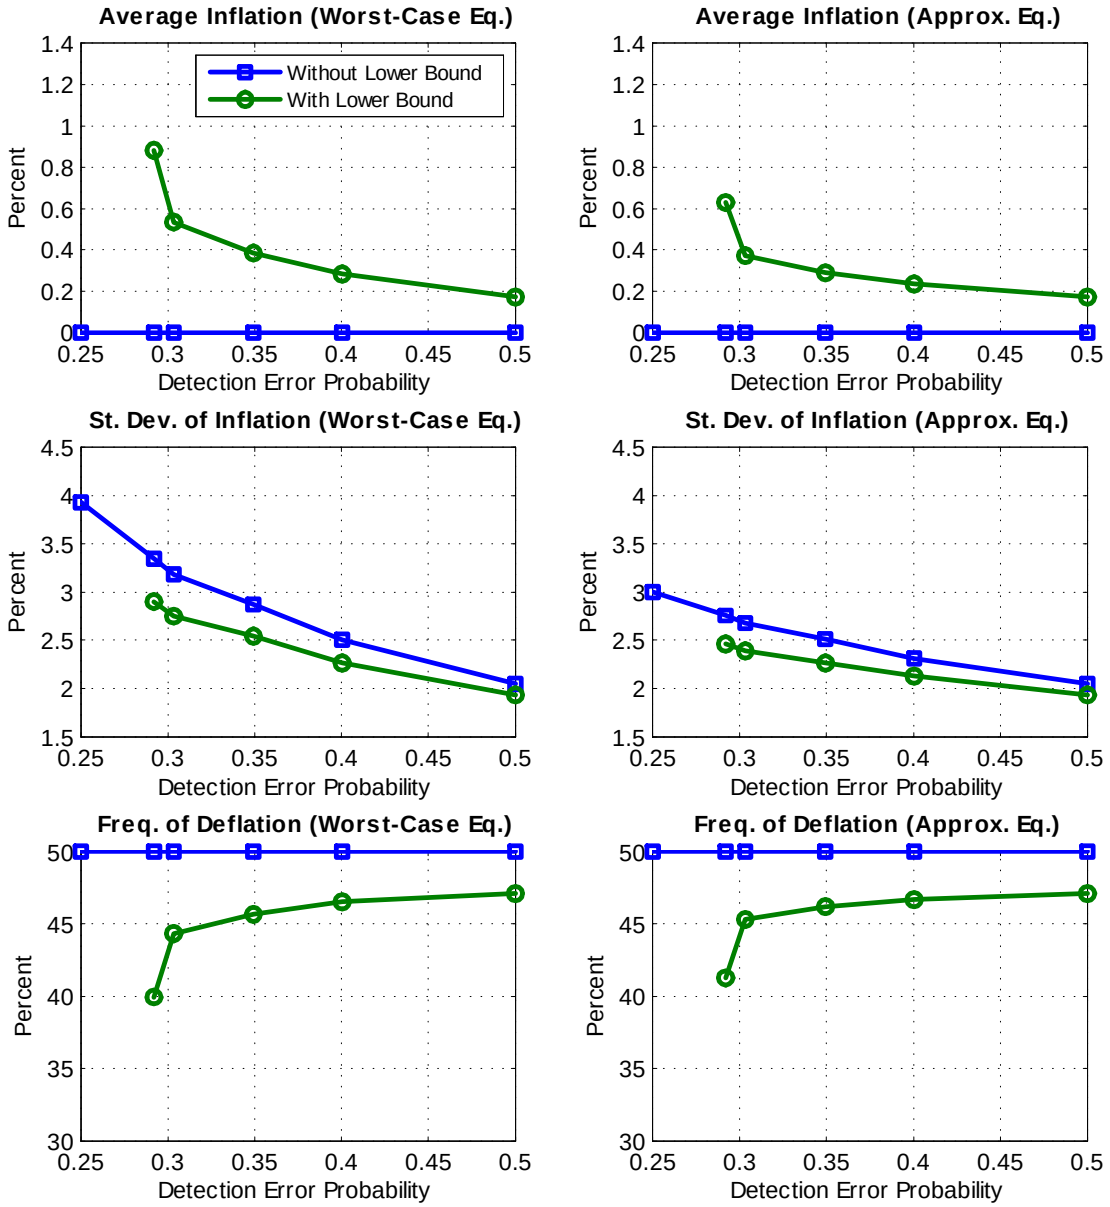

Figure 5: Optimal Inflation and Model Uncertainty (Annual Percent)

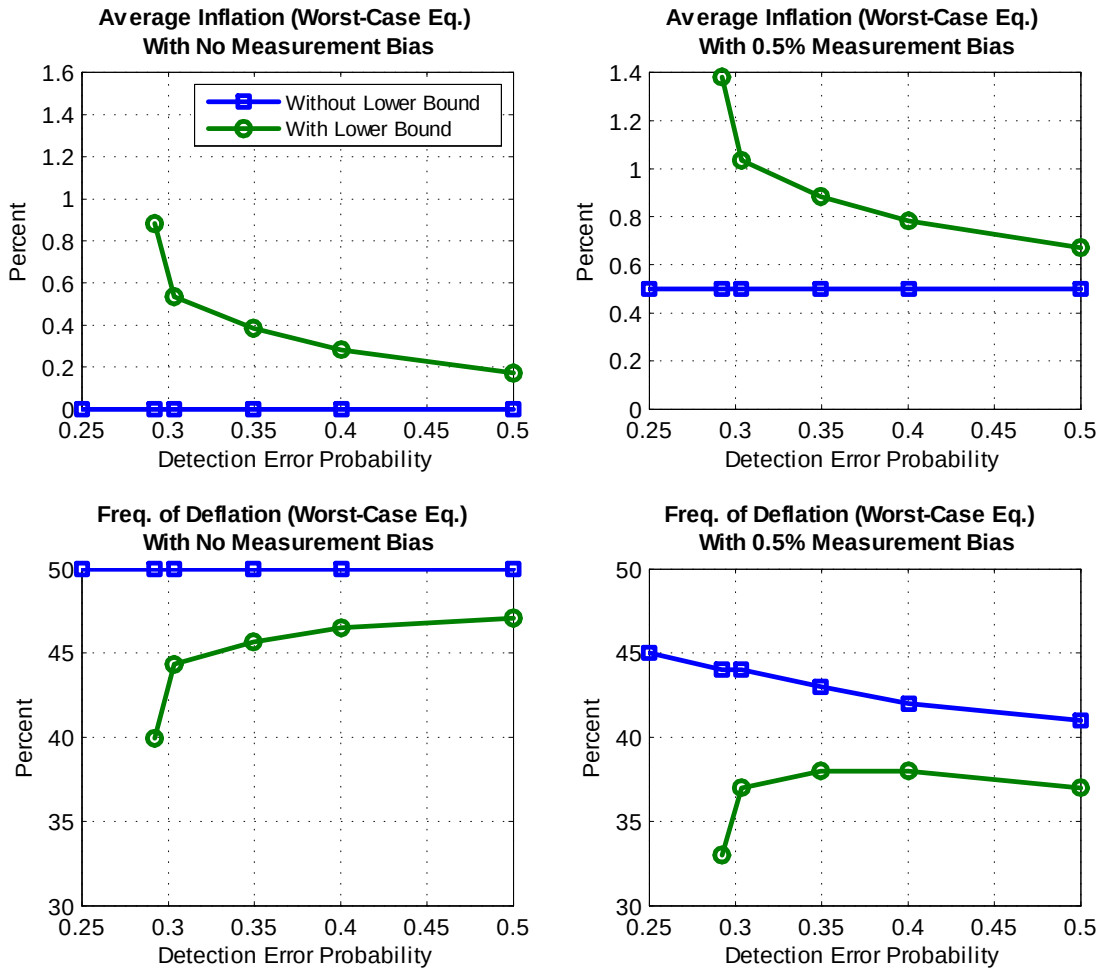

Figure 6: Optimal Inflation, Measurement Bias and Model Uncertainty (Annual Percent)

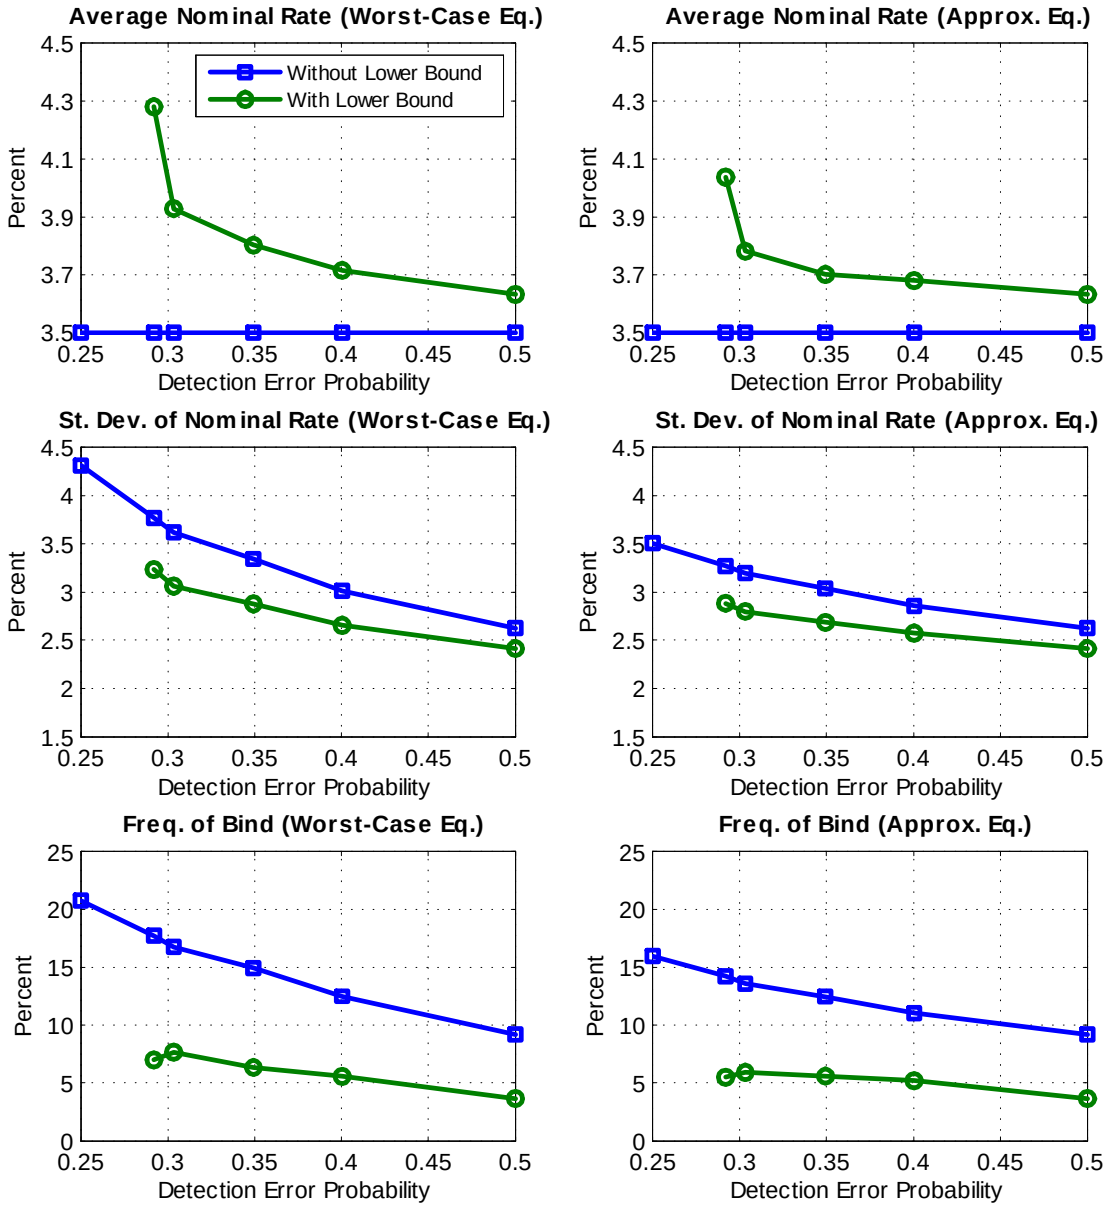

Figure 7: Optimal Nominal Interest Rate and Model Uncertainty (Annual Percent)

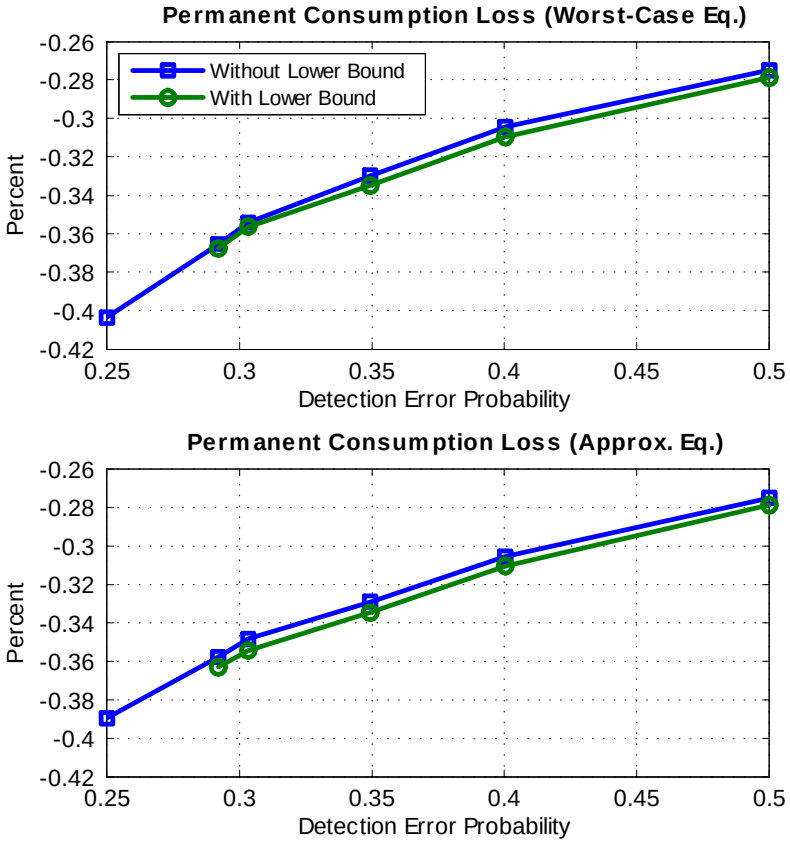

Figure 8: Welfare Cost from the Lower Bound and Model Uncertainty (Annual Percent)

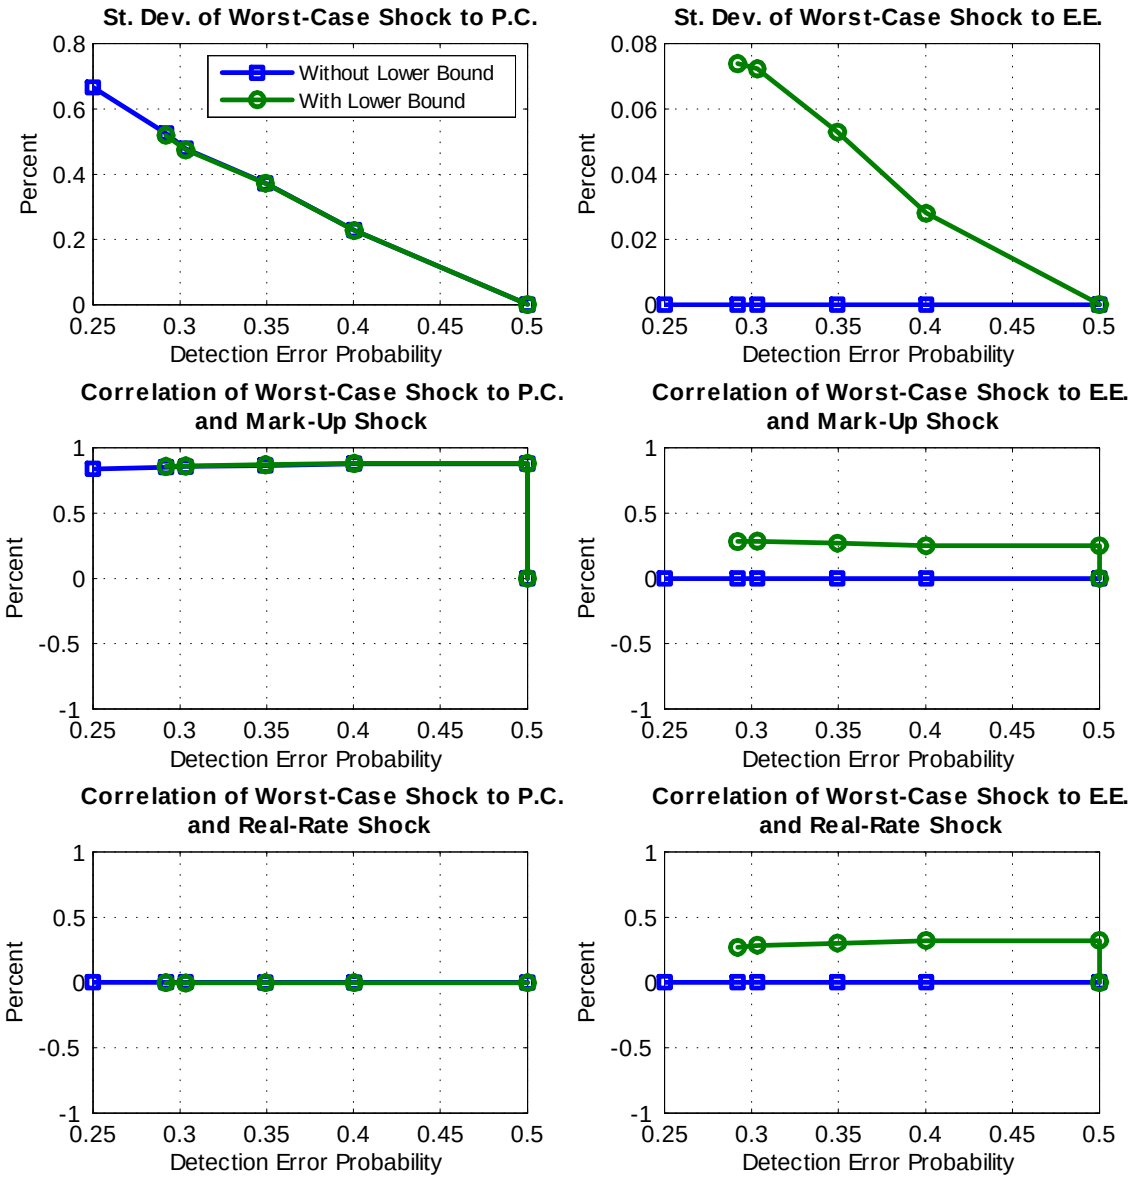

Figure 9: Worst-Case Shocks and Model Uncertainty (Annual Percent)