ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Mertens, Thomas M.; Williams, John C.

Working Paper

What to expect from the lower bound on interest rates: Evidence from derivatives prices

Staff Report, No. 865

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Mertens, Thomas M.; Williams, John C. (2018) : What to expect from the lower bound on interest rates: Evidence from derivatives prices, Staff Report, No. 865, Federal Reserve Bank of New York, New York, NY

This Version is available at:
https://hdl.handle.net/10419/210717

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Federal Reserve Bank of New York
Staff Reports

What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices

Thomas M. Mertens
John C. Williams

Staff Report No. 865
August 2018

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices

Thomas M. Mertens and John C. Williams

Federal Reserve Bank of New York Staff Reports, no. 865
August 2018
JEL classification: E43, E52, G12

Abstract

This paper analyzes the effects of the lower bound for interest rates on the distributions of expectations for future inflation and interest rates. We study a stylized New Keynesian model where the policy instrument is subject to a lower bound to motivate the empirical analysis. Two equilibria emerge: In the "target equilibrium," policy is unconstrained most or all of the time, whereas in the "liquidity trap equilibrium," policy is mostly or always constrained. We use options data on future interest rates and inflation to study whether the decrease in the natural rate of interest leads to forecast densities consistent with the theoretical model. We develop a lower bound indicator that captures the effects of the lower bound on the distribution of interest rates. Qualitatively, we find that the evidence is largely consistent with the theoretical predictions in the target equilibrium and find no evidence in favor of the liquidity trap equilibrium. Quantitatively, while the lower bound has a sizable effect on the distribution of future interest rates, its impact on forecast densities for inflation is relatively modest.

Key words: zero lower bound, inflation expectations, monetary policy, multiple equilibria

---

Williams: Federal Reserve Bank of New York (email: john.c.williams@ny.frb.org). Mertens: Federal Reserve Bank of San Francisco (email: thomas.mertens@sf.frb.org). This is an updated version of a working paper originally issued by the Federal Reserve Bank of San Francisco (Working Paper 2018-03) in January 2018. The authors thank Adrien Auclert, Michael Bauer, Ben Bernanke, Mick Devereux, Stefano Eusepi, Stefan Gerlach, Simon Gilchrist, Yuriy Gorodnichenko, Kevin Lansing, Glenn Rudebusch, Stephanie Schmitt-Grohé, Elmar Mertens, Marek Raczko, Eric Swanson, Andrea Tambalotti, John Taylor, Pablo Winant, and Jonathan Wright as well as the audiences at UBC, the Swiss National Bank's Research Conference 2017, German Economists Abroad 2017, the 2018 AEA meetings, the meeting of the Federal Reserve System Committee on Macroeconomics, the Federal Reserve Board-New York Fed conference "Developments in Empirical Macroeconomics," 2018 NASMES, NBER Summer Institute, the Hoover MMCN conference, the CEBRA conference, and EFA for helpful discussions and suggestions. They also thank Patrick Shultz for outstanding research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve Bank of San Francisco, or the Federal Reserve System.

To view the authors' disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr865.html.

1 Introduction

The lower bound on nominal interest rates has been the subject of extensive study in the academic literature and a key factor in central bank practice over the past two decades. Standard macroeconomic models predict that the lower bound can have profound effects on the behavior of the economy and supply a set of testable empirical predictions. However, the relatively small number of observations during which the lower bound has been binding limits the ability to quantitatively assess these predictions using macroeconomic data.

This paper makes two key contributions to the literature. First, it links the higher-moment predictions of macroeconomic theory to prices of financial market derivatives related to options on interest rates and inflation. Second, it derives and tests hypotheses that distinguish between multiple equilibria in an economy where interest rates are constrained by a lower bound. It uses options data from U.S. financial markets to measure the effects of the lower bound on expectations and thereby the macroeconomy. We compare the forecast densities of future nominal interest rates and inflation rates derived from a theoretical model to those observed in financial markets based on derivatives data. The advantage of this approach is that, unlike macroeconomic data that are buffeted by realizations of shocks, far-ahead expectations should reflect the underlying fundamentals of the economy. Our empirical strategy takes advantage of the significant decline in estimates of the natural rate of interest over the recent past to identify the effects of the lower bound on forecast densities, i.e., distributions of beliefs.

In our theoretical model, inflation is determined by a fundamental shock, expectations about future inflation, as well as the level of the nominal interest rate set by the central bank. Although our model is very simple, the main mechanisms and implications related to the lower bound are common to many more complicated macroeconomic models used in the literature.[1] The central bank optimally sets the interest rate to stabilize the inflation rate and output under discretion. The lower bound on interest rates limits the ability to optimally respond to adverse shocks. Expectations of future inflation also depend on the likelihood with which the lower bound will bind in the future. As a result, a nonlinear feedback between future occurrences of policy being constrained by the lower bound and current inflation and output emerges.

In the deterministic version of the model, two steady-state rates arise, consistent with the findings of Benhabib, Schmitt-Grohé and Uribe (2001). In one equilibrium, which we refer to as the "target equilibrium," the nominal interest rate is strictly above the lower bound and the inflation rate equals the target rate set by

[1] See, for example, Fuhrer and Madigan (1997), Reifschneider and Williams (2000), Eggertsson and Woodford (2003), Evans, Fisher, Gourio and Kane (2015), Reifschneider (2016), and Hamilton, Harris, Hatzius and West (2016).

the central bank. In the second, which we refer to as the "liquidity trap equilibrium," the nominal interest rate is constrained by the lower bound and the inflation rate equals the lower bound less the steady-state real interest rate. Further assumptions are necessary for the selection between these two equilibria.

We extend this analysis to a stochastic environment and analyze how the distributions of inflation associated with each of these two steady states change with the introduction of aggregate uncertainty (see Mendes (2011), Hills, Nakata and Schmidt (2016), and Lansing (2018) for related analysis). Associated with each steady state in the deterministic model, unconditional distributions of interest rates and inflation emerge in the presence of shocks. The two equilibria only differ in the corresponding expectations about future inflation. In the vicinity of the target equilibrium, the interest rate is unconstrained most or all of the time. In contrast, in the distribution near the liquidity trap equilibrium, the interest rate is mostly or always constrained. In contrast to the deterministic case, where the predictions are very stark and clearly at odds with the data in important aspects, the stochastic case is more subtle, with the differences in the equilibria being more a matter of degree.

The existence of the lower bound affects the shapes of the unconditional distribution for interest rates and inflation. In the case of interest rates, the lower bound truncates the distribution from below. For inflation, the presence of the lower bound prevents stabilization in response to all shock realizations and makes the distribution asymmetric. In the distribution associated with the target equilibrium, the presence of the lower bound skews the distribution of inflation to the left and lowers the unconditional median and mean of inflation. In the distribution associated with the liquidity trap equilibrium, the distribution of inflation centers around a lower mean and is truncated at the inflation target such that negative skewness emerges.

We show that, as aggregate uncertainty rises, the two unconditional means of inflation move closer together and eventually are equal. This finding reflects that, with greater variance of shocks, the lower bound constrains less frequently in the distribution associated with the liquidity trap equilibrium but more often in the distribution associated with the target equilibrium. For large enough shock variances, no unconditional mean consistent with the model exists.

In our empirical investigation, we exploit the decrease in the natural rate of interest since the Great Recession (Williams (2017)). In the model, a lower natural rate of interest affects the distributions of interest rates and inflation. In the vicinity of the target equilibrium, a lower natural rate of interest increases the likelihood of being constrained by the lower bound and thus causes expected inflation to decline and otherwise exacerbates the effects of the lower bound on the distributions of inflation and interest rates. In

contrast, in the vicinity of the liquidity trap equilibrium, a lower natural rate of interest causes expected inflation to *increase*, and the effects of the lower bound on the shape of the distributions of inflation and interest rates diminish.

We take these testable implications of our theoretical model to the data to ascertain whether expectations are empirically consistent with the target or liquidity trap equilibrium. We use options data to back out the risk-neutral forecast densities on future nominal interest rates and inflation in the United States. To study the effects of an occasionally binding lower bound on the unconditional distribution, we study the forecast densities of inflation and interest rates over medium-term horizons (see Kitsul and Wright (2013) and Reis (2016)). Looking at a range of options with different strike prices, we can reconstruct the risk-neutral forecast densities of inflation and interest rates at each point in time and study their evolution with a falling natural rate.

The theoretical model suggests an indicator for the severity of the impact that the lower bound has at a given time. This lower bound indicator is defined as the expected value of the interest rate truncation due to the lower bound. That is, it computes how much the lower bound constrains the central bank on average. We show empirically that this indicator summarizes the effects of the lower bound on the forecast densities for interest rates very well and has predictive power for the impact on inflation.

We find clear evidence that financial market participants incorporate the presence of a lower bound in terms of future nominal interest rates, consistent with the predictions associated with the target equilibrium. These findings might be surprising in that the recent episode of a binding lower bound is more likely in the liquidity trap rather than the target equilibrium. By contrast, we find no empirical support for the theoretical implications of the liquidity trap equilibrium. First, the implied probability of a binding lower bound increased during the time when the natural rate of interest fell. Second, the average interest rate fell over the sample period along with the rate of inflation. Third, the forecast density of inflation has shifted to the left. All of these observations are consistent with predictions of the model in the vicinity of the target equilibrium and contradict the predictions of the liquidity trap equilibrium.

Although our findings are qualitatively consistent with the theoretical predictions, the magnitude of the changes in the distribution of inflation expectations are quantitatively small, despite market participants placing a relatively high probability of policy being at the lower bound. This contrasts with results from some studies that suggest very large effects (see Kiley and Roberts (2017)), but is consistent with studies that incorporate a richer set of monetary policy tools and/or fiscal policy that can be effective in putting upward

pressure on prices when short-term interest rates are at the lower bound (see Reifschneider and Williams (2000), Williams (2010), and Reifschneider (2016)).

We point out three caveats in regard to our analysis and results. First, there are relatively few data on inflation and interest rates before the financial crisis, which limits our ability to analyze the behavior of expectations at times when the lower bound was viewed to be less salient. The longest time-series starts in 2006, and we only have full data since 2011. In this regard, comparing data across countries may be useful. Second, we study optimal policy under discretion. If the central bank can commit to future policy actions, its capacity to stabilize expectations and the economy is likely to increase. We leave the full analysis for future research. Third, we study how outcomes of the New Keynesian model change in response to an unanticipated exogenous change in the natural rate of interest. Therefore, we do not address the potential endogeneity or stochastic nature of the decline in the natural rate of interest.

Section 2 presents the key logic that is present in New Keynesian models with a lower bound on interest rates. In this model, we perform comparative statics with respect to a fall in the natural rate of interest. Section 3 discusses the construction of the forecast densities of inflation and interest rates. Section 4 shows the changing forecast densities over the previous years and interprets them through the lens of our theoretical model. Section 5 discusses the robustness of the findings and section 6 concludes.

2 Theoretical Model

We use a textbook New Keynesian model of an economy where the policy instrument is subject to a lower bound to motivate the empirical analysis (Woodford (2003)). Given uncertainty about the modeling of short-run macroeconomic dynamics in the presence of the lower bound, we primarily focus on the ergodic, or unconditional, distribution of inflation and interest rates in the model economy. We are thus able to abstract from model complications and illustrate more clearly the most important theoretical implications of the lower bound for distributions of beliefs.

2.1 The model

The model consists of three equations describing the evolution of three endogenous variables: the inflation rate, π_t, the output gap, x_t, and the short-term nominal interest rate, i_t. Since the focus of the paper is on qualitative changes in the model, we work with the log-linearized version of the standard New-Keynesian model and i.i.d shocks for which we can derive analytical results. The equation describing the behavior of

inflation is given by:

$$\pi_t = \mu_t + \kappa x_t + \beta \mathbb{E}_t \pi_{t+1}, \qquad \mu_t \sim \text{iid}(0, \sigma_\mu^2), \tag{1}$$

where $\mathbb{E}_t$ denotes mathematical expectations based on information at time t, μ_t is a markup shock, $\beta \in (0, 1)$ is the discount factor, and $\kappa > 0$. The equation describing the output gap is given by:

$$x_t = \epsilon_t - \alpha(i_t - \mathbb{E}_t \pi_{t+1} - r^*) + \mathbb{E}_t x_{t+1}, \qquad \epsilon_t \sim \text{iid}(0, \sigma_\epsilon^2), \tag{2}$$

where $\alpha > 0$, r^* is the long-run neutral real rate of interest, and ϵ_t is a demand shock. All agents are assumed to have full knowledge of the model, including the distribution of the shock processes.

The central bank's goal is to keep the output gap near zero and to keep the inflation rate near its target level, which is normalized to zero.[2] Specifically, the central bank chooses its policy instrument, i_t, to minimize the expected quadratic loss:

$$\mathscr{L} = (1 - \beta)\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \lambda x_t^2) \right], \tag{3}$$

where $\lambda \geq 0$ is the relative weight the central bank places on output gap stabilization. The central bank decision for i_t is assumed to occur after the realizations of the shocks in the current period.

The central bank is assumed to lack the ability to commit to future actions; that is, policy is conducted under discretion as in Kydland and Prescott (1977). In addition, the policy action is subject to a lower bound, $i^{LB} < r^*$, that sets a lower limit on i_t for all t.[3] Under these assumptions, combining the equations for inflation and the output gap yields the following expression for the inflation rate (detailed derivations of optimal policy and equilibrium conditions appear in Appendix A):

$$\pi_t = (1 + \alpha\kappa)\mathbb{E}_t \pi_{t+1} + \mu_t + \kappa\epsilon_t - \alpha\kappa(i_t - r^*). \tag{4}$$

Maximizing the objective (3) subject to the equilibrium conditions and lower bound constraint lead to the optimal policy under discretion that depends only on the current state of the economy, which is fully described by the realization of the shocks and the expected value of inflation in the next period and can therefore be

[2]It is straightforward to generalize to a nonzero inflation target by interpreting π_t as the gap between inflation and its target.

[3]In the New Keynesian model, the condition of the lower bound on nominal interest rates being below the natural real rate of interest is a necessary condition for the existence of a steady-state.

written as an interest rate rule:

$$i_t = \max\{r^* + (1 + \frac{1}{\alpha\kappa} - \frac{\lambda\beta}{\alpha\kappa(\kappa^2+\lambda)})\mathbb{E}_t\pi_{t+1} + \frac{1}{\alpha}\epsilon_t + \frac{\kappa}{\alpha(\kappa^2+\lambda)}\mu_t, i^{LB}\}. \tag{5}$$

Plugging in the optimal interest rate rule (5) into (4) leads to the process for inflation. In the following, it is useful to simplify the notation. With $\psi \equiv (1 + \frac{1}{\alpha\kappa} - \frac{\lambda\beta}{\alpha\kappa(\kappa^2+\lambda)})$ and $\gamma \equiv \frac{\kappa}{\alpha(\kappa^2+\lambda)}$, inflation follows

$$\pi_t = \begin{cases} \mu_t + \kappa\epsilon_t - \alpha\kappa(i^{LB} - r^*) + (1+\alpha\kappa)\mathbb{E}\pi, & \text{if } \gamma\mu_t + \frac{1}{\alpha}\epsilon_t \leq i^{LB} - r^* - \psi\mathbb{E}\pi \\ \frac{\lambda}{\kappa^2+\lambda}\left\{\mu_t + \beta\mathbb{E}\pi\right\} & \text{otherwise.} \end{cases} \tag{6}$$

This formulation shows the two different processes for inflation depending on whether the lower bound binds or not. As a result, the deterministic version of the system has two equilibria, one associated with each equation. For the stochastic version of the model, we take the switching between the processes into account.[4]

If the lower bound does not constrain policy in the current period, then optimal policy yields an inflation rate given by the second part of equation (6). Note that the unconstrained optimal policy fully offsets the demand shock ϵ_t. In the special case of $\lambda = 0$, this policy achieves full inflation stabilization, $\pi_t = 0$ for all t, and attains the minimum feasible loss of zero. For the case of $\lambda > 0$, the unconstrained optimal policy balances offsetting markup shocks and deviations of expected future inflation from target against the cost of creating non-zero output gaps. As a result, this policy only partially offsets these two factors that push inflation away from its target value.

We first analyze the deterministic version of the model where σ_ϵ^2 and σ_μ^2 are assumed to equal zero. In that case, the model is characterized by two steady-state values of π. In the one, which we refer to as the "target equilibrium," the steady-state value of the interest rate, denoted by $\bar{i}^u$, equals r^*, and the steady-state value of π, $\bar{\pi}^u$, equals zero. In the second, which we label the "liquidity trap equilibrium," the steady-state value of the interest rate, $\bar{i}^c$, equals the lower bound, and the steady-state value of inflation is given by $\bar{\pi}^c = i^{LB} - r^*$. Without further assumptions, it is not possible to select between these two steady states.[5]

We now extend the analysis to a stochastic environment. As a first step, note that the expected value of π_{t+1} equals its unconditional expectation, denoted by $\mathbb{E}\pi$. In the model, a version of the Fisher equation holds,

[4]Throughout the paper, we abstract from regime switches between different steady states as, e.g. in Lansing (2018).

[5]There are a number of approaches to analyze the stability properties of the two equilibria, which we do not pursue here. That said, it is worth noting that the region of attraction for the target equilibrium is the open unbounded interval to the right of the liquidity trap equilibrium. The region of attraction for the liquidity trap equilibrium is confined to a single point.

whereby the unconditional mean interest rate moves one-for-one with the unconditional mean inflation rate. This is seen by taking the unconditional expectation of equation (4), imposing the steady-state condition, and solving for the expected values:

$$\mathbb{E}\pi = \mathbb{E}i - r^*. \tag{7}$$

We now derive and characterize the equilibria in the stochastic economy. Taking unconditional expectations of both sides of equation (6) yields:

$$\begin{aligned}\mathbb{E}\pi =& \text{Prob}\left(\gamma\mu + \frac{1}{\kappa}\epsilon \leq i^{LB} - r^* - \psi\mathbb{E}\pi\right)\mathbb{E}\left[\mu + \kappa\epsilon - \alpha\kappa(i^{LB} - r^*) + (1+\alpha\kappa)\mathbb{E}\pi \middle| \gamma\mu + \frac{1}{\kappa}\epsilon \leq i^{LB} - r^* - \psi\mathbb{E}\pi\right] \\ &+ \left\{1 - \text{Prob}\left(\gamma\mu + \frac{1}{\kappa}\epsilon \leq i^{LB} - r^* - \psi\mathbb{E}\pi\right)\right\}\mathbb{E}\left[\frac{\lambda}{\kappa^2+\lambda}\left(\mu_+\beta\mathbb{E}\pi\right) \middle| \gamma\mu + \frac{1}{\kappa}\epsilon > i^{LB} - r^* - \psi\mathbb{E}\pi\right].\end{aligned} \tag{8}$$

By construction, the quantity on the right side of the equals sign is non-positive and is strictly negative if the unconditional probability of being at the lower bound is strictly positive.

As can be seen from equation (8), the lower bound binds only occasionally in the environment with sufficient aggregate uncertainty, irrespective of whether the economy is near the target or liquidity trap equilibrium. When the realization of the shocks is sufficiently high, the central bank is unconstrained by the lower bound and can pursue its desired action. Following sufficiently adverse shocks, however, the central bank finds itself constrained by the lower bound, and its inability to sufficiently cut interest rates puts downward pressure on inflation.

An important aspect of this analysis is that we assume that the lower bound always exists and that expectations of future inflation reflect this fact (see Mendes (2011) and Hills, Nakata and Schmidt (2016)). This differs from much of the literature, where expectations are based on the lower bound constraining policy for a finite period in the future (see Fuhrer and Madigan (1997), Reifschneider and Williams (2000), Eggertsson and Woodford (2003), Williams (2010), Evans, Fisher, Gourio and Kane (2015), and Kiley and Roberts (2017)). In the context of our model, such an assumption would imply that $\mathbb{E}_t\pi_{t+j+1} = 0$ for some $j > 1$. If we were to make such an assumption, the ergodic mean of inflation would be unique and closer to the deterministic target equilibrium.

2.2 Specific examples of shock distributions

To analyze the properties of the stochastic model economy, we specify the distribution of the shocks by considering the uniform and normal distributions. While the normal distribution is the most commonly used specification, the case of the uniform distribution has the advantage that it gives rise to an analytical

derivation of the equilibrium conditions.

Example 1: Uniform distribution

We first analyze the model with a uniform distribution of shocks. To simplify the analysis, we assume that there is no demand shock, i.e., $\sigma_\epsilon^2 = 0$. Assume that the markup shock, μ_t, is distributed as a uniform random variable over the interval of $[-\hat{\mu}, \hat{\mu}]$.

The resulting probability that policy is constrained by the lower bound in a given period t is given by:

$$\text{Prob}(\mu_t < \frac{1}{\gamma}(i^{LB} - r^* - \psi\mathbb{E}\pi_{t+1})) = \begin{cases} 1 & \text{if } -\frac{1}{\gamma}(i^{LB} - r^* - \psi\text{E}\pi_{t+1}) \leq -\hat{\mu} \\ \frac{1}{2\hat{\mu}}(\hat{\mu} + \frac{1}{\gamma}(i^{LB} - r^* - \psi\mathbb{E}\pi_{t+1})) & \text{if } -\hat{\mu} < -\frac{1}{\gamma}(i^{LB} - r^* - \psi\mathbb{E}\pi_{t+1}) < \hat{\mu} \\ 0 & \text{if } -\frac{1}{\gamma}(i^{LB} - r^* - \psi\mathbb{E}\pi_{t+1}) \geq \hat{\mu}. \end{cases} \quad (9)$$

In the case of a single shock with a uniform distribution, the probability of policy being constrained by the lower bound is linearly increasing over the support of the distribution, and is either zero or unity otherwise.

The resulting unconditional expectation of inflation in a given period t is given by:

$$\mathbb{E}\pi_t = \begin{cases} -\alpha\kappa(i^{LB} - r^*) + (1 + \alpha\kappa)\mathbb{E}\pi_{t+1} & \text{if } -\frac{1}{\gamma}(i^{LB} - r^* - \psi\mathbb{E}\pi_{t+1}) \leq -\hat{\mu} \\ -\frac{1}{4\hat{\mu}}\left(\frac{\kappa^2}{\kappa^2+\lambda}\left[\frac{1}{\gamma}(i^{LB} - r^* - \psi\mathbb{E}\pi_{t+1}) + \hat{\mu}\right]^2\right) + \frac{\lambda\beta}{\kappa^2+\lambda}\mathbb{E}\pi_{t+1} & \text{if } -\hat{\mu} < -\frac{1}{\gamma}(i^{LB} - r^* - \psi\mathbb{E}\pi_{t+1}) < \hat{\mu} \\ \frac{\lambda\beta}{\kappa^2+\lambda}\mathbb{E}\pi_{t+1} & \text{if } -\frac{1}{\gamma}(i^{LB} - r^* - \psi\mathbb{E}\pi_{t+1}) \geq \hat{\mu}. \end{cases} \quad (10)$$

Note that in the special case of $\lambda = 0$, the optimal policy achieves zero inflation each period (see the last line in equation (10)) except when policy is constrained by the lower bound. In that case, for the intermediate range when policy is constrained some but not all the time, the equation for the expected value of inflation simplifies to: $\mathbb{E}\pi_t = -\frac{1}{4\hat{\mu}}\left(-\hat{\mu} - \alpha\kappa(i^{LB} - r^*) + (1 + \alpha\kappa)\mathbb{E}\pi_{t+1}\right)^2$.

In looking for an unconditional mean, we equate expected values $\mathbb{E}\pi_t$ and $\mathbb{E}\pi_{t+1}$.[6] The set of equations (10) describes the functional relationship between these expected values of π. For a range of intermediate values of $\mathbb{E}\pi$, the relationship is quadratic. Outside of this range, the relationship is piecewise linear, with a slope greater than one for low values of $\mathbb{E}\pi_{t+1}$ and less than one for high values of $\mathbb{E}\pi_{t+1}$.

There are either zero, one, or two values that satisfy the conditions in (10). For small values of $\hat{\mu}$, there are two steady states, as in the deterministic model. Once the degree of uncertainty increases beyond a certain

[6]The same methodology also applies to computing equilibria with a probability of regime switches between the two equilibria computed here.

point, the corner solution where the lower bound is either always or never binding no longer applies. If this occurs in the vicinity of the target equilibrium, the lower bound constrains policy for very low realizations of the shock and, as a result, the unconditional mean of π decreases. If this occurs in the vicinity of the liquidity trap equilibrium, the lower bound does *not* constrain policy in response to very high realizations of the shock, and the unconditional mean of π *increases.*

Figure 1 illustrates the equilibria in the model for different degrees of aggregate uncertainty parameterized by $\hat{\mu}$. To produce Figure 1, we use the following parameter combination:[7] $\alpha = \kappa = 1$, $\beta = 0.99$, $i^{LB} = -0.5\%$, and $r^* = 1\%$. The left panel shows the case of $\lambda = 0$, where the central bank seeks only to stabilize inflation; the right panel shows the case of $\lambda = 0.5$, where the central bank also seeks to stabilize the output gap. In the left panel, the blue line shows the values of $\mathbb{E}\pi_t$ for given values of $\mathbb{E}\pi_{t+1}$ for the deterministic case ($\hat{\mu} = 0$). Note that the relationship is piecewise linear. Two steady states emerge where the function crosses the 45-degree line indicated by the dashed black line. The green line shows the corresponding functional relationship for the case of $\hat{\mu} = 2.25\%$. In this case, the function is quadratic over the relevant range and the quadratic relationship between $\mathbb{E}\pi_t$ and $\mathbb{E}\pi_{t+1}$ crosses the 45-degree line in two places. Note that the unconditional mean in the vicinity of the liquidity trap equilibrium is *higher* than in the deterministic case, and that associated with the target equilibrium is *lower* than in the deterministic case. As the value of $\hat{\mu}$ increases, the two unconditional means move closer together. The red line shows the curve for $\hat{\mu} = 3\%$; in this case, only one equilibrium exists. For values of $\hat{\mu} > 3\%$, no equilibrium exists.

In the graph to the right, the central bank seeks to stabilize both inflation and the output gap. As a result, expected inflation in the current period increases with expected future inflation even in the deterministic relationship. This result comes from the effect that it is no longer optimal for the central bank to keep inflation at zero in response to μ shocks. Again, two equilibria exist and the same pattern emerges when aggregate uncertainty is introduced. The kink in the relationship is smoothed out, the equilibria move closer towards each other, and there is non-existence of a steady-state if uncertainty is sufficiently high.

The lower bound not only affects the mean of inflation, but also the shape of its distribution. For low values of $\hat{\mu}$, in the equilibrium in the vicinity of the liquidity trap equilibrium, policy is always constrained. As a result, the distribution of outcomes echoes the distribution of shocks, centered around the steady-state value. For sufficiently large $\hat{\mu}$, some high positive realizations of the shock are offset by policy action so that

[7] Throughout the paper, we use the model to derive qualitative changes of outcomes in response to a shift in parameters. A different calibration might help in understanding the quantitative aspects of the data but that is not the focus of this paper. In particular, a flatter Phillips curve might lead to smaller changes in moments of inflation.

the distribution of inflation is truncated from above. As $\hat{\mu}$ increases further, this truncation extends further down the distribution of shocks.

The effects of increasing $\hat{\mu}$ on the distribution of π are reversed for the target equilibrium. For low values of $\hat{\mu}$, inflation is always symmetrically distributed around the target. As $\hat{\mu}$ increases beyond a certain point, the lower tail of realizations of the shock cannot be offset by policy actions, and the distribution includes negative outcomes corresponding to situations where policy is constrained by the lower bound. Note that for both equilibria there are two effects on the resulting distribution. First, there is the direct effect of the change in the set of shock realizations that can be offset by policy actions. Second, there is the indirect effect from the resulting shift in the unconditional mean of inflation, which affects the value of current inflation through the expectations channel whenever policy is constrained.

Finally, for $\lambda > 0$, as uncertainty increases the median rate of inflation moves in the same direction as the mean. This occurs because it is costly in terms of the objective function to create a positive output gap to offset the effects when expected inflation differs from the target rate. Thus, a central bank that cares about both the output gap and deviations of inflation from the target will optimally choose to have inflation below target over half of the time. That said, due to the asymmetric nature of the lower bound, the mean inflation rate responds more than the median to uncertainty. As a result, the difference between the mean and median inflation rate increases as uncertainty increases. In the special case of $\lambda = 0$, this tradeoff does not exist and the median inflation rate is generally unaffected by uncertainty.

Example 2: Normal distribution

We now consider the case of the normal distribution of shocks. Qualitatively, the results are very similar to the case of a uniform distribution. The one key difference is that the function mapping $\mathbb{E}\pi_{t+1}$ to $\mathbb{E}\pi_t$ is smooth over the range of relevant values. The general formula for expected inflation with both shocks is

$$\mathbb{E}\pi_t = \frac{1}{2\left(\kappa^2+\lambda\right)}\left(\phi\left(\mathbb{E}_t\pi_{t+1}, i^{\mathrm{LB}}, r^*\right)\left(1+\operatorname{erf}\left(-\frac{\phi\left(\mathbb{E}_t\pi_{t+1}, i^{\mathrm{LB}}, r^*\right)}{\kappa\sigma_{\epsilon\mu}}\right)\right) - \frac{1}{\sqrt{\pi}}\kappa\sigma_{\epsilon\mu}e^{-\frac{\phi\left(\mathbb{E}_t\pi_{t+1}, i^{\mathrm{LB}}, r^*\right)^2}{\kappa^2\sigma_{\epsilon\mu}^2}} + 2\beta\lambda\mathbb{E}_t\pi_{t+1}\right) \tag{11}$$

where $\phi\left(\mathbb{E}_t\pi_{t+1}, i^{\mathrm{LB}}, r^*\right) = \left(\kappa^2+\lambda\right)\left(\mathbb{E}_t\pi_{t+1}(\alpha\kappa+1) - \alpha\kappa(i^{\mathrm{LB}}-r^*)\right) - \beta\lambda\mathbb{E}_t\pi_{t+1}$ and the standard deviation of the linear combination of the two shocks $\sigma_{\epsilon\mu} = \sqrt{2\sigma_\epsilon^2\left(\kappa^2+\lambda\right)^2 + 2\kappa^2\sigma_\mu^2}$.

Figure 2 illustrates the equilibria in the model for different degrees of aggregate uncertainty. The same parameter values are assumed as in the previous figure. Note that the relationship is similar to that for the uniform distribution, except the function is smooth throughout. For values of $\sigma > 1.9\%$, no equilibrium exists

when $\lambda = 0$. The graph to the right shows the plot for positive λ. The same pattern emerges as in Figure 1.

We take away from this section that the specific choice of the distribution is not critical for our empirical analysis. The qualitative results and their underlying reasoning are identical under both the uniform and the normal distribution. Independent of the distribution, there are two possible equilibria if uncertainty is small. Sufficiently large shocks will lead to the possibility of an occasionally binding lower bound in both the target equilibrium and the liquidity trap equilibrium. Therefore, expected values of inflation will move closer together with higher aggregate uncertainty.

2.3 Empirical implications from the theoretical model

This section discusses testable predictions to distinguish whether the economy is in the vicinity of a target or in a liquidity trap equilibrium. After their presentation, we take these predictions to the data. In the following, we focus on *changes* in the unconditional distribution of interest rates and inflation resulting from *changes* in the natural rate of interest, r^*. As this section will make clear, the implied distributions behave differently in the two equilibria when the level of r^* varies. Note that the changes in the implied distribution are derived from comparative statics with respect to r^*, i.e., changes in the parameter that agents in the model treat as constant.

First, consider the relationship between r^* and mean inflation and interest rates. In the deterministic liquidity trap equilibrium, a lower value of r^* *raises* the steady-state value of inflation. The intuition comes from the Fisher equation that relates nominal interest rates to real rates and expected inflation. If real rates fall with r^* and nominal interest rates are constrained by the lower bound, inflation must rise to satisfy the Fisher equation. In the stochastic economy associated with the liquidity trap equilibrium, a lower r^* also increases the range of shocks for which policy is unconstrained leading expected inflation to increase more than one-for-one. The unconditional mean of the interest rate also increases in this case. In contrast, in the deterministic target equilibrium, a decrease in r^* has no effect on the mean inflation rate since policy can be adjusted to offset it. The mean interest rate therefore declines one-for-one with the decline in r^*. In the stochastic economy, a lower value of r^* increases the set of shocks for which policy is constrained, and, as a result, the unconditional mean of inflation declines. In this case, the unconditional mean of the interest rate declines by more than one-for-one with the decline in r^*.

Figure 3 illustrates these effects. This graph, along with the other figures discussed in this section, uses the model with a normal distribution of shocks parameterized as above and assuming $\lambda = 0.5$ and $\sigma_\mu = 1\%$.

The curve for the unconditional mean of inflation is the mirror image from the graph for the mean of interest rates. This relationship arises from equation (7). The graphs show different steady-state outcomes for a range of r^* values where the blue line represents the target and the red line the liquidity trap equilibrium. In the target equilibrium, the means of π and i move in the same direction as the natural rate of interest. For the liquidity trap equilibrium, depicted in red lines, the unconditional means of inflation and interest rates move in the opposite direction as r^*.

Figure 4 depicts the resulting probabilities of being constrained by the lower bound for various values of r^*. The blue line shows how the probability of being constrained by the lower bound rises when the natural rate of interest falls in the target equilibrium. The red line shows that this prediction is reversed in the liquidity trap equilibrium. The reason for this is that, with a lower r^*, the unconditional mean of the inflation rate is higher, as discussed above, and this implies that the lower bound constrains policy less often.

The shapes of the distributions of inflation and interest rates are also affected by the level of the natural rate of interest. Specifically, in the target equilibrium, a lower value of r^* shifts the distribution of interest rates to the left, and the increased probability of hitting the lower bound implies that the asymmetry in the distribution increases. When the interest rate is constrained more frequently, the distribution of inflation moves to the left (if $\lambda > 0$), and the left skewness of the distribution of inflation increases. When r^* is very low, a further reduction leads to less negative skewness such that the overall relationship is U-shaped. In the vicinity of the liquidity trap equilibrium, skewness is negative and further decreases with lower values of the natural rate of interest. Figure 5 illustrates these effects for a range of values of r^*.

Taken together, our simple model yields several testable predictions regarding the responses of the unconditional distributions of inflation and interest rates to a decline in the natural rate of interest. In all but one set of predictions, the responses of the distributions in the target versus the liquidity trap equilibrium are exactly opposite.

3 Construction of Forecasts for Interest Rates and Inflation

This section describes the options data and methodology we use to construct forecast densities for interest rates and inflation. The methodology used in this paper for extracting forecast densities is borrowed from Kitsul and Wright (2013) and Durham (2008). From this procedure, we obtain a time series of forecast densities on a daily frequency. That is, on each day we extract the market-implied forecast densities for inflation and interest rates at various horizons. Since we compare these forecast densities to the unconditional distributions

in our theoretical model, we primarily focus on long-term forecasts.

3.1 Long-term forecast densities for interest rates

We obtain a daily data series of caps on the London Interbank Offered Rate (LIBOR) from Bloomberg that is available since July 2008. The main advantage of using these options is the long horizon of forecasts, available for up to ten years out. An interest rate cap is a series of consecutive European call options, or "caplets," on interest rates that provide the holder with protection against rising interest rates over the life of the contract. For example, the holder of a ten-year interest rate cap on three-month LIBOR will receive a payment at the end of every three-month period over the following ten years if LIBOR exceeds the strike at the beginning of the same three-month period. If the cap is written for a period of ten years, and τ is quarterly, there are 39 potential payoffs made at T_{6M}, T_{9M},..., T_{10Y}. The value of the cap is the sum of all its caplets. This value is quoted in the market as the Black volatility which is the implied volatility from the Black formula that can be used to price the caplet.

Interest rate caps are reasonably liquid with liquidity declining at longer horizons. The contracts are traded over-the-counter but are among the most commonly traded OTC interest rate derivatives. Caps of strikes 1 through 14 percent, in terms of LIBOR, and horizons of up to ten years are used. Prices of caplets are computed by subtracting a shorter maturity cap from the price of a cap with one extra caplet.

We extract forecast densities from options with different different strike prices (Durham (2008)). We recover the forecast density by twice differentiating the Black pricing formula for the jth interest rate caplet, with respect to the strike price k

$$p(k) = \exp(\iota t)\frac{\partial^2}{\partial k^2}\text{Caplet}_j^{\text{Black}}, \tag{12}$$

where $p(k)$ is the density of the underlying interest rate associated with strike k at expiry and ι is the continuously -compounded risk-free interest rate. Therefore, one needs to find a twice differentiable function for the implied volatility to substitute into $\text{Caplet}_j^{\text{Black}}$. This function is given by regressing implied volatility on a quadratic function of the strike price

$$\sigma = \alpha_0 + \alpha_1 k + \alpha_2 k^2 + \nu, \tag{13}$$

where σ is the observed spot implied volatility over a cross-section of caplet prices, and the predicted value is twice differentiable in k. The predicted value from (13) is substituted into Black's formula for caplets, which

then delivers the forecast densities according to equation (12). We map the resulting forecast densities for LIBOR to fed funds rates for which we compute several percentiles.[8] A list of all the steps involved in the procedure is given in Appendix B.

3.2 Inflation forecasts

We use data on index options on the average annual inflation rate over the lifetime of the security from Bloomberg. A cap pays if average (annually compounded) consumer price index (CPI) inflation exceeds the strike rate. The data on caps is available on a daily frequency. Trading in inflation caps started in late 2009. We focus on contracts of maturities five and ten years. Inflation options are traded over-the-counter. Options on inflation have been reasonably liquid between years 2011 and 2017. Since that period, however, liquidity has decreased to the point that Bloomberg no longer provides quotes on these contracts.

The seller of a zero-coupon inflation cap promises to pay a fraction $\max\{(1+\pi_n)^n - (1+k)^n, 0\}$, of a notional underlying principle as a single payment in n years time, where π_n denotes the average annual total CPI inflation from t to $t+n$ and k denotes the strike of the cap. Options prices are obtained from dealer quotes and do not necessarily reflect actual trades. However, Kitsul and Wright (2013) show that inflation forecasts based on these contracts deliver economically reasonable results.

To extract forecast densities for inflation, we follow the methodology in Kitsul and Wright (2013). We apply the method to both five-year and ten-year inflation options. Specifically, we use quotes on inflation caps at different strike prices to find the probability density function for the n-year inflation π_n consistent with observed prices according to

$$P_{t,n}(k) = \exp(-n\iota_n)\,\mathbb{E}^q\left[\max\left\{(1+\pi_n)^n - (1+k)^n, 0\right\}\right], \tag{14}$$

where $P_{t,n}(k)$ denotes the price of an n-year inflation cap with strike price k and ι_n is the continuously compounded n-year zero-coupon bond yield at time t. $\mathbb{E}^q$ denotes expectations under the risk-neutral measure. The risk-free rate is taken from the nominal Treasury term structure as in Gürkaynak, Sack and Wright (2007). Not all prices necessary for the computation are available. Therefore, we approximate the price of an inflation cap at a strike price k' in a neighborhood around k by a local linear function $\beta_0(k)+\beta_1(k)(k'-k)$.

[8] This adjustment allows us to interpret the implied distribution as forecast densities over the central bank's policy rate.

We estimate the coefficients as:

$$\hat{\beta}_0(k), \hat{\beta}_1(k) = \arg \min_{\hat{\beta}_0(k), \hat{\beta}_1(k)} \Sigma_{i=1}^{L} \left\{ y_i - \beta_0(k) - \beta_1(k)(k_i - k) \right\}^2 \frac{1}{h} K\left(\frac{k_i - k}{h}\right), \tag{15}$$

where L is a set of inflation caps, K is the kernel function, and h is a bandwidth. In a last step, we derive the implied forecast density based on the result that the second derivative of the price of a call option with respect to the strike price represents the risk-neutral probability density function. The percentiles of the forecast density can thus be computed from $\hat{\beta}_1^{'}(k)$.

3.3 Conversion to forward rates

There are many different ways to convert forecast densities of average inflation over some time period into forward rates. However, none of them are fully satisfactory. Therefore, we implement a simple mapping that assumes the forward rate maps the nth percentile at time t_1 to the nth percentile at time t_2. Thus, our method for extracting forward rates is given by solving the following set of equations for the different forward rates f^p:

$$\left(1 + \pi_{t_1}^p\right)^{t_1} \left(1 + f^p\right)^{t_2 - t_1} = \left(1 + \pi_{t_2}^p\right)^{t_2}, \tag{16}$$

where $\pi_{t_n}^p$ is the options-implied rate at the pth percentile at time n and f^p is the forward rate that maps the density function at time n to time $n + t$.

Note that the implementation of the formula does not imply that inflation is perfectly persistent. It merely states that the percentiles for the forward rates are averages of percentiles for short-term and long-term forecast densities.

Also note that this methodology does not need to be implemented for interest rates, since the payouts for the underlying options contracts are already constructed as forward rates.

4 Empirical Analysis and Results

In the empirical analysis of this section, we compare distributions of expected interest rates and inflation derived from financial data to the unconditional distributions from our model.[9] We exploit the recent decline

[9] A related strand of the literature uses options to study investors' expectations about interest rates and inflation. Wright (2017) surveys the literature on extracting probability distributions for interest rates. Fleckenstein, Longstaff and Lustig (2017) extract the physical probability distribution from options data. Kitsul and Wright (2013) compute forecast densities of inflation options, and Reis (2016) uses them to investigate unconventional monetary policy.

in the natural rate of interest as a source of variation. We then examine whether the predictions from the model that arise from a decline in r^* match the experience of the U.S. economy. Therefore, we assume that time variation in r^* is the primary driver of changes in the data. As shown before, the model delivers a rich set of predictions for the implied distributions for inflation and interest rates in response to a permanent decline in r^*.

Note that all forecasts are under the risk-neutral probability measure, i.e., the q-measure. First, since the price of risk in the New Keynesian model is constant, the derived moments for interest rates and inflation can be thought of as being under the risk-neutral measure as well. Second, we treat the resulting variation in measured forecasts as primarily stemming from changes in actual forecasts rather than from the variation in risk premia. Variation in risk premia would not be able to explain the observed patterns in the data as we discuss below.

4.1 The decline in the natural rate of interest

Considerable empirical evidence suggests a sizable decline in the longer-run natural real rate of interest in the United States over the past decade (Williams (2017)). Figure 6 shows estimates of the natural real interest rate for 1998 to 2017 from Christensen and Rudebusch (2017).[10] Consistent with other estimates in the literature, measured r^* reached historically low levels in recent years and does not show signs of moving back to previously normal levels despite the fact that the U.S. economy has now fully recovered from the Great Recession. As seen in the figure, there appears to be a break in the series of estimates occurring at the end of 2011, with the mean estimate dropping from about 1 percent over 2008-2011 to 1/4 percent over 2012-2017.

The low level for the natural rate of interest is likely to persist for an extended period of time for at least three reasons. First, measures of historical levels of r^* display a significant amount of persistence. The estimates aim to capture the low frequency component of real short-term interest rates and, as such, are highly persistent. For example, the estimates reported here from Christensen and Rudebusch (2017) correspond to real rates expected to prevail five to ten years in the future. Holston, Laubach and Williams (2017) show that their estimates of the natural rate are nonstationary, reflecting a high degree of persistence. Second, even well into the recovery from the Great Recession of 2007-2009, r^* has not returned to historically normal levels.

[10]Other measures of the natural rate of interest from models of Laubach and Williams (2003), Kiley (2015), Lubik and Matthes (2015), Johannsen and Mertens (2016), Holston, Laubach and Williams (2017), Crump, Eusepi and Moench (2017), and Del Negro, Giannone, Giannoni and Tambalotti (2017) show a consistent picture of a decline in r^*.

Therefore, it is likely that long-term influences are holding the natural rate of interest down. Third, many possible explanations for low r^*, not only in the United States but internationally, reflect highly persistent forces affecting the global supply and demand for savings. For example, one potential explanation for the decline in r^* is a dramatic slowdown in trend real GDP growth in many advanced economies. For a more detailed discussion, see Williams (2017).

For these reasons, it is highly likely not only that the natural rate of interest will remain low for the medium term but also that investors share this expectation. Therefore, we can use the decline in the natural rate of interest as a source of variation to study the behavior of forecast densities.

4.2 Implied distribution for interest rates

For a nonparametric way of looking at the data, we show the raw data for the long-term forecasts in Figure 7. The upper panel shows a 20-day moving average of long-term forecast densities represented by percentiles. The figure confirms that there is considerable uncertainty with the 97.5th percentile ranging up to interest rates above ten percent. It also confirms the presence and importance of a lower bound. The green line in the upper panel of Figure 7, representing the bound on the 15th percentile, converges towards the lower bound. Due to noise in financial market prices, the estimate for the lower bound varies slightly over time but can be placed consistently close to zero. Both the upper panel and the lower left panel in Figure 7 demonstrate that the average long-term forecast of interest rates has decreased over time. The lower right-hand panel shows the difference of percentiles on the upper end of the distribution (blue line), measured as the difference between the 97.5th and 50th percentiles, as well as on the lower end (red line) measured as the difference between the 50th and 2.5th percentiles. While the difference in percentiles on the upper end has stayed roughly constant, the gap has closed on the lower end.

Table 1 summarizes our findings with regards to forecast densities for interest rates. It displays various features consistent with the predictions of the target equilibrium in our theoretical model. First, the average prediction for interest rates falls over the later subsample during which the natural rate of interest was lower. Second, the median fell by slightly more than the mean such that the distribution became more asymmetric. Third, with a higher mass at the lower bound, the variance fell over the latter part of the sample. This evidence is consistent with the economy being in the target equilibrium rather than the liquidity trap equilibrium. And fourth, the skewness of interest rates increased over that subsample.

The patterns in Table 1 are consistent with the target equilibrium of our theoretical model and at odds

Table 1: Summary of interest rates moments

	2008-2011	2012-2017
Mean	3.85%	2.60%
Median	3.45%	2.11%
Std. deviation	2.38%	1.95%
Skew	0.53%	0.72%

with the liquidity trap equilibrium. A fall in the natural rate of interest makes a given stance of policy less effective at boosting the economy. Therefore, the central bank would find it optimal to lower the policy rate in the target equilibrium. As a result, the average interest rate falls and the probability of a binding lower bound would increase. As before, we can use the experiment to distinguish between the two different equilibria. In the liquidity trap equilibrium where interest rates are mostly constrained by the lower bound, a fall in the natural rate would increase average interest rates.

Consistent with the theory for both equilibria, the skewness of interest rates is positive. The two equilibria differ in the direction in which the skewness of interest rates changes when the natural rate of interest declines. In line with the target equilibrium, Table 1 shows an increase in the skewness during the second part of the sample. This view is consistent with the lower bound truncating the range of possible policy rates. Furthermore, this trend suggests that the emerging asymmetry of the distribution is an important factor in the decrease of the variance of forecast densities.

The left panel in Figure 11 depicts these effects by fitting a truncated normal distribution to the percentiles in the two subsamples. The results show that the probability of a binding lower bound in the later part of the sample period is higher than in the earlier part while the perceived lower bound is slightly lower. The asymmetry of the distribution is driven by the truncated part of the distribution.

To summarize, the empirical evidence suggests that the lower bound on interest rates has a sizable effect on expectations of market participants. All the pieces of empirical evidence overwhelmingly suggest that the economy is in the target equilibrium region. All measured changes due to the decrease in the natural rate of interest are at odds with the behavior of interest rates in the liquidity trap equilibrium. However, the theoretical model has predictions not only about the control variable but also inflation. In a following section, we test the predictions with regards to inflation.

4.3 A lower bound indicator

Equation (8) suggests that a single summary statistic should capture most of the effects of the lower bound on interest rates. This statistic, which we call the lower bound indicator, is the expected value of the interest rate truncation due to the lower bound. When the lower bound is binding, the central bank would find it optimal to set interest rates below the lower bound but is unable to do so. The lower bound indicator is defined as how much more the central bank would have liked to cut interest rates on average. Based on this theoretical motivation, we define a lower bound indicator $\mathscr{I}$ at time t as

$$\mathscr{I}_t = \text{Prob}\left[i_t^{opt} < i^{LB} | \mathbb{I}_t\right] \mathbb{E}_t\left[i^{LB} - i_t^{opt} | i_t^{opt} < i^{LB}, \mathbb{I}_t\right], \tag{17}$$

where $\mathbb{I}_t$ is the time t information set and i_t^{opt} is the optimal unconstrained interest rate as defined in the first argument of the max operator in equation (5), $i_t^{opt} = r^* + (1 + \frac{1}{\alpha\kappa} - \frac{\lambda\beta}{\alpha\kappa(\kappa^2+\lambda)})\mathbb{E}_t\pi_{t+1} + \frac{1}{\alpha}\epsilon_t + \frac{\kappa}{\alpha(\kappa^2+\lambda)}\mu_t$.

In our model, the expected interest rate truncation is the driving force for asymmetries in interest rate and inflation densities, where the latter can be seen from equation (8). Therefore, theory tells us that the lower bound indicator in (17) should capture the changes in mean values and asymmetries of both interest rates and inflation. In fact, within our model there is a linear relationship between the lower bound indicator and average interest rates and inflation as the natural rate of interest varies. For the asymmetry of the distribution, the relation is close to being linear.

Guided by the theoretical insights, we obtain a measure of the lower bound indicator $\hat{\mathscr{I}}_t$ for each quarter from the data. Therefore, we assume a truncated normal distribution for interest rates and compute estimates of the mean, standard deviation, and lower bound.[11] Using these three estimates for the forecast density observed within each quarter, we generate a time series of the lower bound indicator $\hat{\mathscr{I}}$. Details about the construction of the lower bound indicator are listed in Appendix C.

Figure 8 demonstrates that lower estimates of the natural rate of interest are associated with elevated levels of the lower bound indicator. Since asymmetries in interest rates and inflation tend to increase with lower levels of r^*, they should also increase with the lower bound indicator. Empirically, this hypothesis holds up. Figure 9 shows that two measures of asymmetry, the difference between mean and median, standardized by the standard deviation, as well as the skewness in interest rate forecast densities increase with higher levels of the lower bound indicator.

[11] An alternative specification where we impose a constant lower bound and only fit mean and variance leads to very similar results.

Taken together, we conclude that the theoretically motivated lower bound indicator also works well in practice to capture the effects of the lower bound on financial markets.

4.4 Implied distribution for inflation

The previous section demonstrates that the dynamics of interest rates are consistent with the model of Section 2. To understand whether the mechanics of the model are in line with empirical estimates, we study long-term forecast densities of inflation.

As for interest rates, the time series of percentiles for forecast densities of inflation deliver a first look at whether the predictions of the model are consistent with the data. The upper panel in Figure 10 shows the time series of long-term inflation forecast densities, again plotted as a 20-day moving average of a five-year forward rate starting in five years. Compared to interest rate forecasts, the distribution looks strikingly symmetric. As shown in the lower left-hand panel in Figure 10, average inflation forecasts decreased during the sample period. The right-hand side shows the difference between the 97.5th and 50th percentiles versus the difference between the 50th and 2.5th percentiles. The distribution became slightly more asymmetric during the sample period. This can be seen from the graph where the red line lies above the blue line in the later part of the sample. Overall however, both lines slope downwards, which is at odds with the theory but attributable to the decrease in the variance.

Table 2 summarizes the moments for subsamples. Consistent with the model predictions for the target equilibrium, the mean of inflation declines, falling from 2.60 percent in the 2011 sample to 2.42 percent for 2012-2017. The density of inflation is slightly skewed to the left. The broader message of a decline in average inflation forecasts is consistent with the target equilibrium of our theoretical model.

Table 2: Summary of inflation moments

	2011	2012-2017
Mean	2.60%	2.42%
Median	2.67%	2.43%
Std. deviation	1.97%	1.29%
Skew	-0.09%	-0.06%

Table 2 also shows that the variance of expected inflation fell in the later part of the sample. This does not appear to be related to the decline in the natural rate of interest, but instead appears to be a reversal of unusually high uncertainty about inflation during and directly following the financial crisis and recession.

In support of this hypothesis, the variance of density forecasts of longer-term inflation expectations in the Survey of Professional Forecasters shows a marked increase in the period 2009-2012, followed by a return to pre-crisis levels in the past few years.

The skewness of inflation is negative in both subsamples, consistent with the theory. The fall in the skewness during the later subsample can only be rationalized by the target equilibrium where the relationship between the skewness of inflation and the natural rate of interest is U-shaped.

The right panel of Figure 11 shows these results graphically. It fits a split-normal distribution, i.e. a distribution consisting of a normal distribution on either side of the mode but with potentially different standard deviations, to the distribution of inflation. With the lower natural rate of interest, the average forecast for inflation fell. Asymmetries in the estimated distributions are present but hardly noticeable. Quantitatively, the effects of a decline in the natural rate of interest on inflation forecasts are modest, a topic we return to later.

Figure 12 shows the link between the lower bound indicator and inflation. The left panel shows that higher levels of the lower bound indicator are, on average, associated with lower levels of inflation, as predicted by the theory. The right panel shows again that the link between interest rates and inflation might be weaker than the theory predicts. The skewness of inflation is negative and increases with higher levels of the lower bound indicator. These facts can be reconciled with the theory where the skewness is negative and U-shaped, and thus ambiguous on the direction, in the target equilibrium.

4.5 Discussion

The results speak very clearly: Qualitatively, the target equilibrium of the New Keynesian model is in almost every aspect consistent with the changes in forecast densities considered here. There is no evidence to support the view that the U.S. economy was in a liquidity trap in the aftermath of the Great Recession.

Two questions arise from our analysis of inflation expectations. First, why are the quantitative effects on inflation expectations so small, despite market participants appearing to place over 30 percent probability of policy being at the lower bound in the future? Second, why is there no convincing evidence of significant asymmetry in the distribution of inflation beliefs, even with very low expectations of future interest rates? For comparison, Kiley and Roberts (2017) find that with a 1 percent natural real rate of interest, the lower bound constrains policy 38 percent of the time, and the distribution of inflation is highly skewed to the left with a mean that is 0.8 percentage points below target.

There are a number of potential explanations for this disconnect between the theoretical predictions and the evidence. One possibility, of course, is that the theoretical model does not capture well the link between interest rates and inflation during times when the lower bound is a constraint. A second possibility is that market participants expect the Federal Reserve or other parts of the government to stimulate the economy when the lower bound constrains. A number of tools can be used to combat low inflation even when interest rates are constrained by the lower bound. First, the central bank can try to commit to future actions, for example, using forward guidance. In theory, this can be a powerful tool that mitigates the effects of the lower bound (for example, see Reifschneider and Williams (2000), Eggertsson and Woodford (2003), Adam and Billi (2006), and Kiley and Roberts (2017)). Second, the central bank can engage in forms of quantitative easing by purchasing longer-term government securities or other assets (see Chung, Laforte, Reifschneider and Williams (2012) and Reifschneider (2016)). Third, the fiscal authority can provide stimulus to the economy (Williams (2010)). In model-based analyses that incorporate these additional policy tools, the effects of the lower bound on the distribution of inflation tend to be relatively modest.

The observed patterns in the changes of forecast densities for inflation and interest rates cannot be explained by changes in risk premia alone. For interest rates, an increased probability of the lower bound binding is consistent with changes in risk pricing if the marginal utility associated with a lower bound increased during our sample period. For inflation, however, while increased marginal utilities during lower bound episodes can explain lower average inflation under the risk-neutral measure, the less negative skewness in the later part of the sample is at odds with a risk-pricing explanation.

Our results are related to the findings of Swanson and Williams (2014a) and Swanson and Williams (2014b), who also find clear evidence that the lower bound on interest rates affects the behavior of expectations for future short-term interest rates but not for longer-term interest rates or the exchange rate. One interpretation of those findings is that market participants expect the central bank to use quantitative easing to push down long rates when short rates are constrained by the lower bound.

5 Robustness: Liquidity in Options Data

For robustness, we look at short-term forecast densities for interest rates based on Eurodollar options. These data have the advantage that the time series is available over a long period and liquidity for these contracts tends to be high. The data, however, have the disadvantage that the information set of investors includes a nontrivial amount of information about the short-run dynamics of interest rates. Therefore, these short-term

forecast densities do not measure the unconditional distribution of interest rates as cleanly as our long-term forecasts. Nevertheless, these forecasts serve well as a robustness check. We discuss the construction of forecast densities in Appendix D.

We perform the analogous estimation we did for long-term interest rates in which we fit a distribution to the forecast densities. Due to the longer time series for short-term forecast densities, we study the pre-crisis time period between September 2006 and the end of 2007 as well as the most recent 2.5 years from 2015-2017.

	2006-2007	2015-2017
Mean	4.33%	1.19%
Median	4.37%	1.03%
Std. deviation	1.27%	0.67%
Skew	-0.07%	0.75%

Table 3: Estimated parameters from short-term forecast densities for a truncated normal distribution for interest rates over different subsamples.

Table 3 shows the results. Prior to the Great Recession, the probability of hitting the lower bound on interest rates was considered to be small. The forecasts have changed substantially. Lower average interest rates have contributed to a higher risk of being constrained in the near future. As a result, the distribution has become asymmetric. These predictions are exactly in line with the target equilibrium of our theory and at odds with the liquidity trap equilibrium.

The time series for short-term forecast densities in Figure 13 confirms the finding that investors put a significant weight on the influence of the lower bound binding after the Great Recession, particularly for the time between 2012 and 2014. After the lift-off of the federal funds rate, the one-year-ahead forecast contains the information that interest rates are unlikely to be cut over this time horizon. As can be seen in the later part of the sample, the probability of a binding lower bound declines significantly.

The moments in the raw data also line up. The lower left panel of Figure 13 shows a decline in the mean during the time the natural rate of interest decreased, consistent with the picture we got from the long-term forecasts. The graph in the lower right-hand panel shows the increase in the asymmetry with the onset of the Great Recession, which trended up until the peak in early 2012. The skewness then plateaued between 2012 and 2014 when the probability of a binding lower bound was high. With the economic recovery and an upwards shift in the forecast density, the skewness came down sharply.

Taken together, the empirical evidence based on Eurodollar options confirms the previously discussed empirical evidence. In addition, it shows that the broad message from the theoretical model applies to the

time period prior to the Great Recession as well. However, this latter point is subject to the caveat that short-term densities vary more with the current state of the economy than long-term forecasts.

6 Conclusion

This paper assesses the impact of a lower bound on the unconditional distribution of interest rates and inflation. We study forecast densities for interest and inflation rates implied by options from U.S. financial markets during a time when the natural rate of interest fell. We compare the changes in these forecast densities to those predicted by a theoretical model. We find clear evidence that financial market participants incorporate the presence of a lower bound on interest rates into their forecasts.

In our model, two equilibria can arise: In a target equilibrium, the central bank largely succeeds in stabilizing the economy, while inflation in a liquidity trap equilibrium fluctuates in response to shocks. We work out the changes in the unconditional distribution of interest rates and inflation rates in response to a fall in the natural rate of interest.

We find that our empirical evidence strongly suggests that the experience of the U.S. economy after the Great Recession is well in line with the target equilibrium of the New Keynesian model. That said, we find quantitatively modest effects on beliefs about the behavior of inflation as measured by options-implied distributions of future inflation. Further research is needed to explore why the effects appear to be so small.

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*, October 2006, *38* (7), 1877–1905.

Benhabib, Jess, Stephanie Schmitt-Grohé, and Martín Uribe, "The Perils of Taylor Rules," *Journal of Economic Theory*, 2001, *96* (1), 40 – 69.

Christensen, Jens H E and Glenn D Rudebusch, "A New Normal for Interest Rates? Evidence from Inflation-Indexed Debt," *Federal Reserve Bank of San Francisco, Working Paper 2017-07*, 2017.

Chung, Hess, Jean-Philippe Laforte, David Reifschneider, and John C. Williams, "Have We Underestimated the Likelihood and Severity of Zero Lower Bound Events?," *Journal of Money, Credit and Banking*, 2012, *44*, 47–82.

Crump, Richard K., Stefano Eusepi, and Emanuel Moench, "The Term Structure of Expectations and Bond Yields," *Federal Reserve Bank of New York Staff Report no. 775*, 2017.

Durham, J Benson, "Implied Interest Rate Skew, Term Premiums, and the "Conundrum"," *Journal of Fixed Income*, 2008.

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

Evans, Charles, Jonas D. M. Fisher, François Gourio, and Spencer Kane, "Risk Management for Monetary Policy Near the Zero Lower Bound," *Federal Reserve Bank of Chicago Working Paper No. 2015-03*, 2015.

Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig, "Deflation Risk," *The Review of Financial Studies*, 2017, *30 (8)*, 2719Ű2760.

Fuhrer, Jeffrey C and Brian F Madigan, "Monetary policy when interest rates are bounded at zero," *Review of Economics and Statistics*, 1997, *79* (4), 573–585.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, "The U.S. Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, 2007, *54* (8), 2291 – 2304.

Hamilton, James D., Ethan S. Harris, Jan Hatzius, and Kenneth D. West, "The Equilibrium Real Funds Rate: Past, Present, and Future," *IMF Economic Review*, Nov 2016, *64* (4), 660–707.

Hills, Timothy S, Taisuke Nakata, and Sebastian Schmidt, "The Risky Steady State and the Interest Rate Lower Bound," *Finance and Economics Discussion Series 2016-009. Board of Governors of the Federal Reserve System*, 2016.

Holston, Kathryn, Thomas Laubach, and John C. Williams, "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics*, 2017, *108* (S1), S59–S75.

Johannsen, Benjamin K. and Elmar Mertens, "A Time Series Model of Interest Rates with the Effective Lower Bound," *Finance and Economics Discussion Series 2016-033, Federal Reserve Board of Governors*, 2016.

Kiley, Michael T., "What Can the Data Tell Us About the Equilibrium Real Interest Rate?," *Finance and Economics Discussion Series 2015-077, Federal Reserve Board of Governors*, 2015.

Kiley, Michael T and John M Roberts, "Monetary policy in a low interest-rate world," *Brookings Papers on Economic Activity*, 2017.

Kitsul, Yuriy and Jonathan H. Wright, "The Economics of Options-Implied Inflation Probability Density Functions," *Journal of Financial Economics*, 2013, *110* (3), 696 – 711.

Kydland, Finn E and Edward C Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *The Journal of Political Economy*, 1977, pp. 473–491.

Lansing, Kevin J., "Endogenous Regime Switching Near the Zero Lower Bound," *Federal Reserve Bank of San Francisco Working Paper 2017-24*, 2018.

Laubach, Thomas and John C. Williams, "Measuring the Natural Rate of Interest," *Review of Economics and Statistics*, 2003, *85* (4), 1063Ű70.

Lubik, Thomas A. and Christian Matthes, "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches," *FRB Richmond Economic Brief 15-10*, 2015.

Mendes, Rhys R, "Uncertainty and the Zero Lower Bound: A Theoretical Analysis," *Bank of Canada*, 2011.

Negro, Marco Del, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti, "Safety, Liquidity, and the Natural Rate of Interest," *Brookings Papers on Economic Activity, Spring*, 2017.

Reifschneider, David L., "Gauging the Ability of the FOMC to Respond to Future Recessions," *Finance and Economics Discussion Series 2016-068. Washington: Board of Governors of the Federal Reserve System*, 2016.

Reifschneider, David L and John C Williams, "Three lessons for monetary policy in a low-inflation era," *Journal of Money, Credit and Banking*, 2000, pp. 936–966.

Reis, Ricardo, "Funding Quantitative Easing to Target Inflation," *Jackson Hole Symposium*, 2016.

Swanson, Eric T and John C Williams, "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates," *American Economic Review*, 2014, pp. 3154–3185.

Swanson, Eric T. and John C. Williams, "Measuring the Effect of the Zero Lower Bound on Yields and Exchange Rates in the U.K. and Germany," *Journal of International Economics*, 2014, *92* (Supplement 1), S2 – S21. 36th Annual NBER International Seminar on Macroeconomics.

Williams, John C, "Heeding Daedalus: Optimal Inflation and the Zero Lower Bound," *Brookings Papers on Economic Activity*, 2010, *2010* (2), 1–37.

Williams, John C., "Three Questions on R-star," *Federal Reserve Bank of San Francisco Economic Letter*, 2017.

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

Wright, Jonathan H., "Forward Looking Estimates of Interest Rate Distributions," *Annual Review of Financial Economics*, 2017, *9.*

A The New Keynesian model

The standard New Keynesian model is given by equations (1) and (2). Solving these equations yields the equation for inflation

$$\pi_t - E_t\pi_{t+1} = \mu_t + \kappa(\epsilon_t - \alpha(i_t - E_t\pi_{t+1} - r^*))) + \beta E_t(\pi_{t+1} - \pi_{t+2}). \tag{18}$$

Under optimal monetary policy with discretion, the final term is zero, so we are left with the equation:

$$\pi_t = (1 + \alpha\kappa)\mathbb{E}_t\pi_{t+1} + \mu_t + \kappa\epsilon_t - \alpha\kappa(i_t - r^*), \tag{19}$$

which is equation (4) in the main body of the paper.

B Extraction of long-term densities for interest rates

The procedure we use involves the following steps:

1. Estimate the surface of flat volatilities using quotes on interest rate caps from Bloomberg;
2. Estimate the three-month forward LIBOR curve, using quotes on Eurodollar futures, forward rate agreement (FRA), and overnight interest rate swaps (OIS);
3. Use the surface of flat volatilities to back out the spot volatility of each underlying caplet;
4. Map the strikes from the LIBOR space into the federal funds rate space using basis swap and map the spot implied volatility on LIBOR into that on the federal funds rate by scaling it by a constant term[12];
5. Regress spot volatilities on the corresponding strikes and strikes-squared, where both volatilities and strike are represented in terms of the federal funds rate. Having recovered this non linear relationship, we are now able to compute a value of volatility corresponding to any strike value, even if the strike is not in our sample;
6. Replace the implied volatility by the quadratic function estimated in the previous step and differentiate the Black formula twice with respect to the strike to obtain the pdf for the future federal funds rate.

[12]Federal funds versus LIBOR basis swap contracts are used to convert the FRA term structure to an expected federal funds rate path. An additional one basis point is subtracted per month to match the Board's term premium assumption.

7. Generate percentiles from the constructed pdf.

C Construction of lower bound indicator

We assume a truncated normal distribution for interest rates. Therefore, we need three inputs: Mean and variance of the underlying normal distribution and the lower bound. We first aggregate the daily time series for the percentiles into a quarterly series by averaging each percentile within the quarter. We take the time series of the median as a measure for the mode of the truncated normal distribution since mean, mode, and median coincide for the normal distribution. To estimate the variance, we convert the difference between the median and the 85th percentile into an estimate of the standard deviation. Therefore, we note that for a normal distribution, the standard deviation is proportional to the difference in the two quantiles, with a proportionality factor of $1/(\sqrt{2}\mathrm{erfc}^{-1}\left(\frac{3}{10}\right))$. With the lower bound being the 2.5th percentile, we have all inputs to construct the truncated normal distribution.

D Short-term forecast densities for interest rate

To ensure that liquidity is not a driver of our findings, we obtain data on options based on Eurodollar futures from Bloomberg for comparison with our longer-term forecasts. Eurodollar futures prices reflect market expectations of interest rates on three-month Eurodollar deposits for specific dates in the future. Thus, the options we use are contracts capturing market expectations of a three-month forward rate n years out. We focus on the 24-month horizon starting in January 2006.

Options on Eurodollar futures are among the most actively traded exchange-listed interest rate options in the world, with an average of over $1.2 trillion trading in notional value per day in 2016. A total of 40 quarterly futures contracts, spanning ten years, plus the four nearest serial (nonquarterly) months are listed at any time. Serial Eurodollar futures are identical to the quarterly contracts except that they expire in months other than those in the March, June, September, and December quarterly cycle. Serial options on Eurodollar futures are American Style and available for the nearest four months. Quarterly contracts are available for the nearest 16 quarterly months.

The forecast densities of the three-month LIBOR are extracted by specifying a functional form for the pdf of the three-month LIBOR and performing nonlinear optimization to obtain parameters that minimize the

average distance between observed options prices and the pdf-implied option prices at various strikes.[13] The implied LIBOR probability density function is then mapped to federal funds rate space.

The model-implied price of a Eurodollar futures call option with strike price K is

$$\hat{c}\left(K,t,T;\theta\right)=\exp\left(-r^{f}\left(T-t\right)\right)\int_{0}^{\infty}\max\left(R-K,0\right)f\left(R;\theta\right)dR, \tag{20}$$

where $\theta=\left\{\phi,\mu_1,\sigma_1,\mu_2,\sigma_2\right\}$, R is the three-month LIBOR in the future, K is the strike price, and f indicates either a mixture of normal or the lognormal distributions. Thus, f is given by :

$$f^{normal}\left(R;\theta\right)=\phi\rho\left(R;\mu_1,\sigma_1\right)+\left(1-\phi\right)\rho\left(R;\mu_2,\sigma_2\right) \tag{21}$$

where $\rho\left(R;\mu_1,\sigma_1\right)=\frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{1}{2}\frac{\left(r-\mu_i\right)^2}{\sigma_i^2}\right)$

$$f^{lognormal}\left(R;\theta\right)=\phi g\left(R;\mu_1,\sigma_1\right)+\left(1-\phi\right)g\left(R;\mu_2,\sigma_2\right) \tag{22}$$

where $g\left(R;\mu_i,\sigma_i\right)=\frac{1}{\sqrt{2\pi}\sigma R}\exp\left(-\frac{1}{2}\frac{\left(\log R-\mu_i\right)^2}{\sigma_i^2}\right)$.

Similarly, the model-implied price of a Eurodollar futures put option with strike price K is given by

$$\hat{p}\left(K,t,T;\theta\right)=\exp\left(-r^{f}\left(T-t\right)\right)\int_{0}^{\infty}\max\left(K-R,0\right)f\left(R;\theta\right)dR. \tag{23}$$

To find the parameters, minimize the sum of squared pricing errors

$$\begin{aligned}\theta^{*}=&\arg\min_{\theta}\{\Sigma_j\left(c_j\left(K_j,t,T\right)-\hat{c}\left(K_j,t,T;\theta\right)\right)^2\\&+\Sigma_k\left(p_k\left(K_k,t,T\right)-\hat{p}\left(K_k,t,T;\theta\right)\right)^2\\&+\left(F\left(t,T\right)-\hat{F}\left(t,T;\theta\right)\right)^2*w\}\end{aligned} \tag{24}$$

using all out-of-the-money options with the same time to maturity $T-t$. The parameter w is chosen large enough to ensure that the futures rate, $F\left(t,T\right)$, corresponds to the mean of the distribution, $\hat{F}\left(t,T;\theta\right)$. The final step consists of adding the LIBOR-OIS spread to map the LIBOR rate to the federal funds rate.

[13]These calculations were conducted by staff at the Board of the Governors of the Federal Reserve System.

Figures

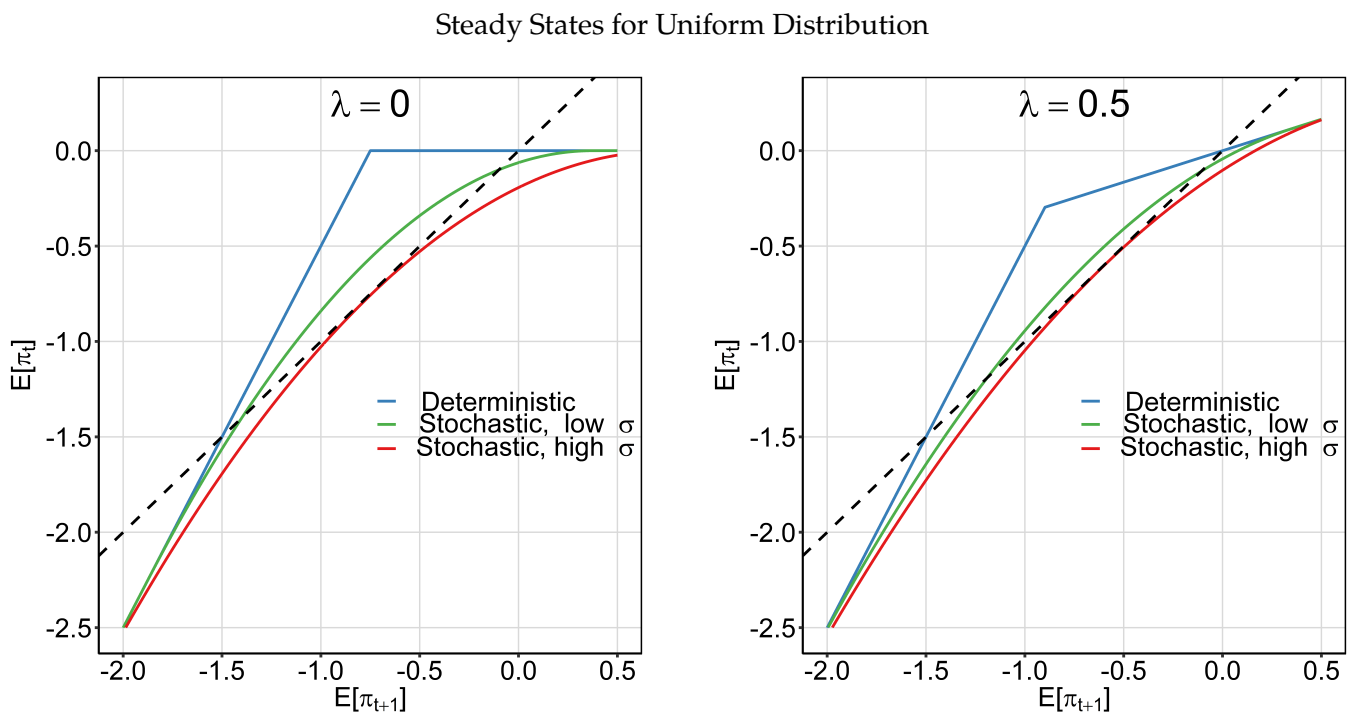

Figure 1: Expected inflation in the current period as a function of expected inflation in the following period assuming a uniform distribution. Parameter values are set to $\alpha = 1$, $\kappa = 1$, $\beta = 0.99$, $r^* = 1\%$, $i^{LB} = -0.5\%$. There are no ϵ-shocks, i.e., $\sigma_\epsilon = 0$. Intersections with the dashed line, the 45-degree line, represent steady states. The lines correspond to different levels of uncertainty parameterized by $\hat{\mu}$, which scales the support for the distribution. Low σ corresponds to $\hat{\mu} = 2.25\%$ for $\lambda = 0$ and $\hat{\mu} = 3.15\%$ for $\lambda = 0.5$. High σ corresponds to $\hat{\mu} = 3\%$ for $\lambda = 0$ and $\hat{\mu} = 3.75\%$ for $\lambda = 0.5$.

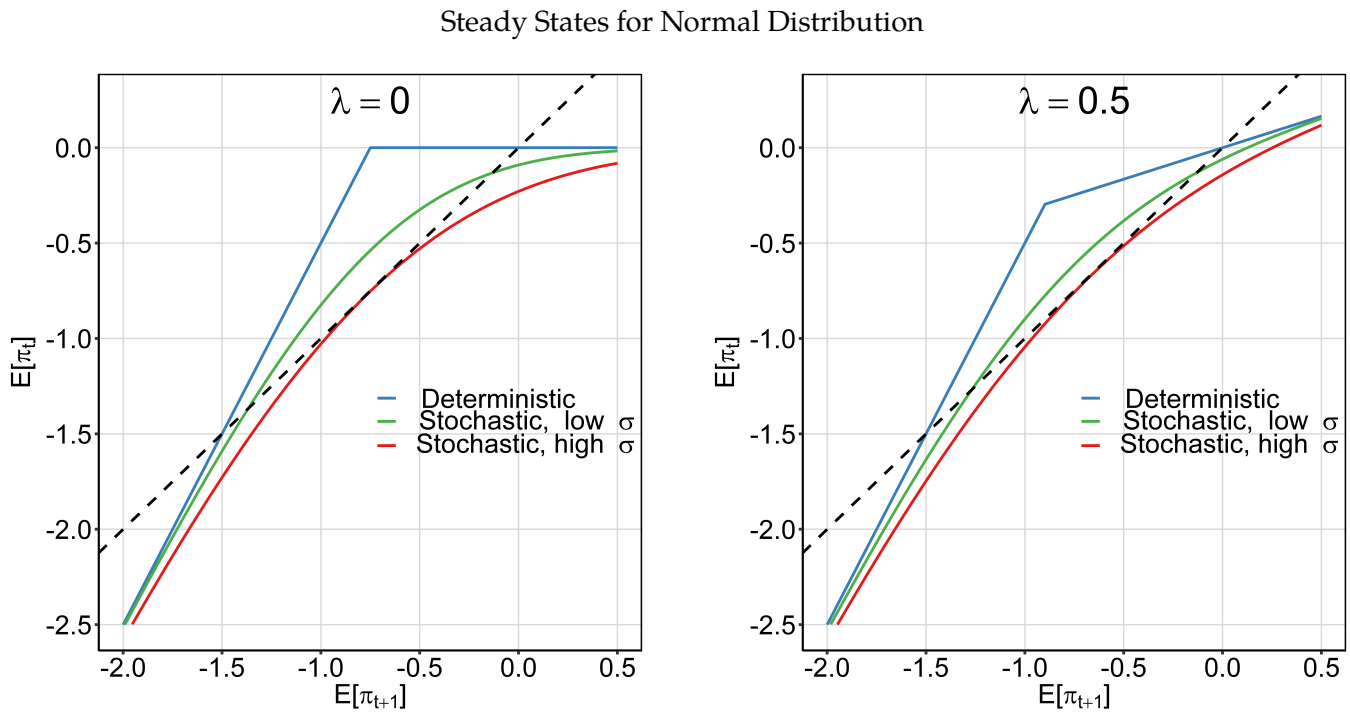

Figure 2: Expected inflation in the current period as a function of expected inflation in the following period assuming a normal distribution. Graphs are analogous to Figure 1 but show values for normally distributed shocks for which uncertainty is parameterized by the standard deviation σ. Low σ refers to a value of $\sigma_\mu = 1.35\%$ for $\lambda = 0$ and $\sigma_\mu = 1.8\%$ for $\lambda = 0.5$. High σ refers to a value of $\sigma_\mu = 1.9\%$ for $\lambda = 0$ and $\sigma_\mu = 2.35\%$ for $\lambda = 0.5$.

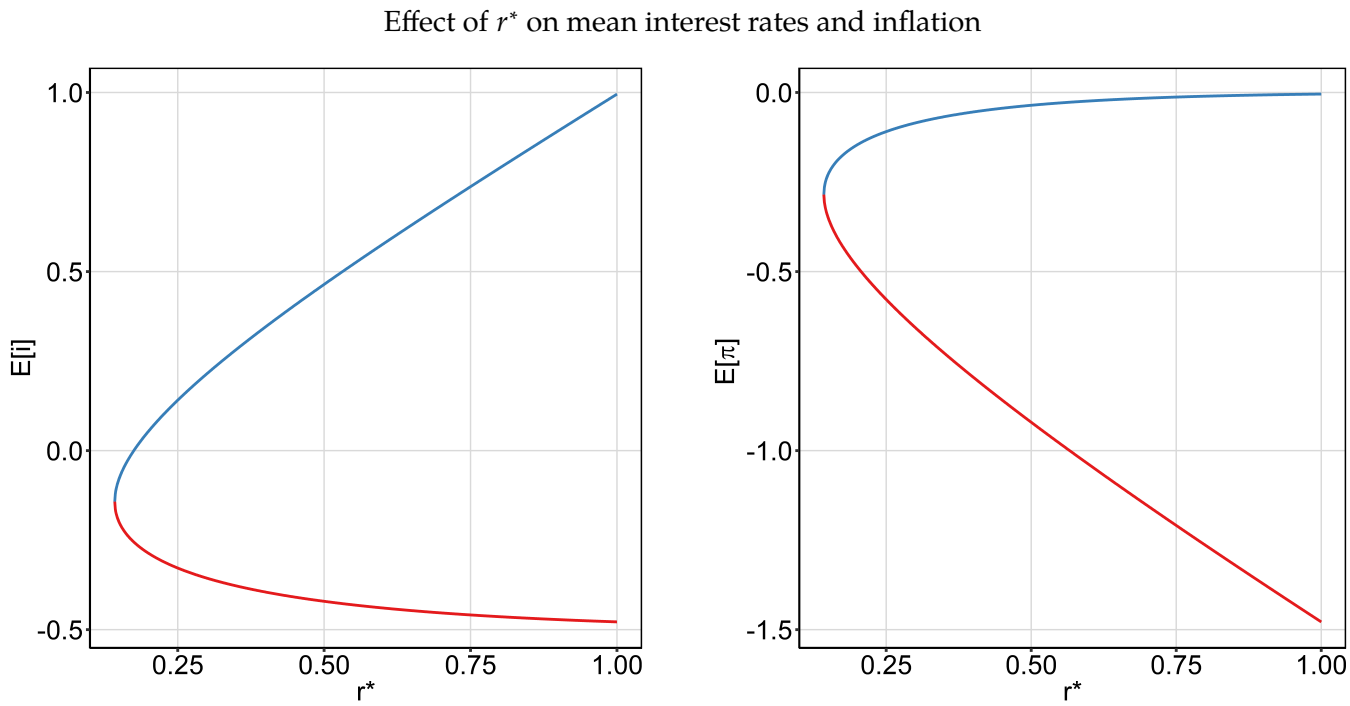

Figure 3: The blue lines show average interest rates (left panel) and average inflation (right panel) for various levels of r^* in the target equilibrium; the red lines display the analogues for the liquidity trap equilibrium. Parameterization is as in Figure 2 with $\sigma_\mu = 1\%$ and $\lambda = 0.5$. Average interest rates rise when $r*$ falls in the liquidity trap equilibrium, but they fall under these conditions in the target equilibrium. Likewise, average inflation rises with $r*$ in the liquidity trap equilibrium, but falls with these conditions in the target equilibrium.

Effect of r^* on probability of a binding lower bound

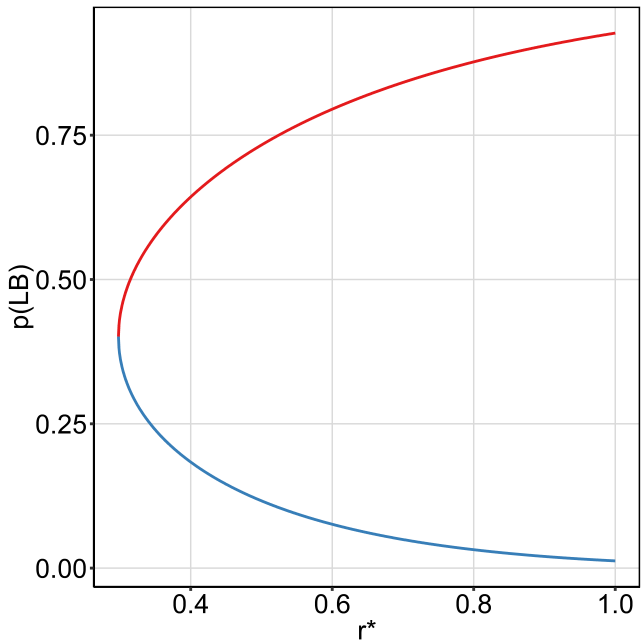

Figure 4: As $r*$ decreases, the probability of hitting the lower bound rises in the target equilibrium (blue line). In the liquidity trap (red line), the probability of being constrained falls. Parameterization is as in Figure 3, again with $\sigma_\mu = 1\%$ and $\lambda = 0.5$.

Effect of r^* on skewness of interest rates and inflation

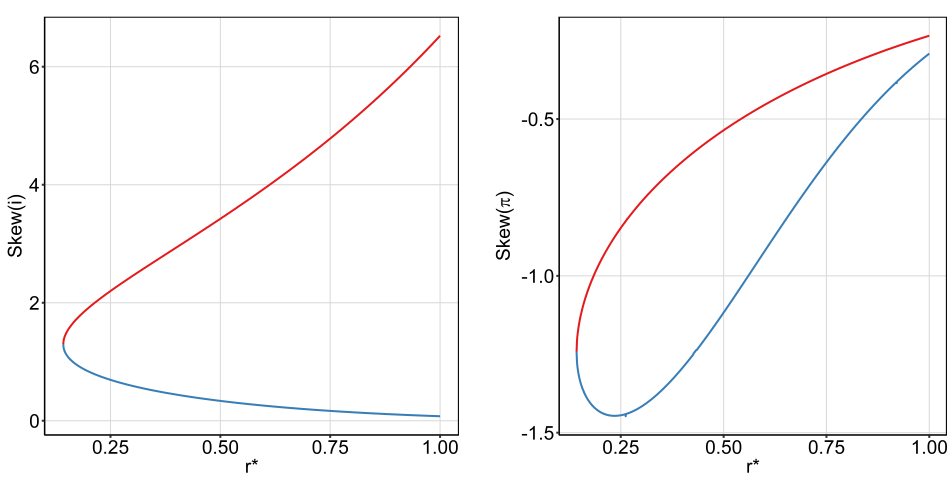

Figure 5: As $r*$ decreases, the distribution for interest rates becomes more asymmetric in the target equilibrium (blue line, left panel) while the skewness of inflation is U-shaped (blue line, right panel). In the liquidity trap equilibrium (red lines), the skewness of interest rates and inflation falls with r^*. Parameterization is as in Figure 3.

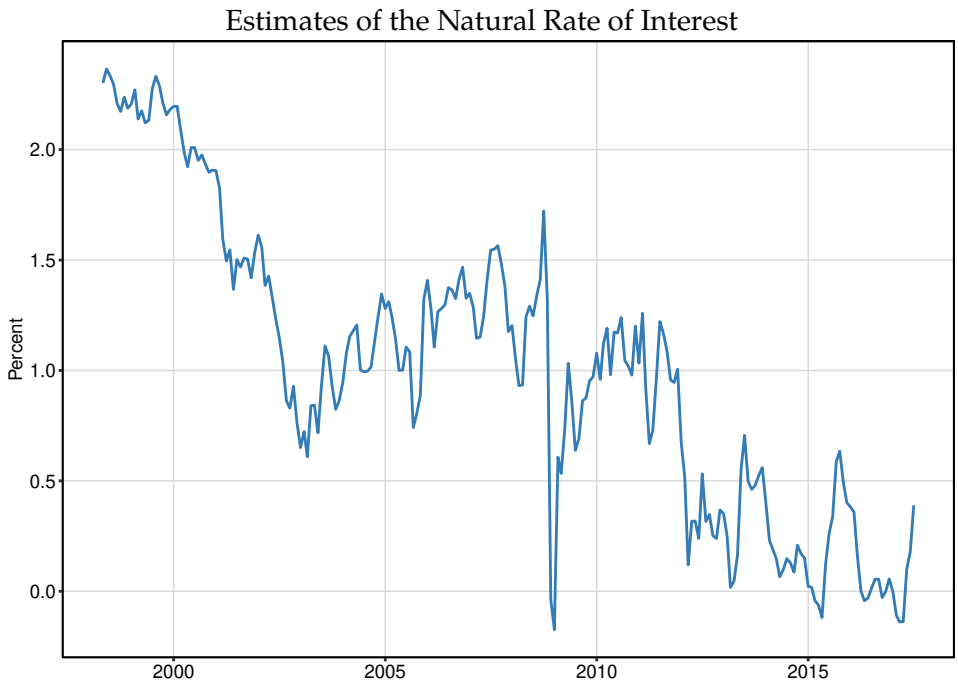

Figure 6: Estimates of the natural rate of interest from Christensen and Rudebusch (2017).

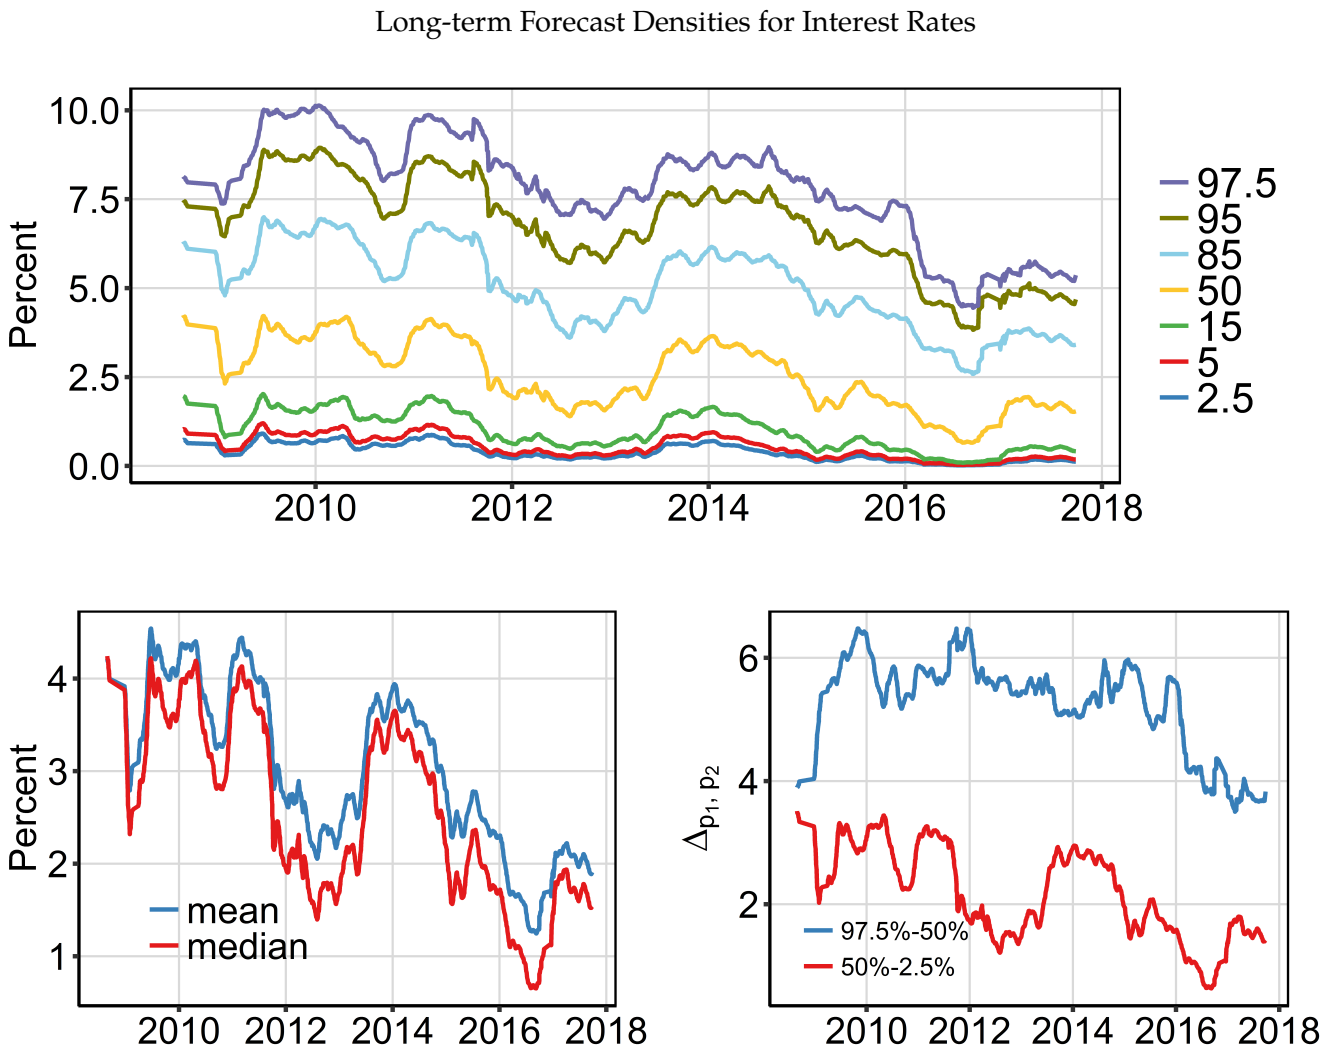

Figure 7: Long-term forecast densities for interest rates measured by the three-month forward rate seven years out. The upper panel shows a 20-day moving average of the time series for various percentiles. The lower panels plot 20-day moving averages for the mean and median of forecast densities and the asymmetry of the distribution via the differences in the 97.5th and 50th percentile versus the 50th and 2.5th percentile (right-hand side).

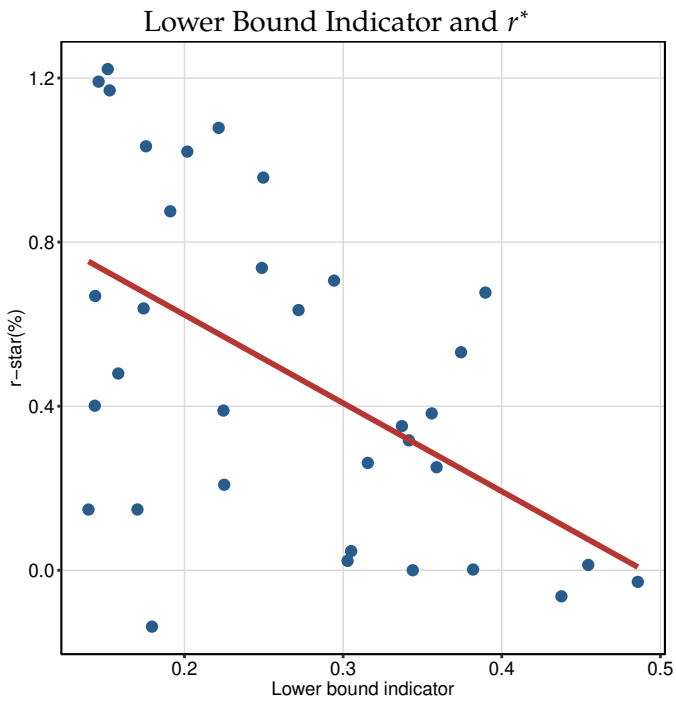

Figure 8: Relationship between the lower bound indicator and r^*. Lower levels of the natural rate of interest are associated with higher measures of our lower bound indicator. The graph displays quarterly data.

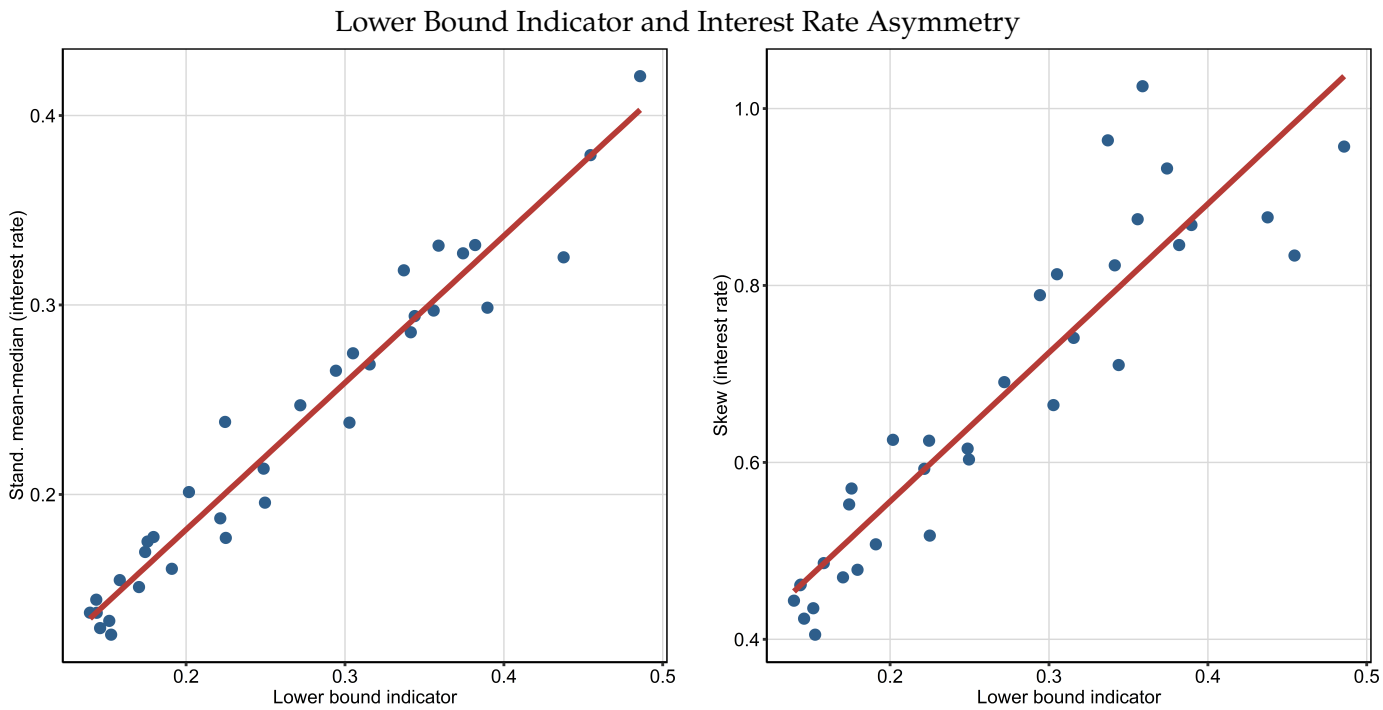

Figure 9: Relationship between the lower bound indicator and the asymmetry of interest rates. This figure shows two measures of asymmetry in interest rate forecast densities, the difference between mean and median as well as the skewness and how they relate to the lower bound indicator at a quarterly frequency.

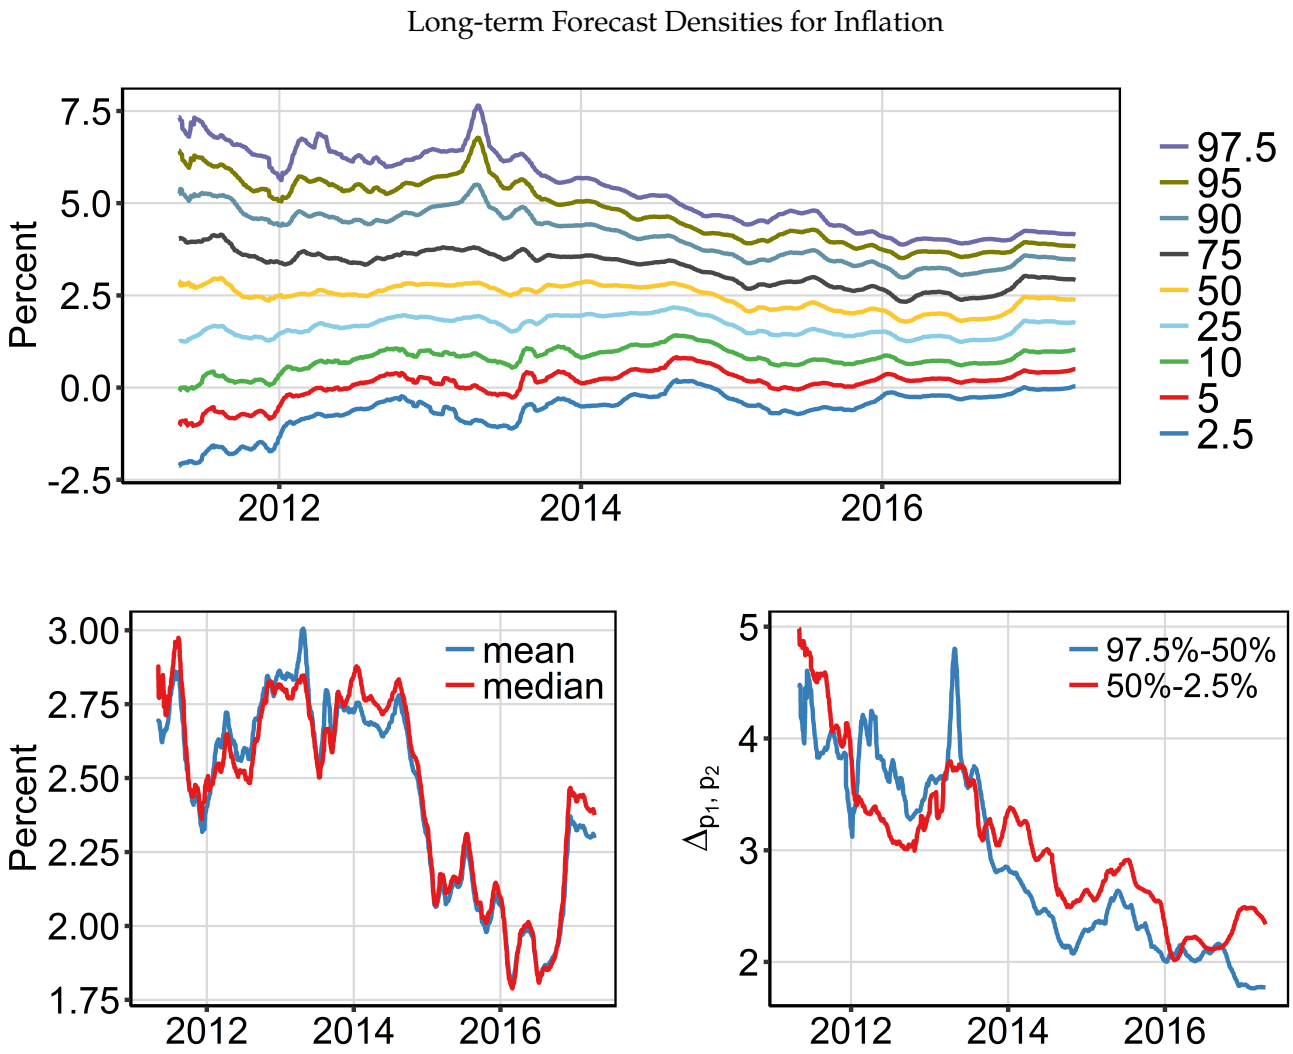

Figure 10: Long-term forecast densities for inflation measured by the five-year forward rate five years out. This figure reports measures of forecast densities analogous to Figure 7 for inflation rates measured by the five-year forward rate five years out. All pictures show 20-day moving averages of the underlying daily time series. The graph displays quarterly data.

Changes in distributions for interest rates and inflation

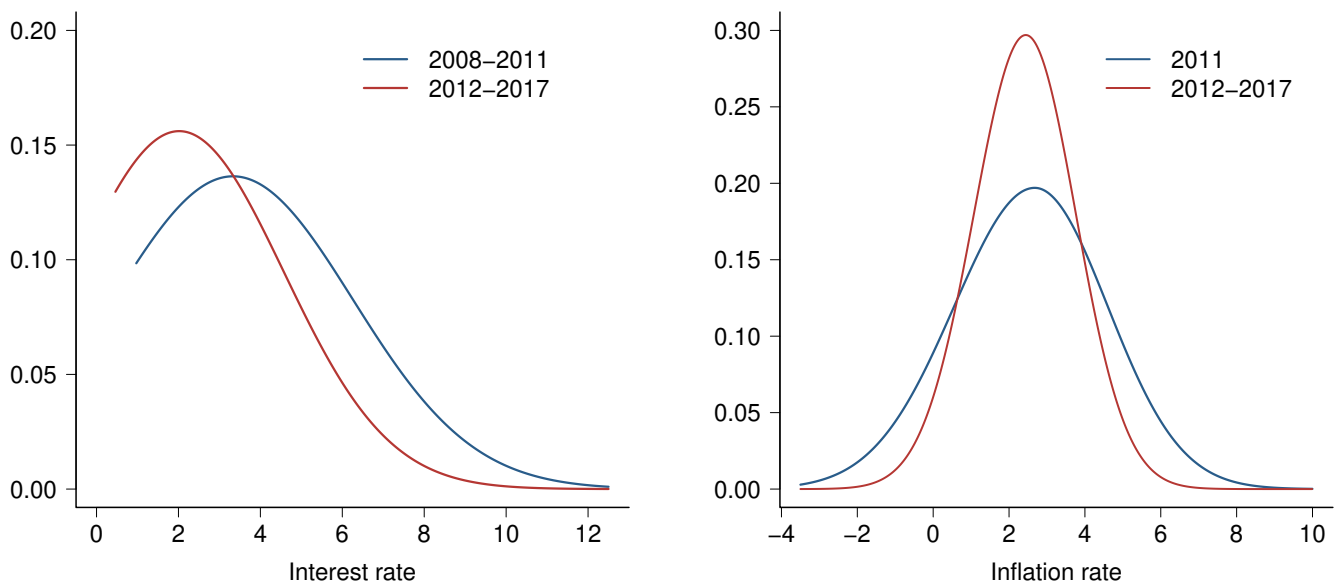

Figure 11: Fitted distributions for interest rates and inflation for the time periods pre- and post-2011. As in the New Keynesian model, the interest rates are represented by a truncated normal distribution. As can be seen from the graph, the probability of a binding lower bound rises in the later part of the sample. Estimated parameters are $\mu = 3.33$, $\sigma = 2.93$, $i^{LB} = 0.96$ in the earlier subsample and $\mu = 2.01$, $\sigma = 2.55$, $i^{LB} = 0.46$ in the later subsample. Inflation densities are approximated by a split-normal distribution that allows for asymmetry. Post 2011, the distribution has a lower mean and variance and is more asymmetric towards low inflation. The parameter values are a mean of $\mu = 2.68$ and standard deviations $\sigma_L = 2.13$ to the left and $\sigma_R = 1.92$ to the right for the earlier subsmaple and $\mu = 2.44$, $\sigma_L = 1.367$, and $\sigma_R = 1.32$ for the later subsample.

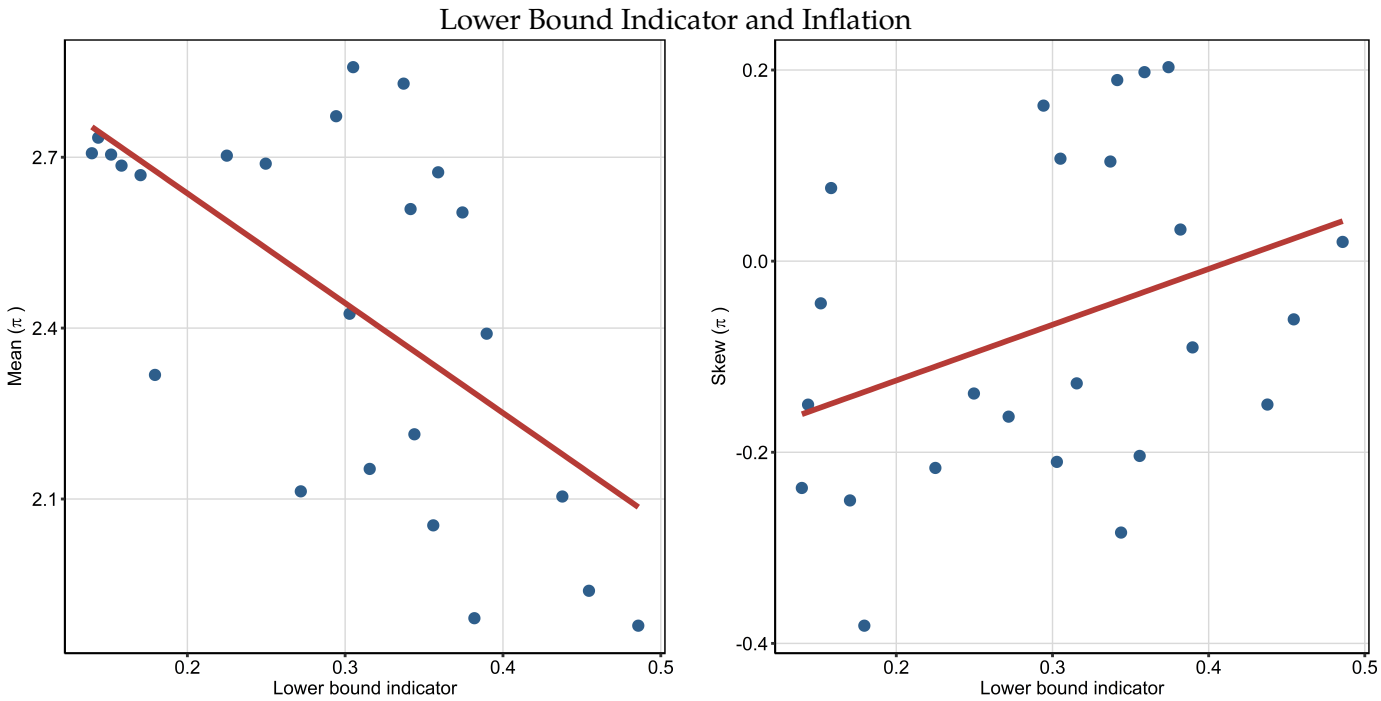

Figure 12: Relationship between lower bound indicator and inflation moments at a quarterly frequency. The scatter plots show quarterly data of mean inflation versus the lower bound indicator (left panel) and the skewness in inflation forecast densities (right panel).

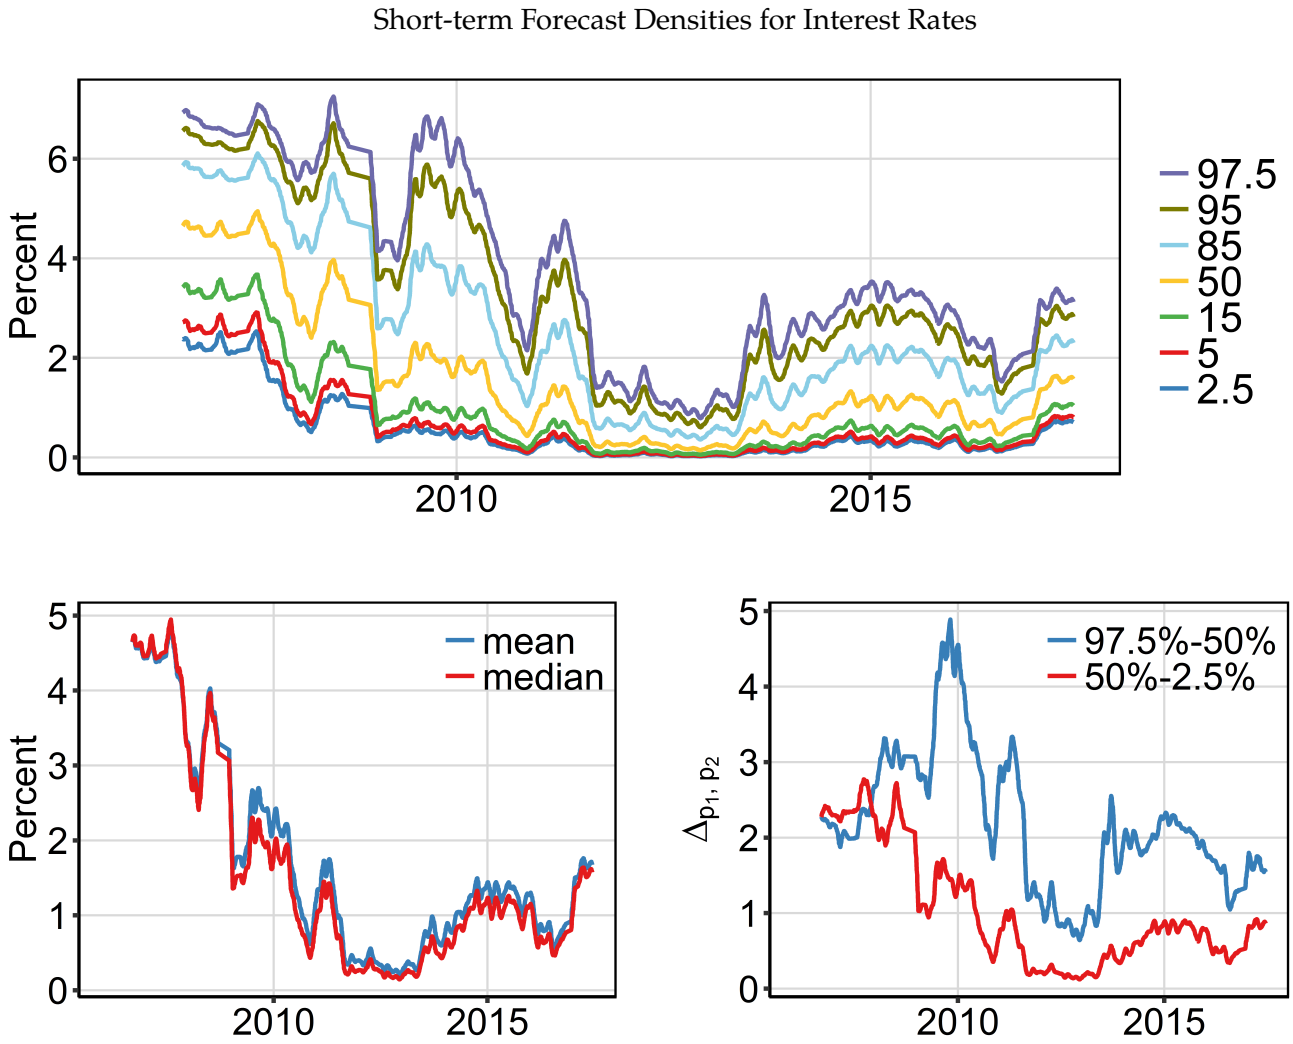

Figure 13: This figure reports the same measures of forecast densities as in Figure 7 for short-term densities. The figures represent 20-day moving averages of percentiles and moments of the distribution of the three-month forward rate 24 months out.