

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

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

Thomas M. Mertens
Federal Reserve Bank of San Francisco

John C. Williams
Federal Reserve Bank of New York

September 2020

Working Paper 2018-03

<http://www.frbsf.org/economic-research/publications/working-papers/2018/03/>

Suggested citation:

Mertens, Thomas M., John C. Williams. 2020. “What to Expect from the Lower Bound on Interest Rates: Evidence from Derivatives Prices,” Federal Reserve Bank of San Francisco Working Paper 2018-03. <https://doi.org/10.24148/wp2018-03>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

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

Thomas M. Mertens[†] and John C. Williams[‡]

September 2, 2020

Abstract

This paper analyzes the effects of the lower bound for interest rates on the distributions of 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 real rate of interest leads to forecast densities consistent with the theoretical model. Qualitatively, we find that the evidence is 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. We develop a lower bound indicator that captures the effects of the lower bound on the distribution of interest rates.

JEL Classification System: E52

*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, Cynthia Wu, and Jonathan Wright as well as the audiences at UBC, Stanford GSB, the Swiss National Bank’s Research Conference, German Economists Abroad, the AEA meetings, the Fed System Meeting on Macroeconomics, FRB Board-NY conference on Empirical Macroeconomics, NASMES, NBER Summer Institute, the Hoover MMCN conference, the CEBRA conference, and EFA for helpful discussions and suggestions. We also thank Emily Martell and Patrick Shultz for outstanding research assistance and Jens Christensen and Glenn Rudebusch for sharing their estimates. The views expressed here are solely those of the authors and do not necessarily represent those of the Federal Reserve Banks of New York and San Francisco or the Federal Reserve System.

[†]Federal Reserve Bank of San Francisco, 101 Market Street, Mailstop 1130, San Francisco, CA 94105; Thomas.Mertens@sf.frb.org.

[‡]Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045; John.C.Williams@ny.frb.org.

I 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. Yet, introducing a lower bound into the New Keynesian model leads to multiple equilibria that differ dramatically in the effectiveness of monetary policy. Therefore, empirical tests of the prevailing equilibrium are critically important for assessing the effects and implications of the lower bound. However, the relatively short period of time during which the lower bound has been relevant limits the ability to quantitatively assess this issue using macroeconomic data alone.

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 far-ahead expectations give a clearer view into the overall effects of the lower bound on the behavior of the economy that is not dependent on a particular realization of shocks or assumptions about short-run dynamics of the economy.

We use a standard New Keynesian model where inflation and output are jointly determined and are affected by expectations of future output and inflation. 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 output and 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

¹See, for example, Fuhrer and Madigan (1997), Reifschneider and Williams (2000), Eggertsson and Woodford (2003), Evans et al. (2015), Reifschneider (2016), Hamilton et al. (2016), and Cuba-Borda and Singh (2020).

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 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. The equilibria differ in expectations about future inflation and output and 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. 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 equilibria are harder to distinguish empirically.

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 sufficiently negative 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 binds 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 real rate of interest, r^* , 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 far-ahead forecast densities on future U.S. nominal interest rates and inflation. 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 Wright (2017) and Kitsul and Wright (2013)). Looking at a range of options with different strike prices, we reconstruct risk-neutral forecast densities of inflation and interest rates under the forward measure 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 interest rate wedge between the policy rate and the unconstrained shadow interest rate. 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. 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. This finding is striking in that short-term interest rates were near zero for seven straight years following the Great Recession of 2007-09.

We further provide evidence on changes in physical expectations by studying survey evidence from the Survey of Primary Dealers. We document that survey respondents' expectations of nominal interest rates declined during the period when the natural rate of interest fell. Consistent with the target equilibrium of our theoretical model, forecasts of inflation declined as well. While changes in risk premia would affect our

estimates from options prices, our survey evidence shows that changes in physical expectations are at least in part driving our results.

While 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, options data for far-ahead interest rates are scarce 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. Our earliest estimates date back to 2002, 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 the two equilibria of our New Keynesian model change in response to an unanticipated exogenous change in the natural rate of interest. Therefore, we neither generalize these findings to alternative models of monetary policy nor address the potential endogeneity or stochastic nature of the decline in the natural rate of interest.

Section II 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 III discusses the construction of forecast densities and tests the predictions from our theoretical model. Section IV discusses the robustness of the findings and section V concludes.

II 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 complications associated with short-run dynamics and illustrate more clearly the most important

theoretical implications of the lower bound for distributions of beliefs.

II.A Model Setup

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:

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

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:

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

where $\alpha > 0$, r^* is the long-run natural real rate of interest, and ϵ_t is a demand shock. All agents, including the central bank, 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:

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

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

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

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

policy and equilibrium conditions appear in Appendix A.A):

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

The current state of the economy is fully described by the realization of the two shocks. The expected value of inflation in the following period arises as an equilibrium outcome and is independent of current shocks due to the i.i.d. assumption. As a result, conditional and unconditional forecasts coincide and remain constant over time. We therefore drop time subscripts for expectations below. Maximizing the objective (3) subject to the equilibrium conditions (1) and (2) leads to the unconstrained optimal policy under discretion i_t that depends only on the current state

$$i_t = r^* + \psi\mathbb{E}\pi + \frac{1}{\alpha}\epsilon_t + \gamma\mu_t, \quad (5)$$

where $\psi \equiv 1 + \frac{1}{\alpha\kappa} - \frac{\lambda\beta}{\alpha\kappa(\kappa^2 + \lambda)} > 1$ and $\gamma \equiv \frac{\kappa}{\alpha(\kappa^2 + \lambda)}$. Solving the maximization problem while imposing the lower bound constraint on the policy rule leads to the optimal policy rule under discretion:

$$i_t = \max\{i_t, i^{LB}\}. \quad (6)$$

Note that the lower bound binds when $\gamma\mu_t + \frac{1}{\alpha}\epsilon_t \leq i^{LB} - r^* - \psi\mathbb{E}\pi$. To simplify notation, we combine the demand and markup shocks into the linear combination on the left-hand side of the inequality $\theta_t = \gamma\mu_t + \frac{1}{\alpha}\epsilon_t$ such that the lower bound binds if the combined shock θ_t falls below the cutoff $\bar{\theta}^{LB} \equiv i^{LB} - r^* - \psi\mathbb{E}\pi$, i.e. $\theta_t \leq \bar{\theta}^{LB}$.

Plugging in the optimal interest rate rule (6) into (4) leads to two different processes for inflation depending on whether the lower bound binds or not:

$$\pi_t = \begin{cases} \mu_t + \kappa\epsilon_t - \alpha\kappa(i^{LB} - r^*) + (1 + \alpha\kappa)\mathbb{E}\pi, & \text{if } \theta_t \leq \bar{\theta}^{LB} \\ \frac{\lambda}{\kappa^2 + \lambda} \{\mu_t + \beta\mathbb{E}\pi\} & \text{otherwise.} \end{cases} \quad (7)$$

If the lower bound does not constrain policy in the current period, optimal policy yields an inflation rate given by the second part of equation (7). Note that the unconstrained optimal policy fully offsets the demand shock ϵ_t . In the special case of $\lambda = 0$, this policy also offsets the markup shock, 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, each of the two parts in equation (7) can be solved for a steady-state values of π . In the one associated with the part where policy is unconstrained, 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 $\bar{\pi}^c = i^{LB} - r^*$ and inflation is thus below target. The only distinguishing features between the two equilibria are expectations about future inflation and output. We would need further assumptions to select between the two steady states.⁴

Monetary policy crucially depends on the equilibrium the economy is in. The optimal policy rule in equation (6) depends on inflation expectations. And because the two equilibria differ in inflation expectations, optimal policy will differ across the two equilibria. In a liquidity trap, the central bank would therefore prescribe lower interest rates than in the target equilibrium.

We now analyze the two equilibria in a stochastic environment. Once we introduce shocks, the two equilibria cannot be clearly linked to either part of equation (7) anymore. The lower bound is occasionally binding in both equilibria. Hence, within each equilibrium, the realization of shocks determines which part of the equation determines inflation. To study this situation, we define a key variable in our analysis, the probability of a binding lower bound P^{LB} , and link it to inflation expectations via

$$\begin{aligned} P^{LB} &= \text{Prob}(\theta \leq \bar{\theta}^{LB}) = \int_{\underline{\theta}}^{\bar{\theta}^{LB}} \phi(\theta) d\theta \\ &= \Phi(\bar{\theta}^{LB}) = \Phi(i^{LB} - r^* - \psi \mathbb{E}\pi), \end{aligned} \tag{8}$$

where $\underline{\theta} \in \mathbb{R}_- \cup \{-\infty\}$ is the lower bound of the support of the combined shock θ , $\phi(\theta)$ denotes its probability density function, and $\Phi(\theta)$ its corresponding cumulative density function. Since the cumulative density function is monotone, the probability of a binding lower bound is inversely related to inflation expectations.

⁴There 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.

If the lower bound always binds, i.e. $P^{LB} = 1$, the first part of equation (7) determines inflation and inflation expectations follow $\mathbb{E}\pi = i^{LB} - r^*$, just like in the deterministic model. If, on the other hand, $P^{LB} = 0$, the second part of the equation holds and inflation expectations are anchored at target. To compute expectations for intermediate cases, we take unconditional expectations of both sides of equation (7) to get

$$\mathbb{E}\pi = P^{LB} \mathbb{E} \left[\mu + \kappa \epsilon - \alpha \kappa (i^{LB} - r^*) + (1 + \alpha \kappa) \mathbb{E}\pi \mid \theta \leq \bar{\theta}^{LB} \right] + (1 - P^{LB}) \mathbb{E} \left[\frac{\lambda}{\kappa^2 + \lambda} (\mu + \beta \mathbb{E}\pi) \mid \theta > \bar{\theta}^{LB} \right]. \quad (9)$$

When the support of the combined shock θ encompasses the cutoff value, the lower bound binds occasionally. This situation arises both irrespective of whether the economy is near the target or liquidity trap equilibrium. When the realization of the shocks is sufficiently high such that $\theta_t > \bar{\theta}^{LB}$, the central bank is unconstrained 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.

The only distinguishing feature between the two equilibria are expectations about future inflation and, implicitly, the output gap. If the lower bound is expected to bind infrequently, the economy is in the target equilibrium and relatively unconstrained due to the forward-looking Phillips curve. Low inflation expectations, on the other hand, imply a lower level of interest rates on average, and therefore the lower bound is more likely to be a binding constraint. Since the cutoff value $\theta_t > \bar{\theta}^{LB}$ depends on inflation expectations, it also differs across the two equilibria.

To further investigate the relation between the probability of a binding lower bound and inflation expectations, we define the interest rate wedge $\Delta i_t = i_t - \bar{i}_t$ as the difference between the constrained (equation (6)) and the unconstrained (equation (5)) optimal policy rates. We define the unconstrained benchmark rate as having the form in equation (5) with inflation expectations that prevail under the constrained optimal policy rule. In this sense, it can be interpreted as a shadow interest rate. As a consequence the wedge between the actual and shadow rates is zero when the central bank is unconstrained and positive when the lower bound binds. With this definition, inflation expectations can be linked to the expected nominal and the natural interest rates via an unconditional version of the Fisher equation.

Proposition 1 (Below-target inflation expectations)

The unconditional mean of inflation equals the mean of nominal interest rates minus the natural rate of interest:

$$\begin{aligned}\mathbb{E}\pi &= \mathbb{E}i - r^* \\ &= -\frac{\alpha\kappa(\kappa^2 + \lambda)}{\kappa^2 + (1 - \beta)\lambda}\mathbb{E}\Delta i.\end{aligned}\tag{10}$$

Appendix A.A contains a derivation of equation (10).

The first part of the equations implies that the real interest rate, $i_t - \mathbb{E}_t[\pi_{t+1}]$, equals the natural real rate of interest r^* on average at each point in time.

The second part of the equation demonstrates a downward bias in average, and thus expected, inflation due to the lower bound, both in the target and the liquidity trap equilibrium whenever the lower bound might bind such that $\mathbb{E}\Delta i > 0$. In this case, average inflation will fall below its target rate.

When the lower bound binds occasionally, equation (1) tells us that expected inflation will be *below* target, and thus its deterministic benchmark, in the target equilibrium due a positive interest rate wedge for some shock realizations. In the liquidity trap equilibrium, occasionally binding constraints imply that the central bank can stabilize inflation in some states of the world. Inflation expectations, while being below target, are *above* their deterministic benchmark.

For a normal distribution of the markup shock, Figure 1 illustrates the two equilibria for different degrees of aggregate uncertainty.⁵ In particular, it shows how aggregate uncertainty lowers inflation expectations in the target equilibrium and raises them in the liquidity trap equilibrium. For a high level of aggregate uncertainty, there is a knife-edge case of a unique equilibrium, and no equilibrium emerges for higher levels of uncertainty. To produce Figure 1, we use the following parameter combination to illustrate the results:⁶ $\alpha = \kappa = 1$, $\beta = 0.99$, $i^{LB} = -0.5\%$, $r^* = 1\%$, and $\sigma_\epsilon = 0$. 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. For values of $\sigma_\mu > 1.9\%$, no equilibrium exists when $\lambda = 0$.

An important aspect of this analysis is that we assume that the lower bound always exists and that

⁵The equilibrium condition (9) for a steady state under the normal distribution is

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

where $v(\mathbb{E}_t\pi_{t+1}, i^{LB}, r^*) = (\kappa^2 + \lambda)(\mathbb{E}_t\pi_{t+1}(\alpha\kappa + 1) - \alpha\kappa(i^{LB} - r^*)) - \beta\lambda\mathbb{E}_t\pi_{t+1}$ and $\sigma_{\epsilon\mu} = \sqrt{2\sigma_\epsilon^2(\kappa^2 + \lambda)^2 + 2\kappa^2\sigma_\mu^2}$.

⁶Throughout 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.

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 et al. (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.

II.B Empirical Implications from the Theoretical Model

This section discusses testable predictions from the theoretical model that allow us to distinguish whether the economy is in the vicinity of a target or in a liquidity trap equilibrium. We then 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 exogenous. For a more concise notation, we define $\xi = 1 + \alpha\kappa + \frac{\alpha\kappa\beta\lambda}{\kappa^2 + (1-\beta)\lambda}$. Note that the previously stated assumptions on the ranges for the underlying parameters imply that $\xi > 1$.

As we saw in the previous section, the probability of a binding lower bound is a key variable of interest that can be linked to the distribution of inflation and interest rates. The following proposition shows how it evolves when the natural rate of interest, r^* , changes.

Proposition 2 (Comparative statics for the probability of a binding lower bound)

The probability of a binding lower bound evolves with the natural rate of interest according to

$$\frac{dP^{LB}}{dr^*} \left(P^{LB} - \frac{1}{\xi} \right) = \frac{1}{\xi} \phi(\bar{\theta}^{LB}(r^*)).$$

The critical value for the probability of a binding lower bound is:

$$\bar{P}^{LB} = \frac{1}{\xi} \in (0, 1).$$

For a proof, see Appendix A.B.

Proposition 2 presents a differential equation for the probability of a binding lower bound. There is a bifurcation at the critical value $\bar{P}^{LB} = 1/\xi$. At the critical value, the derivative of the lower bound probability with respect to the natural rate of interest is unbounded. If the probability of a binding lower bound is below the critical value, it corresponds to the target equilibria for different r^* . A higher probability is associated with the liquidity trap equilibria.

In the case where the natural rate of interest is far above the lower bound on interest rates, the probability of a binding lower bound is close to zero in the target equilibrium since nominal rates are high and the central bank has ample space to cut its policy rate. The probability of a sufficiently negative shock that requires a cut in the nominal rate down to the lower bound is small. In the liquidity trap, on the other hand, the lower bound almost always binds. Inflation expectations are far below the lower bound (close to $i^{LB} - r^*$) and the chance of a sufficiently positive shock to lift the nominal interest rate above the lower bound is small.

According to Proposition 2, reductions in the natural rate of interest lead to a fall in the probability of a binding lower bound in the liquidity trap equilibrium where P^{LB} is above the critical value \bar{P}^{LB} . In the target equilibrium, the lower bound probability lies below the critical value and reductions in r^* raise the lower bound. These results can be seen from the fact that the right-hand side of the equation is nonnegative and the term $P^{LB} - \frac{1}{\xi}$ determines the direction of the change. For a sufficiently small value of r^* , there is a unique equilibrium and no equilibrium for lower values, just like with the comparative statics with respect to uncertainty.

Figure 2 uses the example of a normal distribution to depict 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.

Next, consider the relationship between r^* and expectations of the interest rate. The following proposition shows that there is a direct link between r^* , the mean interest rate, and lower bound risk.

Proposition 3 (Change in level of nominal interest rates)

The probability of a binding lower bound determines how the level of nominal interest rates changes with the natural rate of interest

$$\frac{d\mathbb{E}i}{dr^*} = -\frac{1 - P^{LB}}{\xi P^{LB} - 1}. \quad (11)$$

See Appendix A.C for a proof.

According to Proposition 3, the expected interest rate changes in the same direction as the natural rate of interest if the lower bound probability is below the critical value of $1/\xi$, i.e. in the target equilibrium. In the liquidity trap, the sign is reversed as can be seen from the denominator.

Changes in the natural rate of interest not only affect the probability of a binding lower bound but also alter inflation expectations.

Proposition 4 (Change in inflation expectations)

The comparative statics for the lower bound translate into comparative statics for inflation expectations:

$$\frac{d\mathbb{E}\pi}{dr^*} = -\frac{1}{\psi} \frac{P^{LB}}{P^{LB} - \frac{1}{\xi}}.$$

See Appendix A.D for a proof.

Proposition 4 shows that the change of inflation expectations is positive in response to a fall in r^* in the liquidity trap equilibrium (where P^{LB} is above the critical value) and negative in the target equilibrium. 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 (10). In the stochastic economy associated with the liquidity trap equilibrium, a lower r^* raises inflation through the channel just discussed but then also lowers the cutoff value for the shock, thus limiting the range of shock realizations for which policy is unconstrained. As a result, expected inflation increases 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 inflation rate since policy is free to offset it. Interest rates therefore decline one-for-one with the decline in r^* . In the stochastic economy, a lower value of r^* lowers expected inflation, resulting in a higher cutoff value $\bar{\theta}^{LB}$, and thus increases the set of shocks for which policy is constrained. 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 markup 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 (10). The graphs show different steady-state outcomes

for a range of r^* values, where the blue line represents the target and the red line represents 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^* .

The lower bound on nominal interest rates not only affects average inflation and interest rates but, most profoundly, affects the asymmetry of the distributions. This asymmetry stems from the truncation of the distribution of interest rates that translates into an asymmetry of the distribution for inflation. This latter effect is due to the central bank's inability to stabilize inflation when constrained by the lower bound.

To investigate this channel further, we establish the link between changes in the cutoff value for the realization of shocks, at which point the lower bound starts to bind, and changes in the natural rate.

Lemma 1 (Changes in the cutoff $\bar{\theta}^{LB}$)

$$\frac{d(i^{LB} - r^* - \psi \mathbb{E}\pi)}{dr^*} = \frac{1}{\phi(\bar{x}^{LB}(r^*))} \frac{dP^{LB}}{dr^*} \quad (12)$$

See Appendix A.E for a proof.

Lemma 1 characterizes changes in the range of underlying shocks for which the lower bound will be binding. Specifically, the cutoff inherits its properties from the lower bound probabilities that we analyzed in Proposition 2. Therefore, when the natural rate of interest falls, Proposition 1 implies that the mean moves away (towards) the median in the target (liquidity trap) equilibrium. Furthermore, the distance between the median and a percentile lower than the probability of a binding lower bound moves away (towards) the equivalent measure on the right side of the distribution in the target (liquidity trap) equilibrium.

The intuition for these changes in the asymmetries is that, 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 a more asymmetric distribution. When the interest rate is constrained more frequently, the distribution of inflation shifts 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 natural rates of interest. Figure 4 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 to a change in the natural rate are exactly opposite in the target versus the liquidity trap equilibrium of our theoretical model. To test the predictions, we would ideally like to obtain moments of the unconditional forecast densities of interest rates and inflation. The approach taken in this paper to proxy the true forecast densities with measures we construct from options prices on long-dated options. Long maturities help in mitigating the risk that short-term transition dynamics do not materially affect the estimates.

III Empirical Analysis and Results

In the empirical analysis of this section, we compare forecast densities of interest rates and inflation derived from financial data to the unconditional distributions from our model.⁷ We exploit the evidence of a sizable decline in the natural rate of interest over the period that we study. 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 we compute risk-neutral forecast densities that might contain risk premia. We think the analysis is nevertheless revealing for two reasons. 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 forward 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. In a robustness exercise, we provide direct evidence from survey data that our results hold up for a more direct measures of physical expectations.

III.A Data and 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 the literature. We therefore keep the discussion brief and refer the interested reader to Appendix B for details. For interest rates, we use a parametric approach (see, for example, Wright (2017)). We parameterize the

⁷A 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.

underlying interest rate distribution by a mixture of two lognormal distributions. For inflation, we use the approach in Kitsul and Wright (2013) that uses a nonparametric approach.

Under either approach, the goal is to estimate an underlying distribution $\hat{g}(v_\tau)$ where, in our case, the underlying variable is either the interest rate or rate of inflation at time τ , i.e. $v_\tau = i_\tau, 1 + \pi_\tau$. The distribution is estimated to best fit a series of option prices $p_t(\tau, K_n)$ ($n = 1 \dots N$) on day t associated with different strike prices K_n and the date of maturity τ . Since the payoff to a call option is given by $X_n(v_\tau) = \max\{v_\tau - K_n, 0\}$ and that of a put option by $X_n(v_\tau) = \max\{K_n - v_\tau, 0\}$, the value of the option by taking the expectations under the risk-neutral distribution $g(v_\tau)$, i.e.,

$$p_t(\tau, K_n) = e^{-y_{t,\tau}(\tau-t)} \mathbb{E}^g[X_n(v_\tau)] = e^{-y_{t,\tau}(\tau-t)} \int X_n(v_\tau) g(v_\tau) dv_\tau. \quad (13)$$

$y_t(\tau)$ thereby denotes the discount rate between times t and τ .

The risk-neutral distribution $g(v_t)$ is taken under the forward measure and can thus contain risk premia. It differs from the risk-neutral measure in that payoffs are discounted with a $(\tau - t)$ -period bond instead of the future short-term interest rate (see Kitsul and Wright (2013)).

With a set of prices $p_t(\tau, K_n)$ for each day t , we extract a forecast density for each day that matches the prices and thus obtain a time series of forecast densities on a daily frequency. That is, on each day we extract the market-implied forecast densities for interest rates and inflation at various horizons. Since we compare these forecast densities to the unconditional distributions in our theoretical model, we primarily focus on long-term forecasts.

III.A.1 Forecast densities for interest rates

We obtain a daily data series of caps on the London Interbank Offered Rate (LIBOR) from Bloomberg that is consistently available between January 1, 2007, up to February 28, 2020. All empirical tests with our estimated forecast densities are carried out over this sample period. The main advantage of using these options is the long horizon of forecasts. 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 10-year interest rate cap on three-month LIBOR will receive a payment at the end of every three-month period over the following 10 years that LIBOR exceeds the strike at the beginning of the same three-month period. If the cap is written for a period of 10 years, and τ is quarterly, there are 39 potential payoffs made at quarterly horizons starting in month six up to 40 quarters out. The value of the cap

is the sum of all its caplets.

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. We use caps of strikes 1 through 14 percent, in terms of LIBOR, and horizons of up to 10 years. We compute the prices of caplets by subtracting a shorter maturity cap from the price of a cap with longer maturity.

We extract forecast densities for the three-month interest rate seven years out. Therefore, we parameterize the underlying risk-neutral distribution by a mixture of two lognormal distributions

$$\hat{g}(i_t; \omega) = \omega \rho(i_t; \mu_1, \sigma_1) + (1 - \omega) \rho(i_t; \mu_2, \sigma_2) \quad (14)$$

where $\rho(i_t; \mu_i, \sigma_i) = \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{1}{2} \frac{(\log i_t - \mu_i)^2}{\sigma_i^2}\right)$ and $\omega = \{\phi, \mu_1, \sigma_1, \mu_2, \sigma_2\}$. i_t refers to the three-month (LIBOR) interest rate in the future. With the two means and standard deviations, as well as the mixing weight ω , there are five parameters we can use to fit the distribution. From this distribution, we compute implied prices via equation (13).

For a given day, we extract caplet prices from underlying caps for different strike prices. We then find parameter combinations for the distribution to minimize the fitting errors between observed prices and implied prices. As a result, we obtain long-term forecast densities on a given date. We repeat the procedure for each day in our sample and thus extract a time series of forecast densities. Appendix B.A contains details.

III.A.2 Inflation forecasts

We collect daily data on inflation caps and floors from Bloomberg. An inflation cap pays if the average (annually compounded) consumer price index (CPI) inflation exceeds the strike rate. Trading in inflation caps started in late 2009 in over-the-counter markets and started to become reasonably liquid by 2011. Because liquidity sharply dropped and the market trading essentially ceased by the end of 2016, our sample period ranges from April 1, 2011, to December 31, 2016. We focus on contracts of maturities five and 10 years to extract the five-year forward rate five years out. We complement our analysis of forecast densities for inflation with estimates of forward rates from inflation swaps. Appendix B.B contains details about the data and the construction of forward rates and forecast densities.

The seller of a zero-coupon inflation cap promises to pay a fraction $\max\{(1 + \pi_{t,\tau})^{\tau-t} - (1 + K_n)^{\tau-t}, 0\}$ of a notional underlying principal as a single payment in $\tau - t$ years time, where $\pi_{t,\tau}$ denotes the average annual total CPI inflation from t to τ and K_n denotes the strike rate of the cap. 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) and do not specify a functional form for the underlying distribution. We apply this 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 five-year and 10-year inflation consistent with observed prices, as given in equation (13). We therefore approximate implied volatilities from both caps and floors across strike prices at a given maturity by local regressions. 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 under the forward measure.

We convert our estimates of the distributions for five-year and 10-year inflation rates to five-year forward rates five years out. We therefore use a methodology that converts percentiles of the five-year and 10-year forecast densities into percentiles of the forward rate distribution. This methodology delivers exact results if inflation rates follow an AR(1) process, including special cases of i.i.d. shocks and random walks. A random walk for year-over-year inflation rates has been found to be a reasonable description of the time series for inflation (see Atkeson and Ohanian (2001) and Faust and Wright (2013)). Appendix B.B shows details about the exact calculations for forward rates and proves that the methodology is exact for AR(1) processes. 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.

III.B The Decline in the Natural Rate of Interest

Considerable empirical evidence suggests that the longer-run natural rate of interest in the United States has declined notably over the past several decades (Williams (2017)). Figure 5 shows monthly estimates of the natural interest rate for 1998 to 2020, an updated time series from the model in Christensen and Rudebusch (2019). 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. There appears to be a break in the series of estimates around the end of 2011, with the mean estimate dropping from 1.21 percent over 2007-2011 to 0.56 percent over 2012-2020.

To assess statistical significance of the decline in the natural rate of interest, we test the following null hypothesis.

Null hypothesis: *The natural rate of interest r^* did not change across the two subsamples of 2007-2011 and 2012-2020.*

To test this null hypothesis, we compute the t-statistic

$$t = \frac{\bar{r}_l^* - \bar{r}_e^*}{\sqrt{\text{Var}[\bar{r}_l^* - \bar{r}_e^*]}}$$

where \bar{r}_e^* and \bar{r}_l^* refer to the average of daily estimates of r^* over the early and late subsamples, respectively. The standard deviation in the denominator corrects for serial correlation in the sample (see Wilks (1997) and Appendix C.A). The null hypothesis that r^* has the same means across the two subsamples is clearly rejected at the 1% level with a t-statistic of -11.41.

We assume that a change in the natural rate of interest is permanent. Standard estimates of the natural rate in the literature correspond to the low-frequency component of interest rates that reflects highly persistent influences that are exogenous to the standard New Keynesian model, such as trend productivity growth, demographics, and global demand for safe assets.⁸ In particular, the estimates reported here, from Christensen and Rudebusch (2019), correspond to real interest rates expected to prevail five to 10 years in the future. Similarly, Holston, Laubach, and Williams (2017) show that their estimates of the natural rate of interest are nonstationary. For a more detailed discussion see Williams (2017).

Our theoretical model of Section II suggests that a decline in the natural rate of interest affects the distributions of inflation and nominal interest rates. Rational forecasts should reflect these changes to the distributions. Therefore, we use the decline in the natural rate of interest as a source of variation to study the behavior of forecast densities.

III.C Implied Distribution for Interest Rates

The upper panel in Figure 6 shows a 20-day moving average of daily estimates 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 10 percent while confirming the presence and importance of a lower bound at the lowest percentiles. The green line in the upper panel of Figure 6, representing the bound on the 15th percentile, inches towards the lower bound during the sample period. The 2.5th percentile consistently lies at the lower bound and the mass at the lower bound is increasing over time as the natural rate falls. Both the upper panel and the lower left panel in Figure 6 demonstrate that the average long-term forecast of interest rates has decreased over time. The lower right-hand panel shows the difference of

⁸See, for example, 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 et al. (2017). Estimates from these models for r^* show a decline consistent with that from the measure used in this paper.

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. The difference in percentiles on the upper end stays consistently above the one on the lower end with the gap widening towards the end of the sample indicating an increasing asymmetry.

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, average interest rate forecasts fall 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. And fourth, the skewness of interest rates increased over that subsample.

Table 1: Summary of long-term interest rate moments

	2007-2011	2012-2020
Mean	4.80%	2.88%***
Median	4.54%	2.37%***
Std. deviation	5.54%	3.96%***
Skewness	0.16%	0.45%***

Difference in moments statistically significant at * 10%, ** 5%, *** 1% level (adjusted for serial correlation).

The patterns in Table 1 are consistent with the target equilibrium of our theoretical model and at odds 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 increases. This experiment allows us 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 target equilibrium, the skewness of interest rates is positive. In our model, 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.

All of these changes in the moments of the forecast densities for interest rates are highly statistically significant. This statistical significance is derived from a test that investigates the null hypothesis that a given moment has not changed across the two subsamples. We perform a one-tailed t-test, adjusted for serial correlation as detailed in Appendix C.A, for each moment. The null hypothesis can be rejected at the 1% level, as highlighted in Table 1.

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.

III.D A Lower Bound Indicator

In our theoretical model, the expected interest rate wedge is the driving force behind average interest rates and inflation as well as asymmetries of their distributions. Equation (10) makes this link precise for average inflation. In the following, we refer to this statistic as the *lower bound indicator*. The lower bound indicator is defined as how much more the central bank would have liked to cut interest rates on average,

$$\mathcal{J}_t = \mathbb{E}[\Delta i_t]. \tag{15}$$

Guided by the theoretical insights, we obtain a measure of the lower bound indicator $\hat{\mathcal{J}}_t$ for each month from the data. Consistent with our theoretical model when we use normal distributions for shocks, we assume a generalized rectified Gaussian distribution for interest rates (see Palmer, Hill and Scheduling (2017)). This distribution truncates the left tail of the normal distribution and lumps the truncated mass at the lower bound. We estimate the mean, standard deviation, and lower bound on a monthly frequency.⁹ Using these three estimates for the forecast density observed within each month, we generate a time series of the lower bound indicator $\hat{\mathcal{J}}$. Details about the construction of the lower bound indicator are listed in Appendix C.B.

Figure 7 shows that the lower bound indicator captures 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. The regression coefficients of the lower bound indicator on the natural rate of interest and the asymmetry of forecast densities are significant at the 1% level, using Newey-West adjusted standard

⁹An alternative specification where we impose a constant lower bound and only fit the mean and variance leads to very similar results.

errors. This is not a test of the theory but rather shows the information contained in the lower bound indicator.

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 increase with lower levels of r^* in the target equilibrium, they should also increase with the lower bound indicator. This hypothesis is borne out by the data with a statistical significance at the 1% level. Since the sign of the relationship would be reversed in the liquidity trap equilibrium, these findings provide further evidence in favor of the target equilibrium.

Taken together, we conclude that the theoretically motivated lower bound indicator captures the asymmetries of interest rate distributions due to the lower bound. We can therefore use it to test the theoretical predictions of our model and assess statistical significance of the results. Performing the test for the natural rate of interest shows that there is a significant link between declines in r^* and the lower bound indicator.

III.E Implied Distribution for Inflation

The previous two sections demonstrate that the evolution of the distribution of interest rates are consistent with the model of Section II. Here, we extend the empirical analysis to the theoretical predictions for inflation.

As for interest rates, the time series of percentiles for forecast densities of inflation provide a first look at whether the predictions of the model are consistent with the data. The upper panel in Figure 9 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 9, 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 a 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 2012-2016. The density of inflation is slightly skewed to the left in both subsamples. The broader message of a decline in average inflation forecasts is consistent with the target equilibrium of our theoretical model. We also test this prediction from our model using five-year forward rates on inflation five years out that we obtain from inflation swaps. Consistent with estimates from caps and floors, the forward rate for inflation

fell during the later subsample.

Table 2: Summary of inflation moments

	Forecast densities		Inflation swaps	
	2011	2012-2016	2007-2011	2012-2020
Mean	2.60%	2.42%	2.96%	2.45%***
Median	2.62%	2.43%		
Std. deviation	2.28%	1.44%***		
Skewness	-0.04%	-0.05%		

Difference in moments statistically significant at * 10%, ** 5%, *** 1% level (adjusted for serial correlation).

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 Great Recession. The skewness of inflation is negative in both subsamples, consistent with the theory. Its decline during the later subsample is consistent with the prediction of both equilibria. The same t-test as for interest rates indicates statistically significant results for inflation swaps. There, the drop in forward rates for inflation from 2.96 percent to 2.45 percent is significant at the 1% level. The time series for forecast densities for inflation is too short to have a sufficient amount of power to test the difference in moments across the two subsamples.

Figure 10 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 target equilibrium of our theoretical model. The regression coefficient is significant at the 5% level, using Newey-West adjusted standard errors. The right panel shows again that the changes in the asymmetry of the inflation distribution are less pronounced than for the interest rate distribution. The skewness of inflation is negative and decreases with higher levels of the lower bound indicator, consistent with the more negative skewness in the later subsample in Table 2. All of 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.

Consistent with this evidence from forecast densities, Figure 11 shows a decrease in the swaps-implied five-year forward rate five years ahead for higher values of the lower bound indicator. The regression is in line with the left panel in Figure 10 and statistically significant at the 5% level.

To check whether risk or liquidity premia are the primary drivers of the results, we perform two robustness checks. Section IV.A shows that the same results emerge for more liquid, shorter-term options for interest

rates. In section IV.B, we obtain survey evidence and show that the observed patterns in the changes of forecast densities for inflation and interest rates are unlikely to stem from changes in risk premia alone. The evidence presented there suggests that physical expectations seem to have played a substantial role in the pricing of interest rate and inflation derivatives.

III.F Further Evidence on the Link between Forecast Densities and r^*

The mechanism highlighted in this paper links changes in the natural rate of interest to changes in the distribution, and thus forecast densities, of interest rates and inflation. Thus, the most natural test that also allows us to assess statistical significance of these relationships is to run regressions of moments of the forecast densities on r^* . To this end, we use daily observations of r^* and forecast densities for the entire time series over which we obtain forecast densities. We run regressions of the form

$$m_t = \beta_0 + \beta_1 r_t^* + \varepsilon_t, \tag{16}$$

where m_t refers to a moment of the interest rate or inflation forecast densities.

Table 3: Regressions of long-term interest rate and inflation moments on r^*

	<i>Dependent variable:</i>				
	mean interest rate	skew interest rate	mean inflation (densities)	skew inflation (densities)	mean inflation (swaps)
r^*	2.882*** (0.683)	-0.434*** (0.071)	0.404 (0.348)	-0.068** (0.028)	0.700*** (0.230)
cons.	1.285*** (0.478)	0.690*** (0.058)	2.184*** (0.316)	-0.006 (0.022)	2.077*** (0.277)
R^2	0.789	0.450	0.076	0.040	0.422

Note: *p<0.1; **p<0.05; ***p<0.01. Using Newey-West standard errors.

Table 3 shows the results of these regressions for long-term interest rates. Regression coefficients for interest rate moments are statistically significant at the 1% level, using Newey-West adjusted standard errors. The fact that the regression coefficient for average forecast densities is above one is consistent with the target equilibrium as well. As we established earlier, average forecasts for interest rates fall more than one-for-one with the natural rate, as indicated by the coefficient of 2.88 on r^* . Also consistent with the target equilibrium,

the skewness of interest rates rises when r^* falls, although this finding is statistically not significant.

The regression coefficients for inflation forecast densities are also in line with the target equilibrium. Average inflation falls less than one-for-one with the natural rate, whether measures were derived from forecast densities or swaps. While there is limited power due to the relatively short sample for forecast densities for inflation, the regression coefficient for inflation swaps is significant at the 1% level. The skewness of inflation displays a negative regression coefficient on r^* , meaning that the skewness tends to become less negative when the natural rate of interest falls. While this estimate points to a different comovement on a daily frequency than the evidence from the two subsamples and the lower bound indicator, it is consistent with the target equilibrium of our theoretical model where the relationship between the skewness of inflation and r^* is U-shaped.

III.G Discussion

The results speak very clearly: Qualitatively, the target equilibrium of the New Keynesian model is consistent with the changes in forecast densities considered here. We do not find any 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 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 point 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 the link between interest rates and inflation well 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 binds. 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 et al. (2012), Reifschneider (2016), and Wu and Zhang (2019)). 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.

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.

IV Robustness

IV.A 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 longer time period and liquidity for these contracts tends to be high. With the forecast horizon of 24 months we use here, the data 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. With a sample period from April 4, 2002, until February 28, 2020, we extend our analysis to other time periods. Of particular interest is the time period between 2002-2004 when nominal interest rates neared the lower bound. Analogously to long-term forecast densities, we parameterize the distribution by a mixture of two lognormal distributions. A detailed discussion of the construction of forecast densities can be found in Appendix B.C.

We start with an analysis of subsamples analogous to those in Section III.C. Consistent with the evidence presented before, forecast densities for short-term interest rates shifted down when the natural rate of interest fell. This result can be seen by comparing the first two columns in Table 4. The fall in the median and mean are again statistically significant at the 1% level. The skewness is again positive and rose when the natural rate of interest fell.

Table 4: Summary of moments from short-term interest rate forecast densities

	2002-2011	2012-2020	2002-2004	2005-2007	2009-2016	2017-2020
Mean	3.41%	1.57%***	3.73	4.75	1.35	2.25
Median	3.05%	1.36%***	3.47	4.78	1.07	2.28
Std. deviation	3.04%	1.36%***	3.33	3.73	1.52	1.41
Skewness	0.27%	0.32%*	0.29	-0.04	0.53	-0.04

Difference in moments statistically significant at * 10%, ** 5%, *** 1% level (adjusted for serial correlation).
Significance tests only computed between first two columns.

Prior to the Great Recession the probability of hitting the lower bound on interest rates was considered to be small. These forecasts have changed substantially as average interest rates declined. As a result, the forecast densities have become asymmetric. These predictions are in line with the target equilibrium of our theory and at odds with the liquidity trap equilibrium.

Table 4 also shows moments of short-term forecast densities for interest rates for various subsamples. In particular, early in the 2000s, average forecasts were around 3.73 percent when the federal funds rate was close to the lower bound. In the mid-2000s, before the Great Recession, average forecasts had risen relative to previous levels.

The time series for short-term forecast densities in Figure 12 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 liftoff 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 for short-term forecast densities are also consistent with the evidence in Section III.C. The lower left panel of Figure 12 shows a decline in the mean during the time the natural rate of interest decreased, consistent with the picture in long-term forecasts. The graph in the lower right-hand panel shows the increase in 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.

Table 5 reports regressions of the form in equation (16) as in Section III.F. They show that, consistent with evidence for the longer term, average forecasts for interest rates in the near term on average fell with the natural rate of interest while the distribution became more symmetric. The table also shows that the

distribution shifted down at every percentile. The coefficients and R^2 are thereby larger for higher percentiles. The lower percentiles were constrained by the lower bound for some time during the sample, thus limiting the explanatory power for those percentiles.

Table 5: Regressions of short-term interest rate moments on r^*

	<i>Dependent variable:</i>								
	mean	skew	2.5	5	15	50	85	95	97.5
r^*	2.269*** (0.713)	-0.024 (0.101)	0.856*** (0.144)	0.991*** (0.207)	1.369*** (0.444)	2.121*** (0.724)	3.174*** (0.426)	3.995*** (0.366)	4.497*** (0.264)
cons.	0.398 (0.715)	0.321** (0.130)	-0.246*** (0.082)	-0.208 (0.135)	-0.045 (0.373)	0.360 (0.723)	0.774 (0.626)	1.151** (0.470)	1.418*** (0.301)
R^2	0.422	0.001	0.211	0.208	0.206	0.324	0.553	0.732	0.801

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Using Newey-West standard errors.

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 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.

IV.B Robustness to Physical versus Forward Measure: Survey Data

A common concern with using financial market data is that estimates might be driven by changes in risk premia rather than changes in expectations. They can, instead, reflect movements in risk premia. To address this concern, we obtain data from the Survey of Primary Dealers. This survey collects information eight times per year from primary dealers. Among other questions, the survey asks about expectations for the average federal funds rate over the next 10 years and the long-run federal funds rate. It also collects information about the distribution of inflation. Appendix C.C contains details about the data used in this section.

We first investigate the link between average expected federal funds rates under the physical measure and our estimates of the natural rate of interest. We therefore use survey responses on the average federal funds rate over the following 10 years. Our sample period ranges from June 2012 until March 2020. We run a regression with the mean of the responses as the dependent variable and the estimate of the natural rate of interest on the survey date as the independent variable, analogous to regressions in Section III.F.

Column 1 of Table 6 shows the estimates from the regression. The regression coefficient shows that

average federal funds rates fell statistically significantly (at the 5% level) with r^* . It suggests that decreases in the natural rate of interest comove close to one-for-one with estimates of the federal funds rate over the following 10 years. Column 2 studies the analogous link between the long-run federal funds rate and r^* . The results for the analysis are consistent with those for the 10-year average of the federal funds rate. Both results are only consistent with the target equilibrium and at odds with the liquidity trap equilibrium.

Table 6: Regressions of Survey Responses from the Survey of Primary Dealers on r^*

	<i>Dependent variable:</i>		
	10-year FFR	Long-run FFR	inflation expect.
r^*	0.887** (0.390)	0.914* (0.535)	0.051** (0.023)
cons.	2.012*** (0.286)	2.705*** (0.450)	2.088*** (0.025)
R ²	0.365	0.019	0.114

Note: *p<0.1; **p<0.05; ***p<0.01. Using Newey-West standard errors.

For inflation, we collect answers about the forecast densities of the five-year five-year forward rate of CPI inflation from the Survey of Primary Dealers. The sample period ranges from January 2011 until March 2020. We compute the average expected forward rate and run the regression analogous to those for interest rates.¹⁰

Column 3 shows the regression results for inflation. Average forecasts for inflation fall under the physical measure just as they did under the forward measure, consistent with the target equilibrium in our theoretical framework. The regression coefficient, thereby, is statistically significant at the 5% level.

In sum, survey evidence is consistent with the findings from derivatives prices suggesting that a significant portion of options prices are driven by changes in physical expectations when the natural rate of interest declines.

V 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

¹⁰Further details on the construction of the dependent variables are listed in Appendix C.C.

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 analytically derive the relationships between the natural rate of interest and the unconditional distributions of interest rates and inflation.

Our empirical evidence strongly supports the hypothesis 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.

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.
- Atkeson, Andrew and Lee E. Ohanian**, “Are Phillips Curves Useful for Forecasting Inflation?,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2001, 25 (1), 2–11.
- 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,” *The Review of Economics and Statistics*, 2019, 101 (5), 933–949.
- 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.
- Cuba-Borda, Pablo and Sanjay R. Singh**, “Understanding Persistent ZLB: Theory and Assessment,” *mimeo*, UC Davis, 2020.
- 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.
- Faust, Jon and Jonathan H. Wright**, “Forecasting Inflation,” in “Handbook of Economic Forecasting,” Vol. 2A of *Handbook of Economic Forecasting*, Elsevier, 2013, pp. 2 – 56.
- Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig**, “Deflation Risk,” *The Review of Financial Studies*, 2017, 30 (8), 2719–2760.
- 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.
- **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), 1063–1070.
- 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.
- Palmer, Andrew W., Andrew J. Hill, and Steven J. Scheding**, “Methods for Stochastic Collection and Replenishment (SCAR) optimisation for persistent autonomy,” *Robotics and Autonomous Systems*, 2017, 87, 51–65.
- 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.
- **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.
- **and –**, “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.
- Wilks, Daniel S.**, “Resampling Hypothesis Tests for Autocorrelated Fields,” *Journal of Climate*, 1997, 10 (1), 65–82.
- Williams, John C.**, “Heeding Daedalus: Optimal Inflation and the Zero Lower Bound,” *Brookings Papers on Economic Activity*, 2010, 2010 (2), 1–37.
- , “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.
- Wu, Jing Cynthia and Ji Zhang**, “A shadow rate New Keynesian model,” *Journal of Economic Dynamics and Control*, 2019, 107.

A The New Keynesian Model with a Lower Bound

A.A Model Setup

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

$$\pi_t - E_t \pi_{t+1} = \mu_t + \kappa(\epsilon_t - \alpha(i_t - E_t \pi_{t+1} - r^*)) + \beta E_t(\pi_{t+1} - \pi_{t+2}). \quad (17)$$

Under optimal monetary policy with discretion and i.i.d. shocks, the final term is zero, so we are left with the equation:

$$\pi_t = (1 + \alpha\kappa)E_t \pi_{t+1} + \mu_t + \kappa\epsilon_t - \alpha\kappa(i_t - r^*), \quad (18)$$

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

To derive the Fisher equation, take the unconditional expectation of equation (18) and observe that $E\pi_{t+1} = E\pi_t = E\pi$ due to the temporary nature of shocks that does not require state variables for the model. Solving for inflation expectations results in the Fisher equation

$$E\pi_t = Ei_t - r^*.$$

Plugging in that $i_t = \mathbf{i}_t + \Delta i$ and using the fact that $Ei_t = r^* + \left(1 + \frac{1}{\alpha\kappa} - \frac{\lambda\beta}{\alpha\kappa(\kappa^2 + \lambda)}\right)E\pi_t$, we get

$$E\pi = -\frac{\alpha\kappa(\kappa^2 + \lambda)}{\kappa^2 + (1 - \beta)\lambda}E\Delta i.$$

A.B Proof of Proposition 2

Proof: As defined in the main text, the cutoff for the shocks at which the lower bound starts to bind, $\bar{\theta}^{LB}(r^*)$, is the realization of the combined shock such that the unconstrained policy coincides with the lower bound $\mathbf{i}_t = i^{LB}$, i.e. $\bar{\theta}^{LB}(r^*) = i^{LB} - r^* - \psi E\pi$. Furthermore, recall the lower bound for the shock θ_t is $\underline{\theta} \in \mathbb{R}_- \cup \{-\infty\}$. Since there are no endogenous state variables, we drop the time subscripts for a more concise notation.

We start from the definition of the probability of a binding lower bound

$$P^{LB} = \int_{\underline{\theta}}^{i^{LB} - r^* - \psi E\pi} \phi(\theta) d\theta.$$

Using this definition, we compute its derivative via Leibniz's rule

$$\frac{dP^{LB}}{dr^*} = - \left(1 + \psi \frac{d\mathbb{E}\pi}{dr^*} \right) \phi(\bar{\theta}^{LB}(r^*)).$$

Now combine the previous expression with equation (20)

$$\frac{dP^{LB}}{dr^*} = - \left(1 - \xi \frac{d\mathbb{E}\Delta i}{dr^*} \right) \phi(\bar{\theta}^{LB}(r^*)). \quad (19)$$

Having linked the change in the probability to changes in the expected interest rate wedge, we show that it can be linked instead to the probability of a binding lower bound directly.

The following lemma links changes in the expected interest rate wedge to the probability of a binding lower bound.

Lemma 2 (Expected wedge for nominal interest rates)

$$\frac{d\mathbb{E}[\Delta i]}{dr^*} = \frac{P^{LB}}{\xi P^{LB} - 1}$$

whenever $P^{LB} \neq \frac{1}{\xi}$.

Proof: Applying Leibniz's rule to the average wedge in interest rates, we get

$$\frac{d\mathbb{E}[\Delta i]}{dr^*} = \left(-1 - \psi \frac{d\mathbb{E}\pi}{dr^*} \right) \int_{\underline{\theta}}^{\bar{\theta}^{LB}(r^*)} \phi(\theta) d\theta = -P^{LB} \left(1 + \psi \frac{d\mathbb{E}\pi}{dr^*} \right).$$

Using (20), write the previous condition as

$$(1 - \xi P^{LB}) \frac{d\mathbb{E}[\Delta i]}{dr^*} = -P^{LB}$$

and the lemma follows. □

Plugging the equation from Lemma 2 into equation (19) proves the first part of the proposition. The critical value is obtained by setting the second factor on the left-hand side to zero. □

As a direct consequence of Proposition 1, we show that inflation expectations and expected wedges in nominal interest rates have to change in opposite directions in response to a change in r^* .

Corollary 1 (Change in inflation expectations)

$$\psi \frac{d\mathbb{E}\pi}{dr^*} = -\xi \frac{d\mathbb{E}\Delta i}{dr^*}. \quad (20)$$

Proof: Differentiating equation (10) from Proposition 1 with respect to the natural real rate of interest, we obtain Corollary 1. \square

A.C Proof of Proposition 3

Proof: To prove Proposition 3, we start from the decomposition of nominal interest rates into the shadow rate and the interest rate wedge

$$\frac{d\mathbb{E}i}{dr^*} = \frac{d\mathbb{E}(\Delta i + i)}{dr^*} = \frac{d\mathbb{E}\Delta i}{dr^*} + \frac{d\mathbb{E}i}{dr^*}.$$

Now we use the definition of the shadow rate i and take unconditional expectations on both sides to get

$$\frac{d\mathbb{E}i}{dr^*} = 1 + \psi \frac{d\mathbb{E}\pi}{dr^*} = 1 - \xi \frac{d\mathbb{E}\Delta i}{dr^*},$$

where the second equation follows from Corollary 1. Plugging this expression into the decomposition for the nominal rate results in

$$\frac{d\mathbb{E}i}{dr^*} = 1 + (1 - \xi) \frac{d\mathbb{E}\Delta i}{dr^*}.$$

Replace the change in the expected interest rate wedge from Lemma 2 and the proposition emerges. \square

A.D Proof of Proposition 4

Proof: Use Corollary 1 and replace the expected interest rate wedge by the expression from Lemma 2. Proposition 4 emerges. \square

A.E Proof of Lemma 1

Proof: Start from the definition of the probability of a binding lower bound

$$P^{LB} = \int_{\underline{\theta}}^{\bar{\theta}^{LB}(r^*)} \phi(\theta) d\theta$$

and differentiate the expressions on both sides with respect to r^* using Leibniz's rule

$$\frac{dP^{LB}}{dr^*} = \frac{d\bar{\theta}^{LB}(r^*)}{dr^*} \phi(\bar{\theta}^{LB}(r^*)).$$

As a result, the lower bound probability and the cutoff change in the same direction. □

B Appendices on Construction of Forecast Densities

B.A Extraction of Long-Term Forecast Densities for Interest Rates

This section describes the details of the algorithm for the extraction of long-term forecast densities for interest rates used in this paper. It is analogous to other approaches used in the literature (see, e.g. Wright (2017)). To be able to use equation (13) for option prices, we need to obtain a discount rate $y_{t,\tau}$ as well as data on option, or caplet, prices.

We start with the discount rate. Since we obtain forecast densities seven years out, we need to construct a discount curve for at least that time period. We use two different sets of data for the short end and one for the long end. For the short end, we collect daily data on the three-month LIBOR rate from the St. Louis Fed via FRED and Eurodollar futures with quarterly maturities from one to eight quarters out via Bloomberg to take advantage of the high liquidity in these contracts.¹¹ For the longer end, we use interest rate swaps with maturities 1 to 10 years, again obtained from Bloomberg.¹²

To get a discount rate, we start by constructing zero coupon bond yields on the short end. Therefore, we fit a Svensson yield curve through the Eurodollar futures and extract the zero coupon bond yields. Next, we fit a Svensson yield curve through the swap rates on the longer end and extract the zero coupon yields from it. Therefore, notice that the discounted cashflows on the fixed (with rate $i_{t,\tau}^{\text{fixed}}$) and floating legs of a swap with T payments are equated such that

$$i_{t,\tau}^{\text{fixed}}(Z_{t,t+1} + Z_{t,t+2} + \dots + Z_{t,\tau}) = i_t Z_{t,t+1} + \mathbb{E}^g[i_{t+1}]Z_{t,t+2} + \dots + \mathbb{E}^g[i_{\tau-1}]Z_{t,\tau},$$

where the price for a zero-coupon bond between time t and τ is denoted by $Z_{t,\tau}$ and the expectation of i_s under the distribution g is $\mathbb{E}^g[i_s]$ is the risk-neutral expectation of the interest rate in period s , i.e. the forward rate. Now observe that the forward rate can be expressed by zero-coupon bond prices via $\mathbb{E}^g[i_s] = Z_s/Z_{s+1} - 1$.

¹¹Ticker symbols are EDx Comdty where x needs to be replaced by the number of quarters out.

¹²We use the following ticker symbols USSWx Curncy where x should be replaced by the number of years out.

As a result, we obtain the zero-coupon bond prices from swaps via

$$Z_{t,\tau} = \frac{1 - i_{t,\tau}^{\text{fixed}} \sum_{s=t+1}^{\tau-1} Z_{t,s}}{1 + i_{t,\tau}^{\text{fixed}}},$$

starting from the short end of the curve where the zero-coupon price is known. Lastly, fitting a Svensson curve through the zero coupon yields delivers the discount rate for the seven-year time horizon. Furthermore, we can compute forward rates from this curve.

Having obtained the discount rate, we turn to extracting the forecast density. Therefore, we use interest rate caps data from Bloomberg where we obtain the full matrix of caps with maturities of 1 to 10 years and strikes from 1% to 10%.¹³ We choose a start date of January 1, 2007, due to missing data before this date.

We first compute the prices for caplets that are used on the left-hand side of equation (13). Therefore, we compute a Nelson-Siegel curve through the flat volatilities for each strike price on each date.

Next, we convert the caps value (quoted in terms of a flat volatility) into a spot volatility. The flat volatility $\sigma_{t,\tau}^{\text{cap}}$ is the volatility that equates the value of the sum of all caplets evaluated under the flat volatility with the sum of all caplets evaluated under the spot volatility σ_t^{cap}

$$\sum_{s=t+1}^{\tau} \text{Caplet}_s(K_n, \sigma_{t,\tau}^{\text{cap}}) = \sum_{s=t+1}^{\tau} \text{Caplet}_s(K_n, \sigma_{s-1}^{\text{cap}}). \quad (21)$$

We start with the very short end of the curve where the flat and spot volatilities are equal, i.e. $\sigma_{t,t+2} = \sigma_{t+1}$. The reason for this equality stems from the fact that between period t and payoffs in $t + 2$, only the interest rate in $t + 1$ is unknown. With the short end fixed, we can now solve for spot volatilities forward by solving equation (21) for the next time horizon. Using the Black pricing formula, we obtain the corresponding value of the caplet $\hat{p}_t(\tau, K_n)$. Equipped with these caplet prices for different strikes, we now aim to extract a forecast density that best matches the implied prices.

We parameterize the underlying distribution by a mixture of two lognormal distributions \hat{g} of equation (14). In a robustness test, we change the functional form of the distribution to a rectified Gaussian distribution. This functional form truncates the normal distribution and concentrates the truncated mass of the distribution at the truncation point. In that case, we fix the lower bound at zero and estimate the mean and the standard deviation.

¹³Ticker symbols on Bloomberg are USCFx0y where x corresponds to the strike and y to the number of years out. y should be replaced by X for the 10 year cap with a strike of 10%. We use the Bloomberg default as the pricing source and supplement it with CMPN data where data is missing.

To find the parameters of the distribution for a given day, and thus the forecast density, we minimize the sum of squared pricing errors over the parameters of the distribution

$$\hat{\omega}^* = \arg \min_{\omega} \sum_{n=1}^N (p_t(\tau, K_n) - \hat{p}_t(\tau, K_n))^2 + w \left(\mathbb{E}^{\mathcal{S}}[i_{\tau}] - \mathbb{E}^{\hat{\mathcal{S}}}[i_{\tau}] \right)^2, \quad (22)$$

where we obtain $\mathbb{E}^{\mathcal{S}}[i_{\tau}]$ from forward rates, $\mathbb{E}^{\hat{\mathcal{S}}}[i_{\tau}] = \int i_{\tau} \hat{g}(i_{\tau}; \omega) di_{\tau}$ is computed from the estimated distribution, and w is chosen to take a value of ten such that the mean of the distribution is aligned with forward rates. We solve this minimization using a global optimization routine. Specifically, we loop over all dates and employ for each daily observation a generalized simulated annealing algorithm that does not require any initial guesses. For the rectified Gaussian distribution, the set of parameters we are extracting is limited to the mean and the standard deviation of the underlying normal distribution.

Lastly, we compute the time series of forecast densities from the parameters for the mixture of lognormal distributions or rectified Gaussian distributions, respectively.

B.B Construction of Inflation Forecast Densities

This section lays out the method we use to extract inflation forecast densities and the data we use for our analysis. We complement our analysis of forecast densities for inflation by studying forward rates from inflation swaps. Therefore, we obtain daily data for the five-year forward rate five years out by taking twice the 10-year swap rate and subtracting the five-year swap rate.¹⁴ The sample period from January 1, 2007, up to February 28, 2020, for this data matches that of long-term forecast densities for interest rates.

As for long-term forecast densities for interest rates, we obtain daily data. To extract forecast densities for inflation, we follow the methodology in Kitsul and Wright (2013) and twice differentiate the pricing formula for inflation derivatives.

To discount expected payoffs in equation (13), we use GSW yields (see Gürkaynak, Sack and Wright (2007)) available on the website of the Board of Governors of the Federal Reserve System. We compute the forward rate from inflation swaps available on Bloomberg.¹⁵

We furthermore collect data on inflation caps and floors at maturities of five and 10 years from Bloomberg. The strike rates for caps are between -1% and 6% (with 0.5% increments) and between -3% and 4.5% for floors.¹⁶

¹⁴We obtain the data from Bloomberg with ticker symbols USSWIT5 Curncy and USSWIT10 Curncy. The construction of the forward rate follows the calculation in Bloomberg for the five-year-five-year forward rate available via FWISUS55.

¹⁵The ticker symbols are USSWIT5 Curncy and USSWIT10 Curncy.

¹⁶The ticker symbols for these contracts are USIZCx and USIZFx where x specifies the strike and m the maturity.

We obtain a daily time series over the time period from January 1, 2007, up to December 31, 2016. Due to a sharp decline in liquidity in the inflation caps and floors markets, we end the sample period at the end of 2016 despite Bloomberg providing quotes beyond this date.

Using the Black formula with the respective strike rates, the forward rate from the inflation swap, and the interest rate from the GSW yield, we compute implied volatilities for caps and floors on a given day.

To obtain an estimate of the implied volatility around a given strike, we run a local linear regression between implied volatilities and strike prices. Using the local linear regression line, we take the second derivative of Black's formula for the option price to get an estimate of the forecast density at a given strike. We repeat this procedure for different strikes to back out the forecast density for a particular day. We repeat these calculations for all days in our sample to end up with the time series of forecast densities for five-year and 10-year inflation.

Lastly, we turn these estimates into annualized five-year five-year-forward rates. We define the forward rate of inflation between h_1 and h_2 periods out as

$$f_{t+h_1,t+h_2} = \frac{1}{h_2 - h_1} \sum_{s=h_1+1}^{h_2} \pi_{t+s}, \quad (23)$$

where π_t denotes the rate of inflation as measured by the change in log prices. We denote the standard deviation of the forward rate at time t by σ_{f,t,h_1,h_2} .

The method to extract the percentiles of the distribution for forward rates works as follows: With the estimates of the percentiles of the five-year and 10-year inflation rates from the data, we compute the mean, standard deviation, and the daily time series of Z-scores corresponding to each percentile. The Z-score, here written for forward rates, is defined as the number of standard deviations a percentile is distanced from the mean

$$Z_f^p = \frac{\hat{f}_{t+h_1,t+h_2}^p - \mathbb{E}_t[f_{t+h_1,t+h_2}]}{\sigma_{f,t,h_1,h_2}}. \quad (24)$$

We then match the Z-scores of the forward rate distribution with those from the five-year and 10-year distributions that we calculate for the various percentiles from the analogous definitions to equation (24). With estimates of the mean and standard deviations for forward rates, we recover the percentiles from the Z-scores. Note that this procedure allows for asymmetric distributions where the Z-scores for percentiles on either side of the mean can differ in magnitude.

More specifically, we first estimate the mean of the forward-rate distribution via

$$\mathbb{E}_t[f_{t+5,t+10}] = 2\mathbb{E}_t[\bar{\pi}_{t,10}] - \mathbb{E}_t[\bar{\pi}_{t,5}],$$

where $\bar{\pi}_{t,h}$ refers to the annualized inflation rate over the next h years. This relation holds since the 5-year inflation rate and the forward rate average out to the 10-year inflation rate. We further compute the standard deviation via

$$\sigma_{f,t,5,10} = (4\text{Var}_t[\bar{\pi}_{t,10}] + \text{Var}_t[\bar{\pi}_{t,5}] - 4\text{Cov}_t[\bar{\pi}_{t,5}, \bar{\pi}_{t,10}])^{\frac{1}{2}}.$$

To compute the standard deviation, we choose a random walk for inflation such that the covariance term in the previous equation equals the variance of the five-year inflation rate. Note that only the calculation of the standard deviation of the forward rate is affected by this choice. We apply estimates of Z-scores from the five-year and 10-year distributions to back out percentiles of the forward rate. Therefore, we match the Z-score of the forward rate with an average of the estimated Z-scores, i.e. $\hat{Z}_f^p = (Z_{\bar{\pi},5}^p + Z_{\bar{\pi},10}^p)/2 \equiv \bar{Z}_{\bar{\pi}}^p$. Plugging this expression into (24), we obtain the p -th percentile

$$\hat{f}_{t+5,t+10}^p = \mathbb{E}_t[f_{t+5,t+10}] + \bar{Z}_{\bar{\pi}}^p \sigma_{f,5,10}.$$

In the following, we show that this procedure precisely recovers the distribution of the forward rate when inflation follows an AR(1) process with mean reversion ρ

$$\pi_{t+1} = \pi_\infty + \rho(\pi_{t-1} - \pi_\infty) + \nu_{t+1}, \quad (25)$$

where π_∞ denotes the unconditional mean of inflation and ν_t is normally distributed with mean zero and variance σ_ν^2 . In particular, a random walk ($\rho = 1$) is a common choice for the year-over-year inflation rate. Inflation at time $t + h$ is given by

$$\pi_{t+h} = \pi_\infty + \rho^h(\pi_t - \pi_\infty) + \sum_{s=1}^h \rho^{h-s} \nu_{t+s},$$

and thus, based on time- t information, conditionally normally distributed with mean $\pi_\infty + \rho^h(\pi_t - \pi_\infty)$ and variance

$$\text{Var}_t[\pi_{t+h}] = \sum_{s=1}^h \rho^{2(h-s)} \sigma_\nu^2.$$

Lemma 3 (Z-scores of forward rates)

Let inflation follow an AR(1) process with mean reversion $\rho \in [0, 1]$. The Z-scores Z_f^p for all percentiles of the forward rate between h_1 and h_2 periods out, $\hat{f}_{t+h_1,t+h_2}^p$, coincide with the Z-scores of the average inflation rates for horizons h_1 and h_2 for all percentiles p , independent of ρ .

Proof: Given its definition in equation (23), the forward rate $f_{t+h_1,t+h_2}$ is normally distributed provided that inflation follows an AR(1) process, as in (25).

The p -th percentile of the forward rate is $\hat{f}_{t+h_1,t+h_2}^p = \mathbb{E}_t[f_{t+h_1,t+h_2}] - \sqrt{2}\sigma_{f,t,h_1,h_2} \text{erfc}^{-1}[2p]$ due to its normal distribution. erfc thereby refers to the complementary error function and erfc^{-1} to its inverse. Plugging the expression for the p -th percentile into the definition of the Z-score in (24), we obtain

$$Z_f^p = -\sqrt{2}\text{erfc}^{-1}[2p]. \quad (26)$$

Note that the Z-score only depends on the percentile p and is independent of the mean of the forward rate, its standard deviation, and mean reversion ρ . The average inflation rate over the next h periods, $\bar{\pi}_{t,h} = \frac{1}{h} \sum_{s=1}^h \pi_{t+s}$, is normally distributed. Following the same logic as for the forward rate, the Z-scores of the average inflation rates corresponding to the p -th percentile of the distribution, $Z_{\bar{\pi},h}^p$, is thus also described by $-\sqrt{2}\text{erfc}^{-1}[2p]$, as in equation (26). As a result, we get $Z_f^p = Z_{\bar{\pi},5}^p = Z_{\bar{\pi},10}^p$. \square

Lemma 3 shows that the Z-scores for the forward rate coincide with those for the five-year and 10-year rates. As a result, we can recover the percentiles of the forward rate from either the Z-score or a linear combination thereof. The method used to extract forward rates thus delivers exact results for any AR(1) process when the standard deviation of the forward rate is computed as

$$\sigma_{f,5,10} = (4\text{Var}_t[\bar{\pi}_{t,10}] + (1 - 4\rho^5)\text{Var}_t[\bar{\pi}_{t,5}])^{\frac{1}{2}}.$$

B.C Short-term Forecast Densities for Interest Rates

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 two years out. We focus on the 24-month horizon. Due to the lack of data availability, our sample period is limited to the sample

period starting on April 4, 2002, and ending on February, 28, 2020.

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. Quarterly contracts are available for the nearest 16 quarterly months.

The forecast densities of the three-month LIBOR from Eurodollar futures options are extracted, analogous to long-term forecast densities, by specifying a mixture of two lognormal distributions for the three-month interest rates and performing nonlinear optimization to obtain the parameters. In our description of the algorithm, we follow the outline of Appendix B.A.

We start by collecting data on the discount rate as well as the forward rate. For the former, we again use GSW yields as in the extraction of inflation forecast densities. For forward rates, we use the data on Eurodollar futures described in Appendix B.A.

We obtain a set of Eurodollar futures options from the Chicago Mercantile Exchange (CME). We collect call and put prices with strikes from \$91 and above in 1/8 increments to just below \$100 for maturities between 21 and 27 months. To update the data set and fill in for missing data, we supplement this data with prices obtained from Bloomberg.¹⁷ We only use out-of-the-money call and put prices that mature at a horizon closest to 24 months out in our estimation. We end up with a time series of option prices that is consistent from April 4, 2002, on.

To set up the implied prices from our underlying distribution, we again use equation (13) for calls and puts. Under the parameterized distribution, we obtain implied prices $\hat{p}_t(\tau, K_n)$ with which we are trying to match the set of observed market prices. Therefore, we again perform the global minimization of equation (22) for prices observed on each day, using a weight of one.¹⁸

C Appendices on Empirical Tests

C.A Adjusting for Serial Correlation

Wilks (1997) adjusts the standard error for serial correlation of various lags ρ_k . Therefore, we first compute

¹⁷Ticker symbols for the data follows the convention ED(month)(year)(call/put) (strike) Comdty. Months are labeled as H (March), M (June), U (September), and Z (December). Thus, a Eurodollar futures call expiring in March 2022 with a strike of 99.875 would have the ticker EDH2C 99.875 Comdty.

¹⁸Using the same value for the weight as for long-term densities leads to consistent results.

the scaling factor

$$\vartheta = 1 + 2 \sum_{k=1}^{N-1} \left(1 - \frac{k}{N}\right) \rho_k,$$

with N being the sample size. Second, the approach adjusts for a potential bias in estimating the scaling factor by specifying the modified scaling factor

$$\hat{\vartheta} = \vartheta e^{2\frac{s}{N}}.$$

Applying this method to data for an early, X^e , and late, X^l , subsample, we compute the variance of the difference of means via

$$\text{Var} [\bar{X}^l - \bar{X}^e] = \hat{\vartheta}^l \frac{s_l^2}{N_l} + \hat{\vartheta}^e \frac{s_e^2}{N_e},$$

where s^2 refers to the sample variance of the subsample.

C.B Construction of Lower Bound Indicator

We assume a generalized rectified normal distribution for interest rates that concentrates the truncated mass at the lower bound (see Palmer, Hill and Scheduling (2017)). The most natural parameterization of our theoretical model of Section II with a normal distribution of shocks would result in such a 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 monthly time series by averaging each percentile within the month. We take the time series of the median as a measure for the mode of the underlying 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 $1/(\sqrt{2}\text{erfc}^{-1}(\frac{3}{10}))$. The 2.5th percentile serves as the lower bound. With these inputs we compute the lower bound indicator as the expected interest rate wedge where the underlying shadow interest rate follows a distribution with these parameters.

C.C Survey Evidence

We use data on expectations of average 10-year federal funds rates, long-run federal funds rate, and forecast densities. The data are accessible from the Federal Reserve Bank of New York's website. We use answers to the following question: "In addition, provide your estimate of the longer run target federal funds rate and your expectation for the average federal funds rate over the next 10 years." Respondents can then input a number under "Longer run" and one under "Expectation for average federal funds rate over next 10 years". For inflation data, we use responses to the question "For the outcomes below, provide the percent chance* you attach to the annual average CPI inflation rate from April 1, 2025 - March 31, 2030 falling in each of the following ranges. Please also provide your point estimate for the most likely outcome." (here from the March 2020 survey) where respondents give a distribution according to pre-specified bins.

Figures

Steady States for Normal Distribution

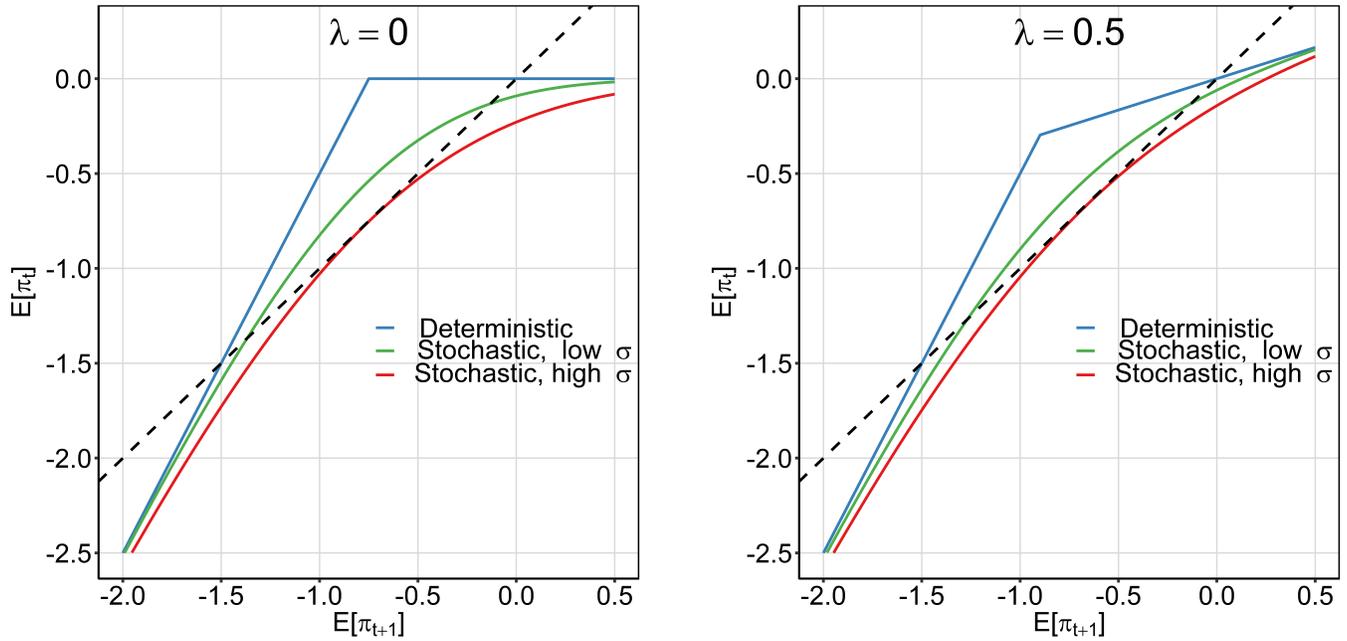


Figure 1: Expected inflation in the current period as a function of expected inflation in the following period assuming a normal distribution. Parameter values are set to $\alpha = 1$, $\kappa = 1$, $\beta = 0.99$, $r^* = 1\%$, and $i^{LB} = -0.5\%$. There are no ϵ -shocks, i.e., $\sigma_\epsilon = 0$. Intersections with the dashed line, the 45-degree line, represent steady states. Graphs 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 $\lambda = 0.5$.

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

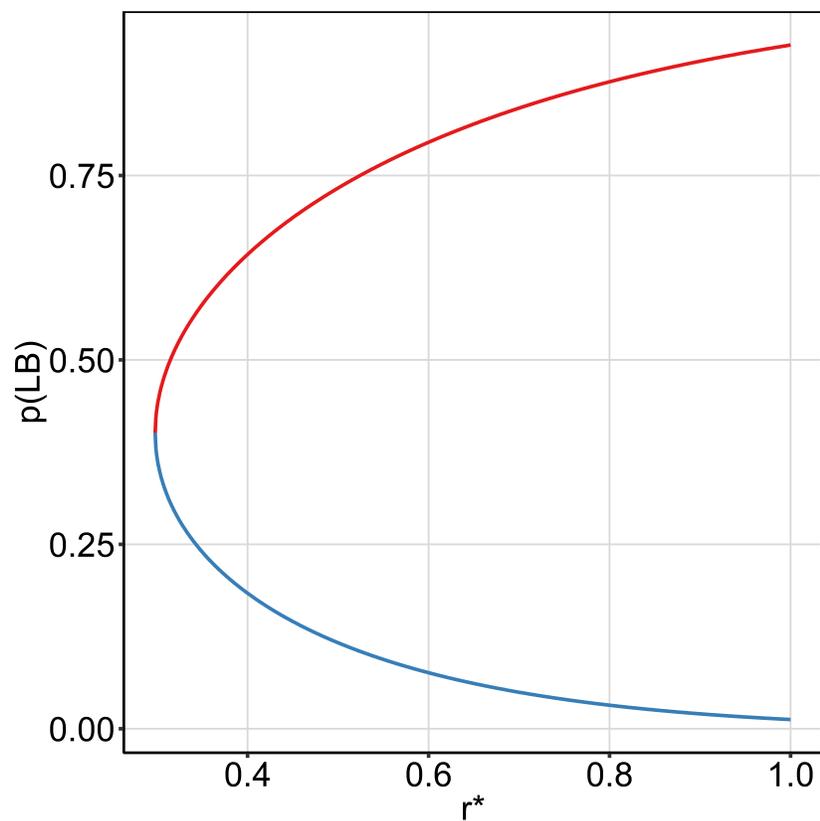


Figure 2: As r^* decreases, the probability of hitting the lower bound rises in the target equilibrium (blue line). In the liquidity trap (red line), the probability of being constrained falls. Parameterization is as in Figure 1, again with $\sigma_\mu = 1\%$ and $\lambda = 0.5$.

Effect of r^* on mean interest rates and inflation

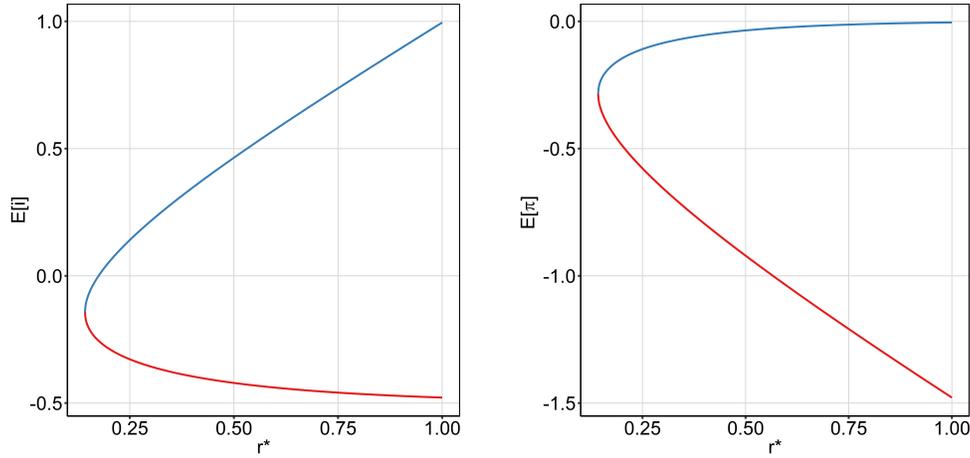


Figure 3: The blue lines show average interest rates (left panel) and average inflation (right panel) for various levels of r^* in the target equilibrium; the red lines display the analogues for the liquidity trap equilibrium. Parameterization is as in Figure 1 with $\sigma_\mu = 1\%$ and $\lambda = 0.5$. 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 skewness of interest rates and inflation

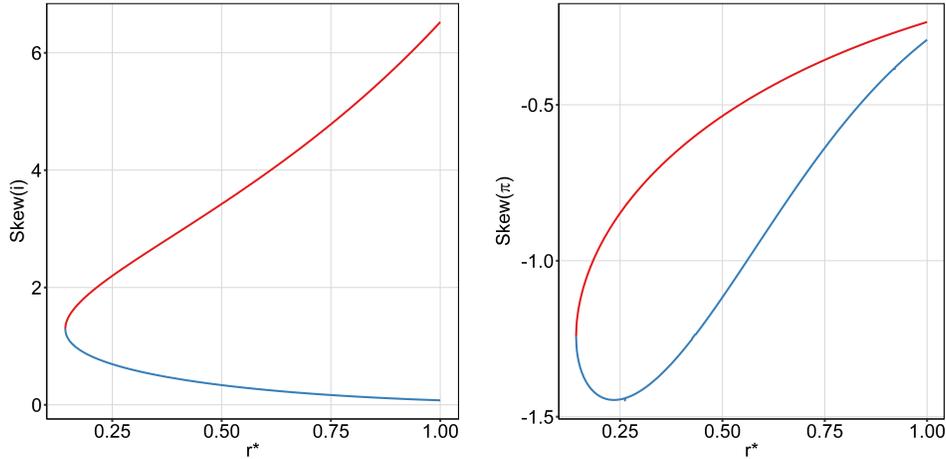


Figure 4: As r^* decreases, the distribution for interest rates becomes more asymmetric in the target equilibrium (blue line, left panel) while the skewness of inflation is U-shaped (blue line, right panel). In the liquidity trap equilibrium (red lines), the skewness of interest rates and inflation falls with r^* . Parameterization is as in Figure 1.

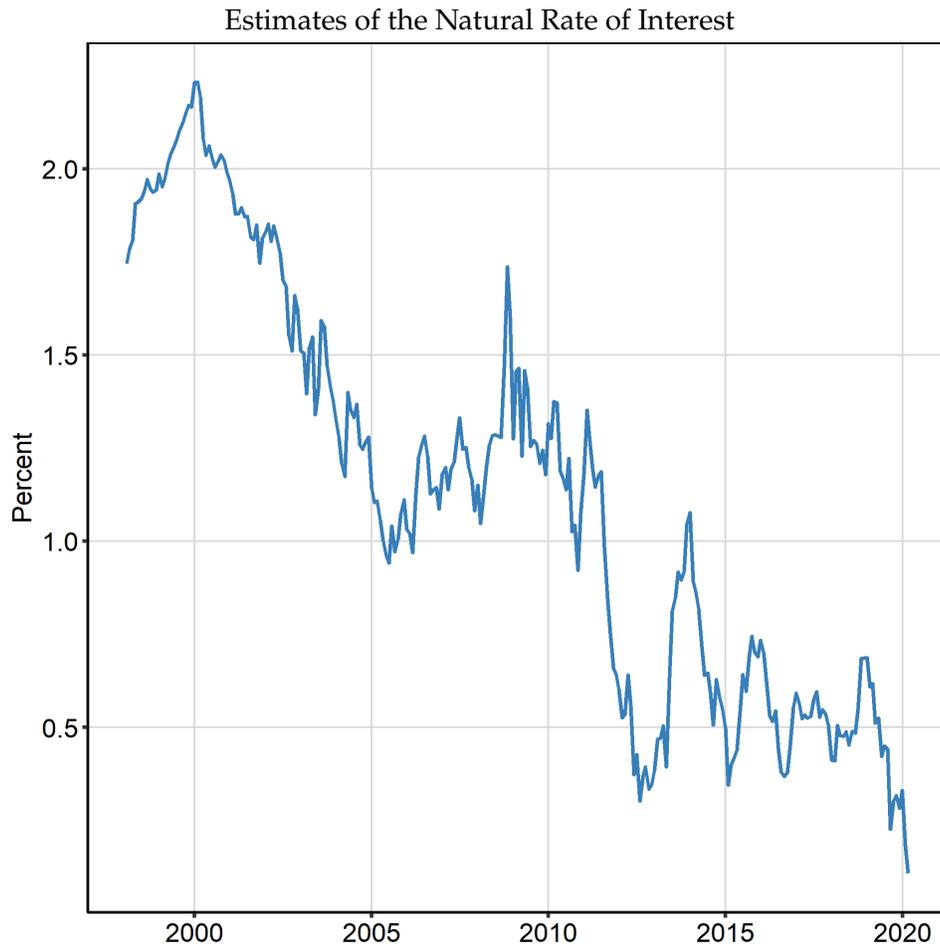


Figure 5: Monthly estimates of the natural rate of interest from Christensen and Rudebusch (2019).

Long-term Forecast Densities for Interest Rates

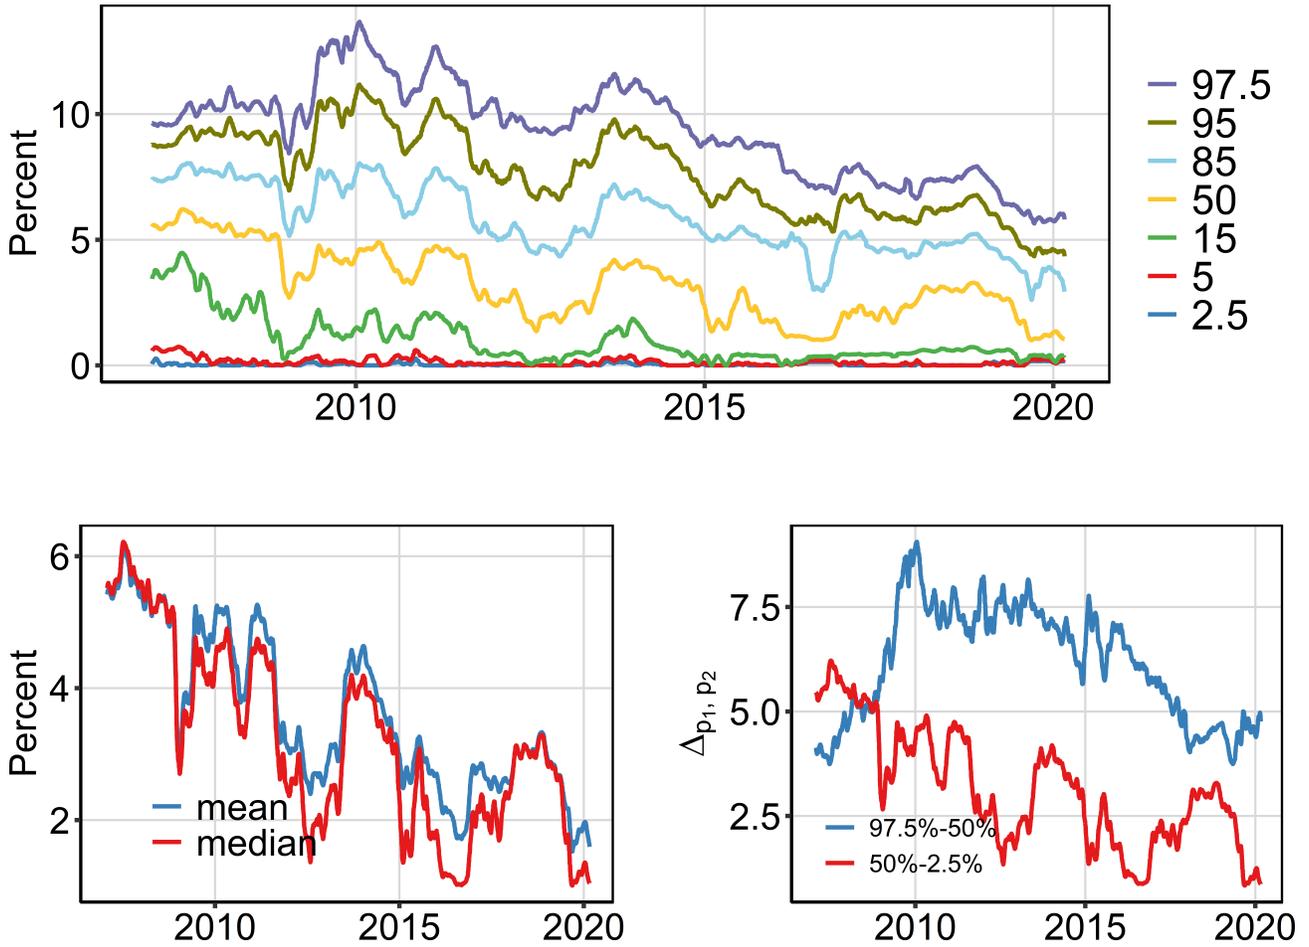


Figure 6: Long-term forecast densities for interest rates measured by the three-month forward rate seven years out. The upper panel shows a 20-day moving average of the daily time series for various percentiles. The lower panels plot 20-day moving averages for the mean and median of forecast densities 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).

Lower Bound Indicator and Interest Rate Asymmetry

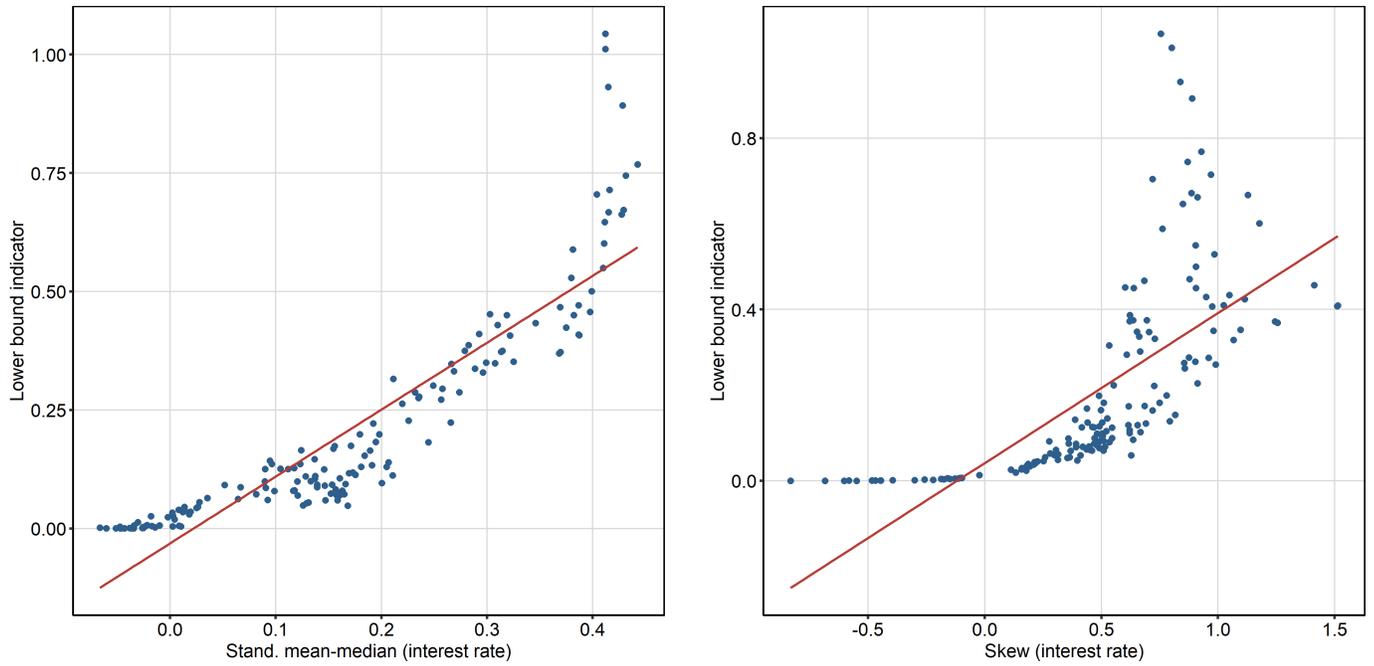


Figure 7: Relationship between the lower bound indicator and the asymmetry of interest rates. This figure shows two measures of asymmetry in interest rate forecast densities, the difference between mean and median as well as the skewness and how they relate to the lower bound indicator at a monthly frequency. The regression line in the left panel has an intercept of -0.03 (0.04) and a slope of 1.41 (0.25) where standard errors are Newey-West adjusted. The slope is significant at the 1% level. In the right panel, the intercept is 0.04 (0.05) and the slope 0.35 (0.07). The slope is again significant at the 1% level.

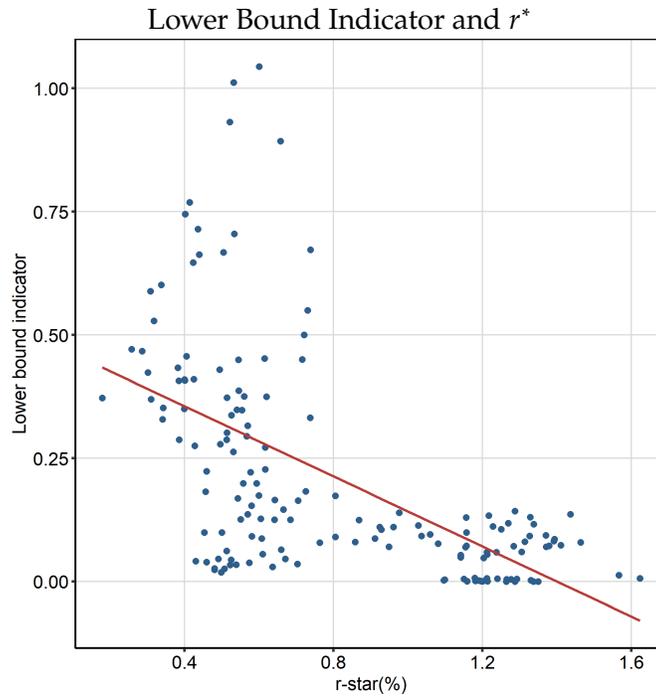


Figure 8: Relationship between the lower bound indicator and r^* . Lower levels of the natural rate of interest are associated with higher measures of our lower bound indicator. The graph displays monthly data. The regression line has an intercept of 0.49 (0.13) and a slope of -0.36 (0.10) where standard errors are Newey-West adjusted. Both are significant at the 1% level.

Long-term Forecast Densities for Inflation

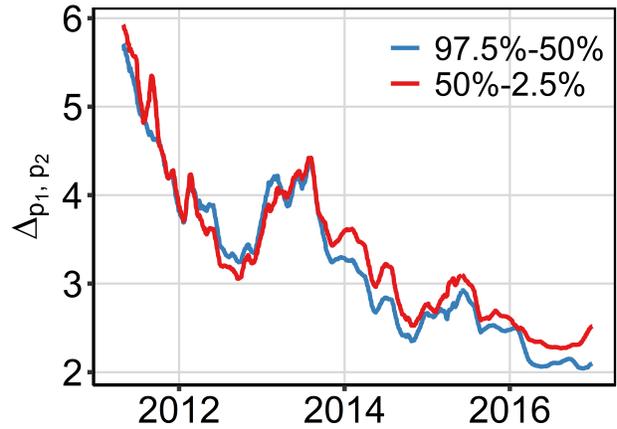
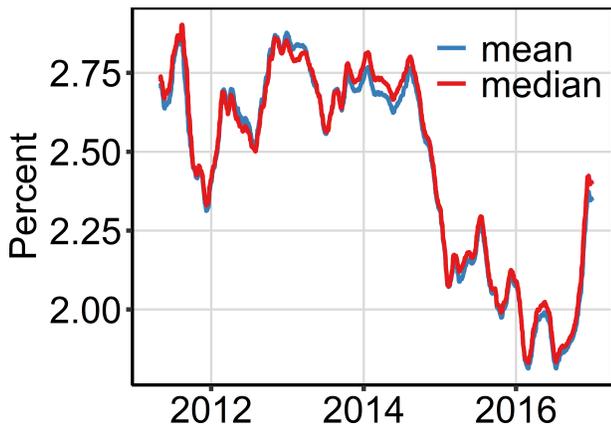
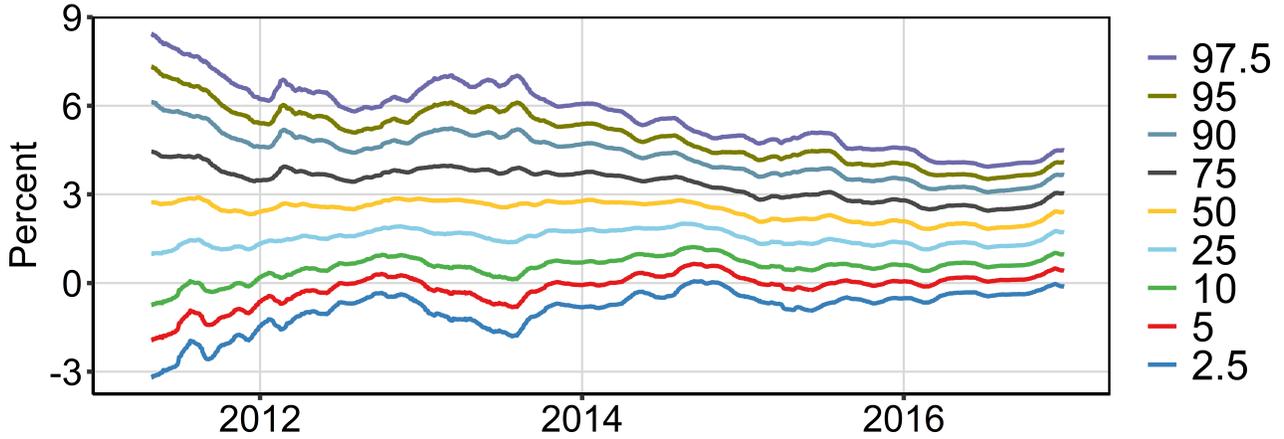


Figure 9: Long-term forecast densities for inflation. This figure reports measures of forecast densities analogous to Figure 6 for inflation rates measured by the five-year forward rate five years out. All pictures show 20-day moving averages of the underlying daily time series.

Lower Bound Indicator and Inflation Forecast Densities

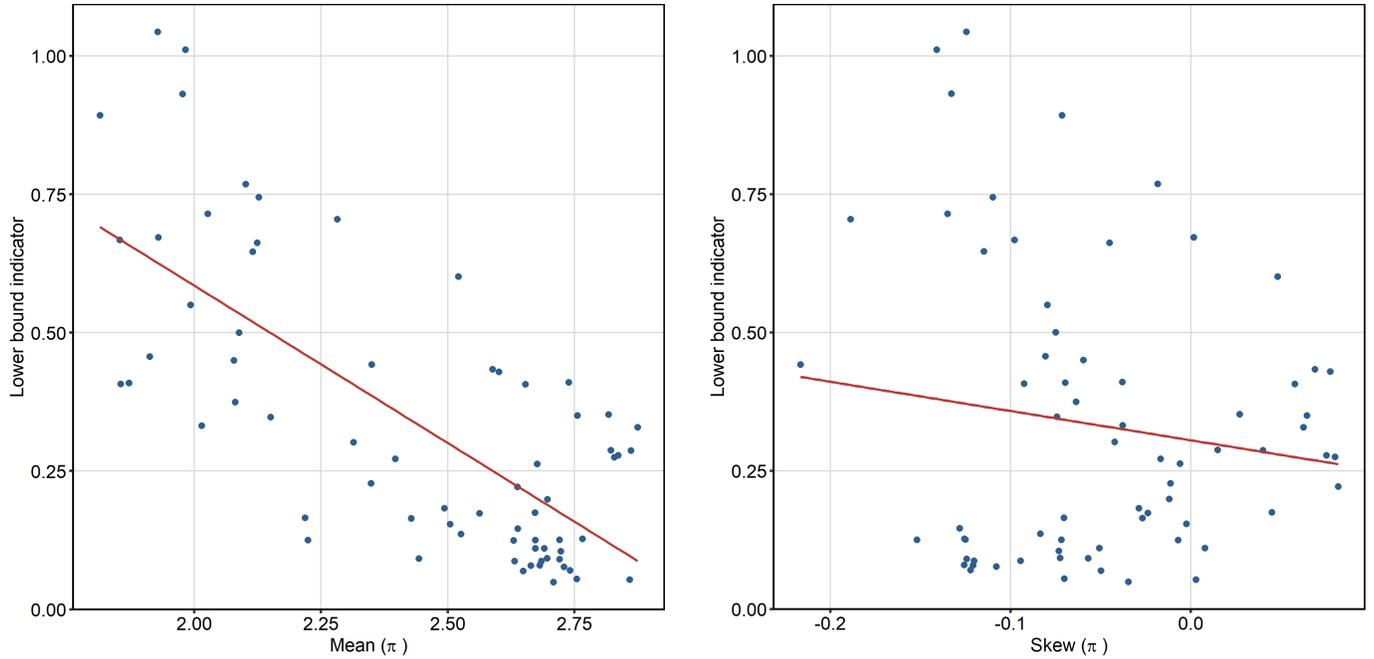


Figure 10: Relationship between lower bound indicator and inflation moments at a monthly frequency. The scatter plots show monthly data of mean inflation versus the lower bound indicator (left panel) and the skewness in inflation forecast densities (right panel). The regression line in the left panel has an intercept of 1.72 (0.58) and a slope of -0.57 (0.23) where standard errors are Newey-West adjusted. The intercept is significant at the 1% level and the slope at the 5% level. In the right panel, the intercept is 0.31 (0.06) and the slope -0.53 (1.16). The intercept is statistically significant at the 1% level while the slope is statistically insignificant.

Lower Bound Indicator and Inflation Rates from Swaps

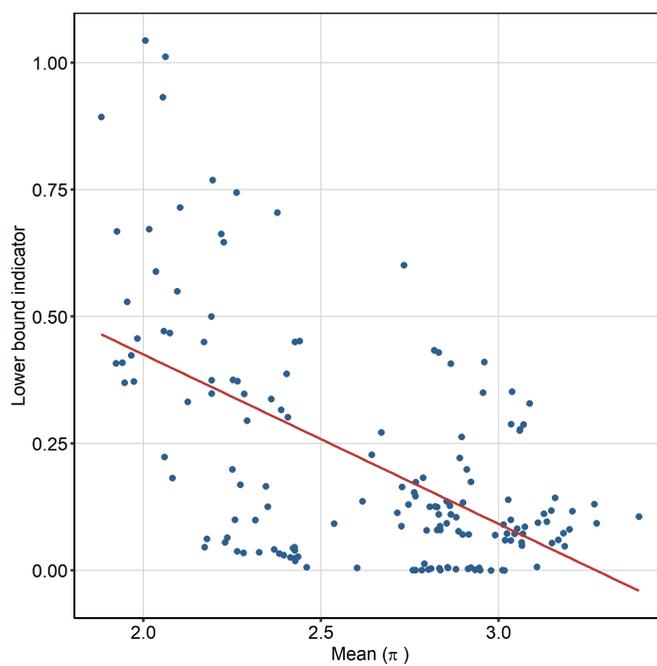


Figure 11: Relationship between lower bound indicator and the swaps-implied five-year forward rate of inflation five years ahead. The regression line has an intercept of 1.09 (0.47) and a slope of -0.33 (0.16) where standard errors are Newey-West adjusted. The intercept and slope are significant at the 5% level.

Short-term Forecast Densities for Interest Rates

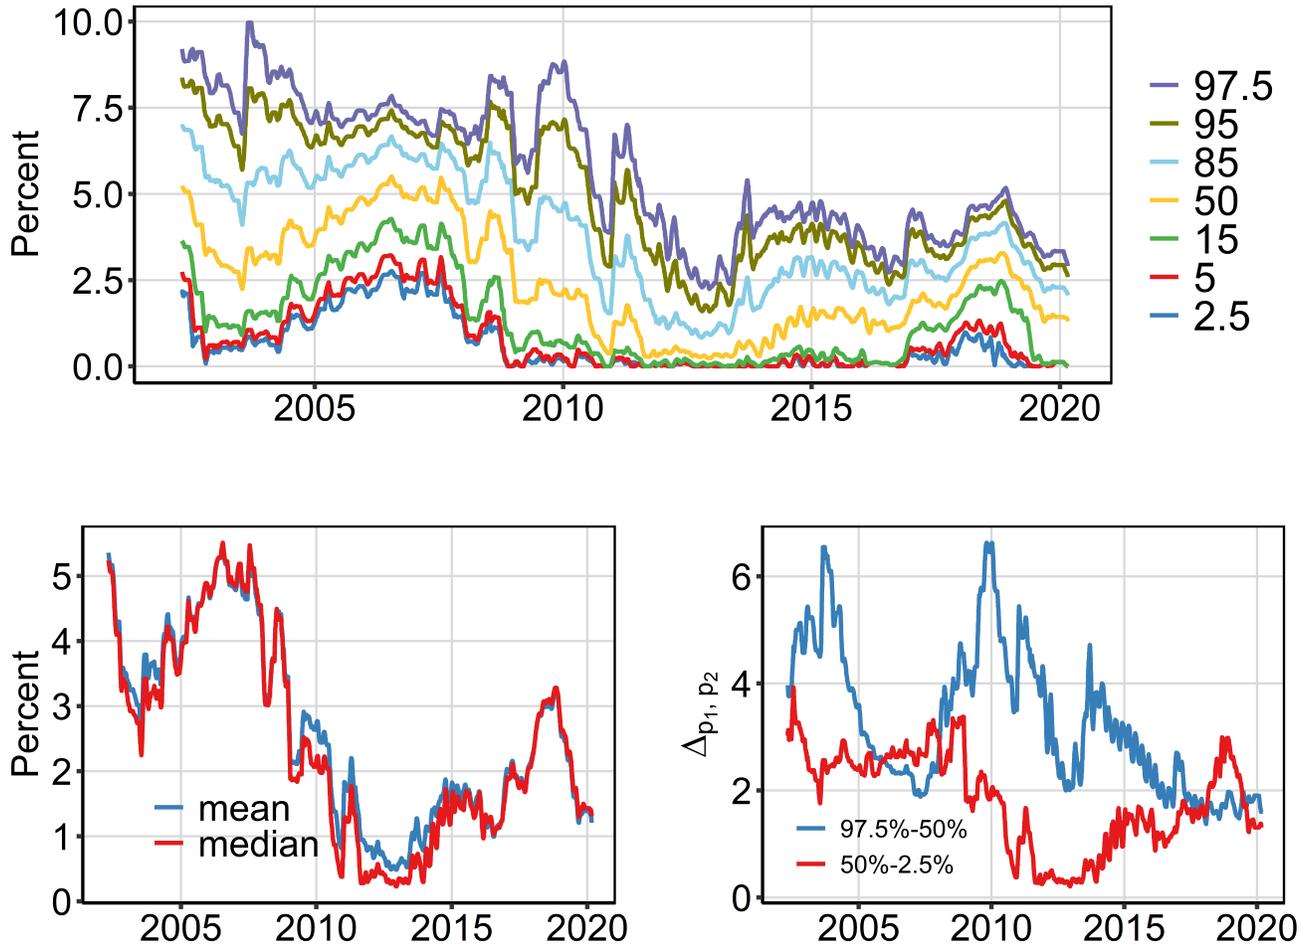


Figure 12: This figure reports the same measures of forecast densities as in Figure 6 for short-term densities. The figures represent 20-day moving averages of percentiles and moments of the distribution of the three-month forward rate 24 months out.