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

Thomas M. Mertens and John C. Williams

Online Appendix

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

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

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

$$\pi_t = (1 + \alpha\kappa)\mathbb{E}_t\pi_{t+1} + \mu_t + \kappa\epsilon_t - \alpha\kappa(i_t - r^*), \tag{18}$$

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 $\mathbb{E}\pi_{t+1} = \mathbb{E}\pi_t = \mathbb{E}\pi$ 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

$$\mathbb{E}\pi_t = \mathbb{E}i_t - r^*.$$

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

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

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

$\mathtt{i}_t = i^{LB}$, i.e. $\bar{\theta}^{LB}(r^*) = i^{LB} - r^* - \psi\mathbb{E}\pi$. 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

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

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)

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

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

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

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)

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

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

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

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

Now we use the definition of the shadow rate $\mathtt{i}$ and take unconditional expectations on both sides to get

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

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. □

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. □

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 (see **?**) and Eurodollar futures with quarterly maturities from one to eight quarters out via **?** 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

[11] Ticker symbols are EDx Comdty where x needs to be replaced by the number of quarters out. Used with the Permission of Bloomberg.

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

with T payments are equated such that

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

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 $\mathbb{E}^g[i_s] = Z_s/Z_{s+1} - 1$. 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 $\sigma_{t,\tau}^{cap}$ 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}

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

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

[13]Ticker 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.

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.

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

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

where we obtain $\mathbb{E}^g[i_\tau]$ from forward rates, $\mathbb{E}^{\hat{g}}[i_\tau] = \int i_\tau \hat{g}(i_\tau; \varpi) di_\tau$ 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 on the domain with boundaries 10^{-6} for all variables in the vector $(\omega, \mu_1, \sigma_1, \mu_2, \sigma_2)$ on the low end and $(1, 20, 10, 20, 10)$ for the respective variables on the upper end. Note that the algorithm does not require an initial guess. 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.

[14]We 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.

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 **?**) 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] 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

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

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

[15]The ticker symbols are USSWIT5 Curncy and USSWIT10 Curncy.

[16]The ticker symbols for these contracts are USIZCxm and USIZFxm where x specifies the strike and m the maturity.

written for forward rates, is defined as the number of standard deviations a percentile is distanced from the mean

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

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). Assuming a random walk for inflation, we compute the standard deviations for forward rates. Using this standard deviation together with the mean of 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. The distribution of forward rates thereby inherits asymmetries from inflation forecast densities at the five-year and 10-year horizons.

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 average inflation rate over the next h years. This relation holds since the five-year inflation rate and the forward rate average out to the 10-year inflation rate. We further compute the standard deviation via

$$\begin{aligned}\sigma_{f,t,5,10} &= \left(4\mathrm{Var}_t[\bar{\pi}_{t,10}] + \mathrm{Var}_t[\bar{\pi}_{t,5}] - 4\mathrm{Cov}_t[\bar{\pi}_{t,5}, \bar{\pi}_{t,10}]\right)^{\frac{1}{2}} \\ &= \left(4\mathrm{Var}_t[\bar{\pi}_{t,10}] - 3\mathrm{Var}_t[\bar{\pi}_{t,5}]\right)^{\frac{1}{2}},\end{aligned}$$

where the second line uses the assumption of a random walk for inflation. 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. $\hat{Z}_f^p = (Z^p_{\bar{\pi},5} + Z^p_{\bar{\pi},10})/2 \equiv \bar{Z}^p_{\bar{\pi}}$. Plugging this expression into (24), we obtain the p-th percentile

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

In the following, we show that this procedure precisely recovers the distribution of the forward rate when inflation follows a Gaussian random walk process

$$\pi_{t+1} = \pi_t + \nu_{t+1}, \tag{25}$$

where ν_{t+1} is normally distributed with mean zero and variance σ_ν^2. Inflation at time $t+h$ is thus, based on time-t information, conditionally normally distributed with mean π_t and variance $\text{Var}_t\left[\pi_{t+h}\right] = h\sigma_\nu^2$.

Lemma 3 (Z-scores of forward rates)

Let inflation follow a Gaussian random walk process. The Z-scores Z_f^p for all percentiles of the forward rate between h_1 and h_2 periods out, $\hat{f}^p_{t+h_1,t+h_2}$, coincide with the Z-scores of the average inflation rates for horizons h_1 and h_2 for all percentiles p.

Proof: Given its definition in equation (23), the forward rate $f_{t+h_1,t+h_2}$ is conditionally normally distributed provided that inflation follows a Gaussian random walk, as in (25).

The p-th percentile of the forward rate is $\hat{f}^p_{t+h_1,t+h_2} = \mathbb{E}_t[f_{t+h_1,t+h_2}] - \sqrt{2}\sigma_{f,t,h_1,h_2}\text{erfc}^{-1}[2p]$ 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

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

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

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. Using the average of the Z-scores, as described above, thus delivers exact results if inflation rates follow a Gaussian random walk and serves as an approximation otherwise.

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 (see **?**).[17] 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.[18] 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.[19]

[17]The market data is the property of CME Group Inc. All rights reserved.

[18]Ticker 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.

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

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 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

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

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

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

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 et al. (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 $1/(\sqrt{2}\mathrm{erfc}^{-1}\left(\frac{3}{10}\right))$. 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 (see **?**). 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.

D Example with uniform distribution

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

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

$$\text{Prob}(\mu_t < \frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1})) = \begin{cases} 1 & \text{if } \; -\frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1}) \leq -\hat{\mu} \\ \frac{1}{2\hat{\mu}}(\hat{\mu} + \frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1})) & \text{if } \; -\hat{\mu} < -\frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1}) < \hat{\mu} \\ 0 & \text{if } \; -\frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1}) \geq \hat{\mu}. \end{cases} \tag{27}$$

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

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

$$\mathbb{E}\pi_t = \begin{cases} -\alpha\kappa(i^{LB} - r^*) + (1 + \alpha\kappa)\mathbb{E}\pi_{t+1} & \text{if } \; -\frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1}) \leq -\hat{\mu} \\ -\frac{1}{4\hat{\mu}}\left(\frac{\kappa^2}{\kappa^2+\lambda}\left[\frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1}) + \hat{\mu}\right]^2\right) + \frac{\lambda\beta}{\kappa^2+\lambda}\mathbb{E}\pi_{t+1} & \text{if } \; -\hat{\mu} < -\frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1}) < \hat{\mu} \\ \frac{\lambda\beta}{\kappa^2+\lambda}\mathbb{E}\pi_{t+1} & \text{if } \; -\frac{1}{\gamma}(i^{LB} - r^* - \psi \mathbb{E}\pi_{t+1}) \geq \hat{\mu}. \end{cases} \tag{28}$$

Note that in the special case of $\lambda = 0$, the optimal policy achieves zero inflation each period (see the last line in equation (28)) except when policy is constrained by the lower bound. In that case, for the intermediate range when policy is constrained some but not all the time, the equation for the expected value of inflation simplifies to: $\mathbb{E}\pi_t = -\frac{1}{4\hat{\mu}}\left(-\hat{\mu} - \alpha\kappa(i^{LB} - r^*) + (1 + \alpha\kappa)\mathbb{E}\pi_{t+1}\right)^2$.

In looking for an unconditional mean, we equate expected values $\mathbb{E}\pi_t$ and $\mathbb{E}\pi_{t+1}$. The set of equations (28) describes the functional relationship between these expected values of π. For a range of intermediate values of $\mathbb{E}\pi$, the relationship is quadratic. Outside of this range, the relationship is piecewise linear, with a slope greater than one for low values of $\mathbb{E}\pi_{t+1}$ and less than one for high values of $\mathbb{E}\pi_{t+1}$.

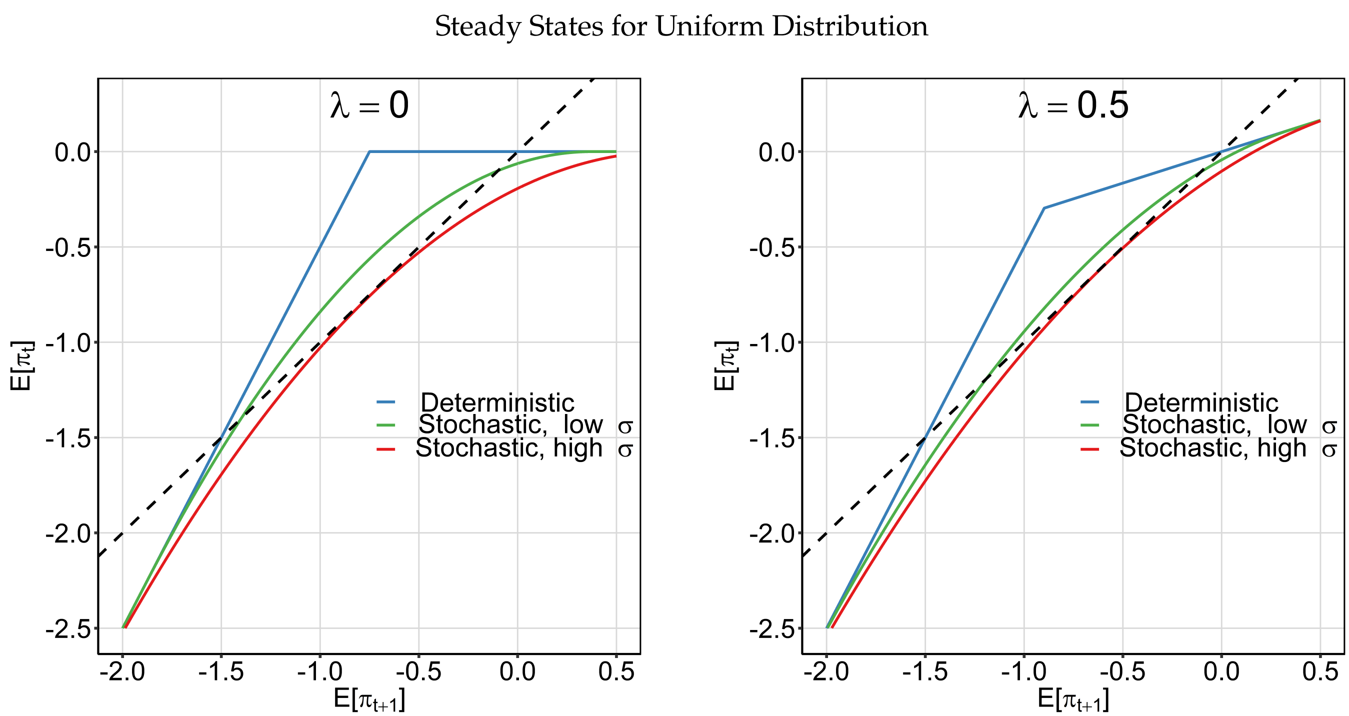

Online Appendix Figure 1: Expected inflation in the current period as a function of expected inflation in the following period assuming a uniform distribution. Parameter values are set to $\alpha = 1$, $\kappa = 1$, $\beta = 0.99$, $r^* = 1\%$, $i^{LB} = -0.5\%$. There are no ϵ-shocks, i.e., $\sigma_\epsilon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

There are either zero, one, or two values that satisfy the conditions in (28). For small values of $\hat{\mu}$, there are two steady states, as in the deterministic model. Once the degree of uncertainty increases beyond a certain point, the corner solution where the lower bound is either always or never binding no longer applies. If this occurs in the target equilibrium, the lower bound constrains policy for low realizations of the shock and, as a result, the unconditional mean of π decreases. If this occurs in the liquidity trap equilibrium, the lower bound does *not* constrain policy for high shock realizations, and the unconditional mean of π *increases.*

Figure 1 illustrates the equilibria in the model for different degrees of aggregate uncertainty parameterized by $\hat{\mu}$. To produce Figure 1, we use the following parameter combination: $\alpha = \kappa = 1$, $\beta = 0.99$, $i^{LB} = -0.5\%$, and $r^* = 1\%$. The left panel shows the case of $\lambda = 0$, where the central bank seeks only to stabilize inflation; the right panel shows the case of $\lambda = 0.5$, where the central bank also seeks to stabilize the output gap. In the left panel, the blue line shows the values of $\mathbb{E}\pi_t$ for given values of $\mathbb{E}\pi_{t+1}$ for the deterministic case ($\hat{\mu} = 0$). Note that the relationship is piecewise linear. Two steady states emerge where the function crosses the 45-degree line indicated by the dashed black line. The green line shows the corresponding functional relationship for the case of $\hat{\mu} = 2.25\%$. In this case, the function is quadratic over the relevant range and the quadratic relationship between $\mathbb{E}\pi_t$ and $\mathbb{E}\pi_{t+1}$ crosses the 45-degree line in two places.

E Robustness to alternative underlying distributions

This section shows an alternative construction of long-term and short-term forecast densities for interest rates. The methodology in the main body of the paper relies on a mixture of lognormal distributions. Here, instead, we rely on a more parsimonious distribution, the rectified Gaussian distribution where we allow for a mass point at the zero lower bound.

To construct the forecast densities under the rectified Gaussian distribution, we replace the mixture of lognormals assumption in the estimation by an underlying normal distribution that gets truncated at the zero lower bound. We find the optimal parameters for the rectified distribution by minimizing the squared distance between the implied prices and observed prices. We thus estimate the parameters directly from the data. This step is analogous for caps and Eurodollar futures options and is explained in the corresponding sections in Appendix B.

The resulting time series of forecast densities for long-term interest rates are plotted in Figure 2. Similarly to the ones for the mixture of lognormal distributions, average forecast densities have been trending down during the sample period. In comparison, however, the mass at the lower bound is higher under the rectified Gaussian distribution as can be seen from the time series of the 15th-percentile.

The lower left-hand panel shows both mean and median fell during the time period when the natural rate, r^* declined. Therefore, the evidence is again consistent with the target equilibrium of our theoretical model and at odds with the liquidity trap equilibrium.

The path for short-term forecast densities for interest rates is also consistent with the estimates from the

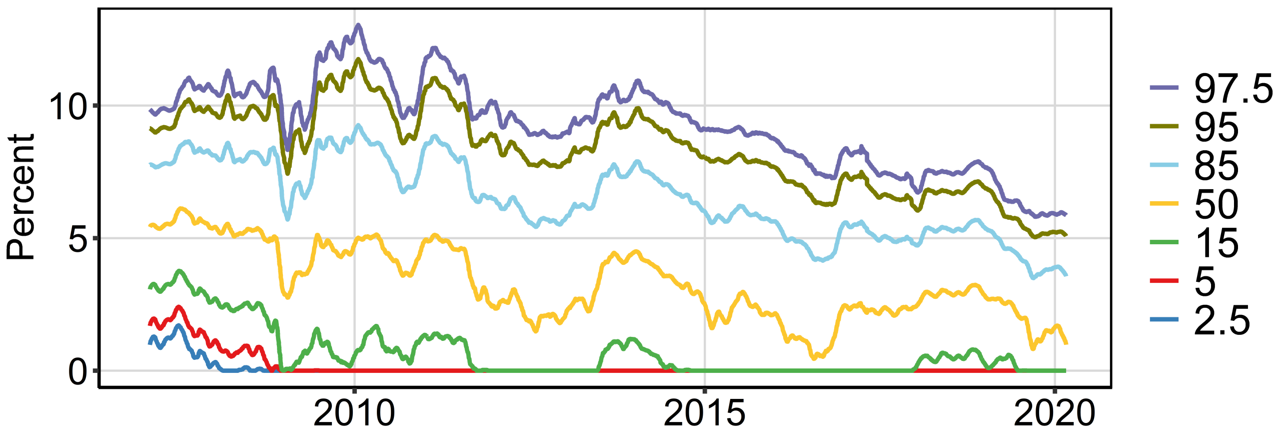

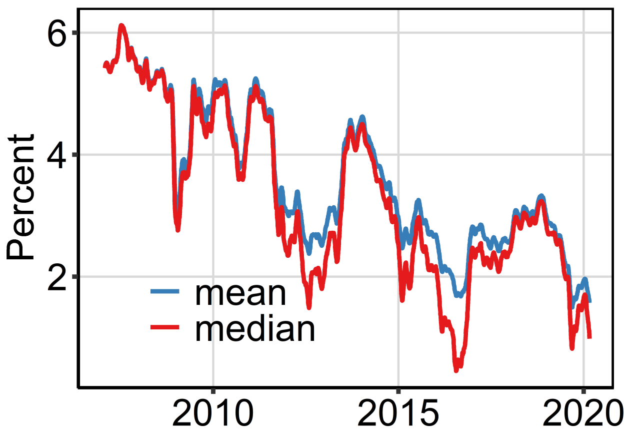

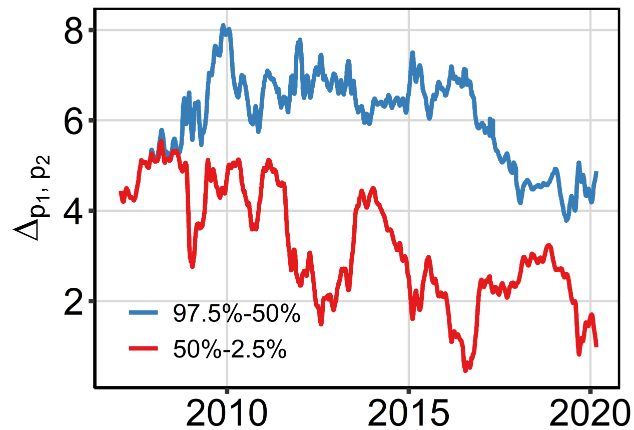

Online Appendix Figure 2: Long-term forecast densities for interest rates measured by the three-month forward rate seven years out assuming an underlying rectified Gaussian distribution. The upper panel shows a 20-day moving average of the daily time series for various percentiles. The lower panels plot 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).

mixture of lognormal distributions. Figure 3 plots our estimates. The general patterns are the same as under the estimates in the main body of the paper.

While the switch to the truncated distribution might alter estimates of the level of forecasts and uncertainty, the general pattern of changes of these quantities remained the same over the time period when the natural rate of interest fell.

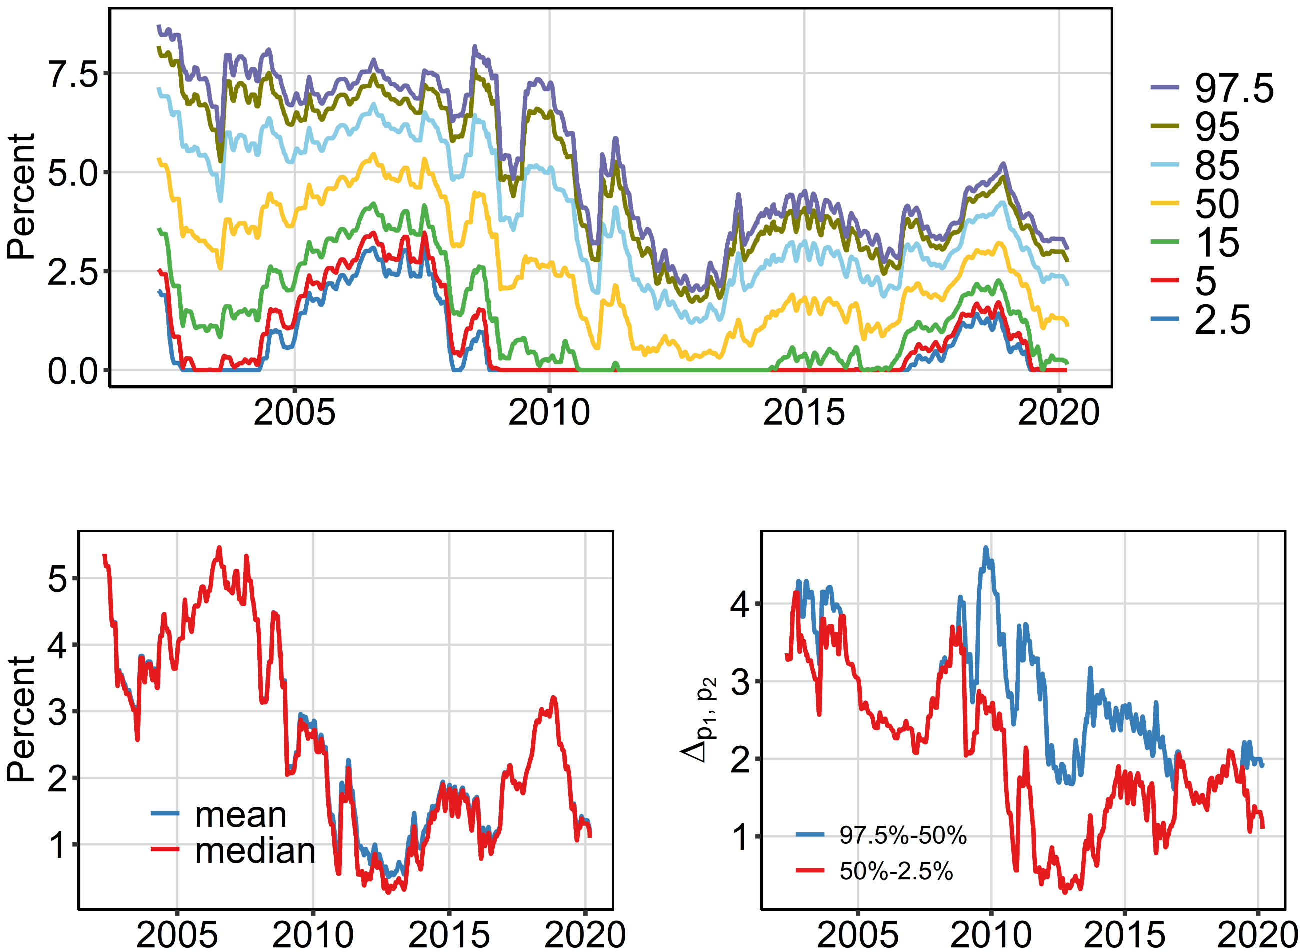

Online Appendix Figure 3: Short-term forecast densities for interest rates measured by the three-month forward rate two years out assuming an underlying rectified Gaussian distribution. The upper panel shows a 20-day moving average of the daily time series for various percentiles. The lower panels plot 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).

F Robustness to alternative subsamples

This section investigates the robustness of the main results to a different cutoff. We therefore choose the end of 2013 as an alternative date to break the sample into two subsamples. The mean estimate for r^* falls from 1.03 percent over 2007-2013 to 0.54 percent over 2014-2020. 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 -3.44. The less negative t-statistic, compared to the number in Section III.B for a cutoff in 2011, indicates that the break in the time series of r^* more likely occurred in 2011.

The results we get from altering the cutoff are qualitatively consistent. For the interest rate distribution, there is a marked decline in the mean and the median. The skewness is positive and rising over the second subsample but is not as highly statistically significant as in the results reported in Table 1.

Online Appendix Table 1: Summary of interest rates moments, end-of-2013 cutoff

	2007-2013	2014-2020
Mean	4.37%	2.74%***
Median	4.01%	2.26%***
Std. deviation	5.29%	3.74%***
Skewness	0.25%	0.44%*

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

The picture for inflation is similar. Comparing Online Appendix Table 2 to Table 2 in the main text, the results are qualitatively and quantitatively similar. The regression coefficients tend to be statistically more significant in this alternative specification.

Online Appendix Table 2: Summary of inflation moments, end-of-2013 cutoff

	Forecast densities		Inflation swaps	
	2011-2013	2014-2016	2007-2013	2014-2020
Mean	2.67%	2.23%***	2.94%	2.30%***
Median	2.67%	2.25%***		
Std. deviation	1.87%	1.17%***		
Skewness	0.00%	-0.10%***		

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

G Robustness: Extended Time Series for Inflation Forecast Densities

In this subsection, we use the quotes for inflation caps and floors provided by Bloomberg beyond the end of the sample period in the main body of the paper. Here, we construct forecast densities for inflation for the time period from January 1, 2007, up to February 28, 2020.[20]

The upper panel in Online Appendix Figure 4 shows the time series of a 20-day moving average of long-term inflation forecast densities, where the measure of inflation is a five-year forward rate five years out, as

[20]There are 12 days in the sample for which the implied variance of forward rates is negative. We drop those data points from the time series.

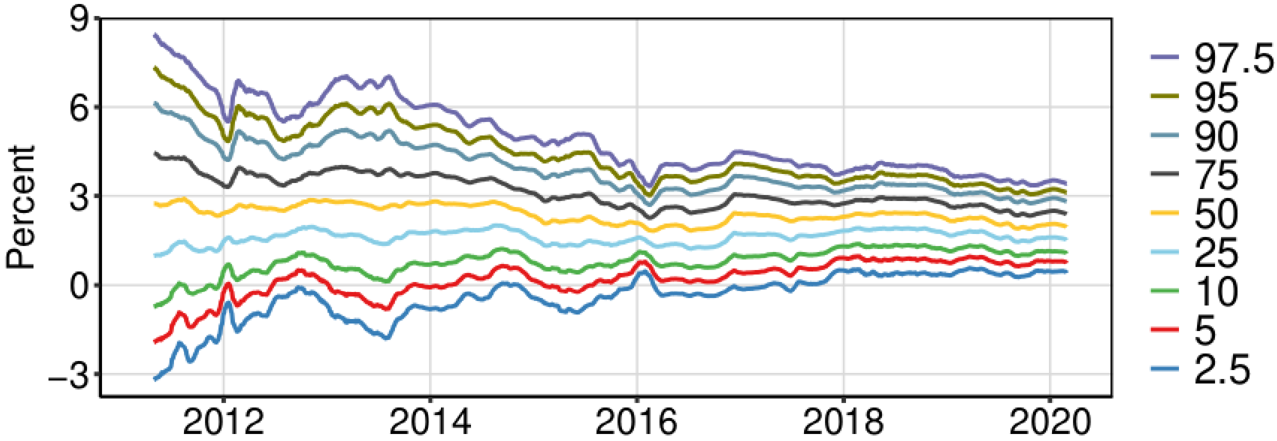

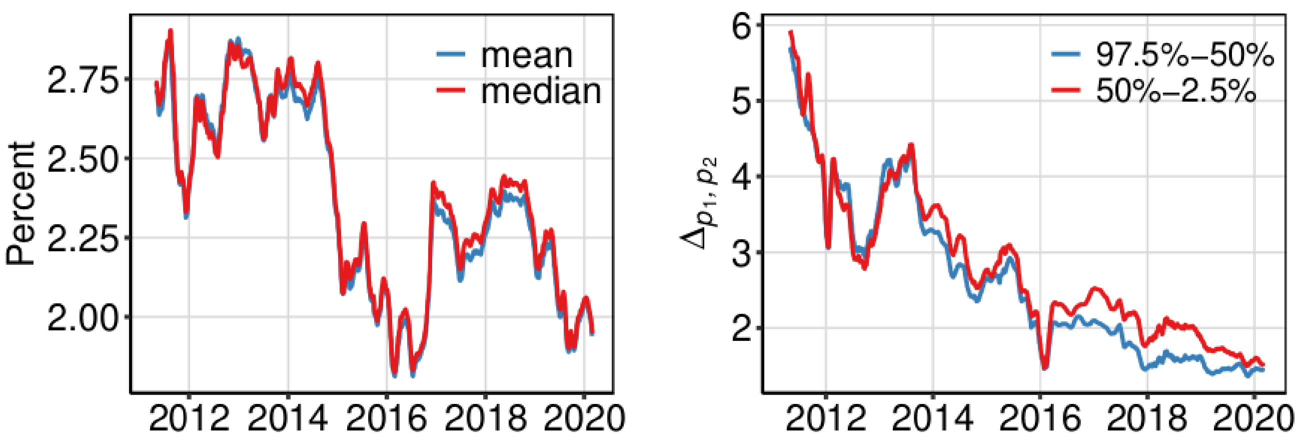

Online Appendix Figure 4: Long-term forecast densities for inflation measured by the five-year forward rate five years out. This figure reports measures of forecast densities analogous to Figure 6 for inflation rates measured by the five-year forward rate five years out. All pictures show 20-day moving averages of the underlying daily time series.

in the main body of the paper.

Online Appendix Table 3 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 2.33 percent for 2012-2020. The density of inflation is slightly skewed to the left in both subsamples with the skewness increasing in the later subsample.

Online Appendix Figure 5 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 right panel shows again

Online Appendix Table 3: Summary of inflation moments, extended time series

	2011	2012-2020
Mean	2.60%	2.33%**
Median	2.62%	2.35%**
Std. deviation	2.26%	1.18%***
Skewness	-0.04%	-0.11%

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

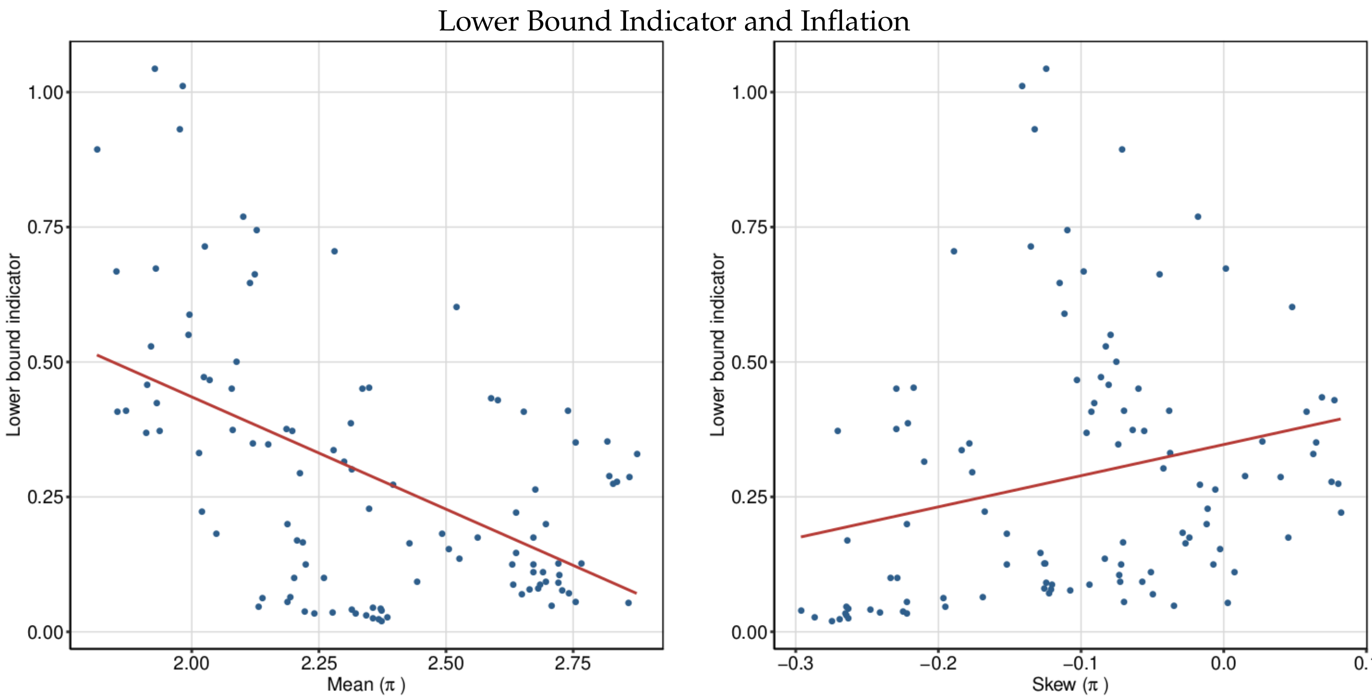

Online Appendix Figure 5: Relationship between lower bound indicator and inflation moments at a monthly frequency. The scatter plots show monthly data of mean inflation versus the lower bound indicator (left panel) and the skewness in inflation forecast densities (right panel). The regression line in the left panel has an intercept of 1.27 (0.56) and a slope of −0.42 (0.23) where standard errors are Newey-West adjusted. The intercept is significant at the 5% level and the slope at the 10% level. In the right panel, the intercept is 0.35 (0.07) and the slope 0.58 (0.39). The intercept is statistically significant at the 1% level and the slope statistically insignificant.

that the changes in the inflation distribution are less pronounced than for the interest rate distribution. The skewness of inflation is negative and increases with higher levels of the lower bound indicator. The upward sloping nature of the link between the asymmetry of the inflation distribution and the lower bound indicator at the monthly level is statistically insignificant, yet still interesting. While on average, a lower r^* has led to a more negative skewness in the later subsample in Online Appendix Table 3, higher values for the lower bound indicators are associated with less negative skewness on a monthly frequency. 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.

H Robustness to specification of forward rates for inflation

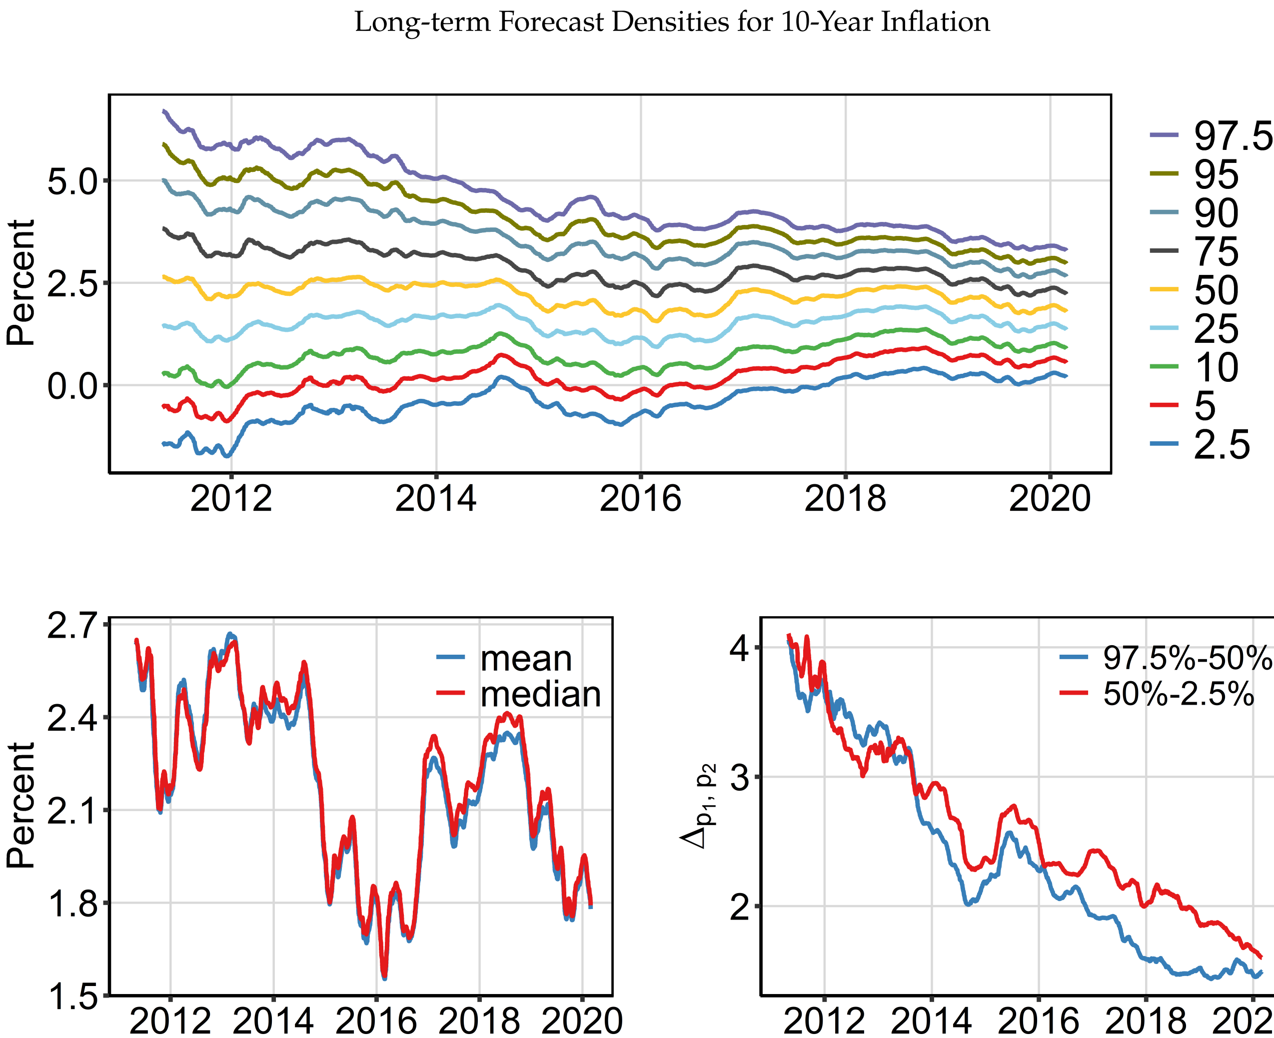

Online Appendix Figure 6: Long-term forecast densities for inflation over the following 10 years. All lines show 20-day moving averages of the underlying daily time series.

There are many different ways to convert forecast densities of average inflation over some time period into forward rates. In the main body of the paper, we start with an AR(1) model for inflation rates where the i.i.d. model and the random walk model are nested as special cases. As a robustness check, this section shows that forecast densities of the average inflation rate over the next 10 years display the same patterns as the five-year forward rate five years out. We thus conclude that the results in the main body of the paper do not rely on the specific method for constructing forward rates.

Online Appendix Table 4: Summary of inflation moments, inflation rates over next 10 years

	2011	2012-2016
Mean	2.36%	2.18%
Median	2.39%	2.20%
Std. deviation	1.78%	1.26%***
Skewness	-0.05%	-0.06%

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

Online Appendix Figure 6 shows the time series for forecast densities of 10-year average inflation rates analogous to the one for forward rates in Figure 9 in the main body of the paper. The mean and median of the distribution drift downwards during the sample period with the distribution becoming skewed towards the downside.

Online Appendix Table 4 shows these trends for the two subsamples used in the main body of the paper. Consistent with the evidence for forward rates, the skewness is negative and falls in the later part of the sample period. The declines in mean, median, and skewness are thereby of very similar magnitude.

I Model fit under parameterized distributions for interest rates

In this section, we show the fit of model-implied option prices to the data under the functional form we chose. The mixture of lognormal distributions for short-term interest rates gives us five parameters to match the out-of-the-money calls and puts on Eurodollar futures.

Therefore, we calculate option prices by computing the discounted value of the expected payoff under the mixture of lognormal distributions, as in equation (13) for calls and analogously for puts by discounting the payoff $\max\{K_n - i_t, 0\}$.

Figure 7 shows the fit of the model to the data (the dots in the chart). The graph is produced with data from the last day in our sample period. The fit on other days looks similar. The orange line shows the fit through put prices while the green line shows the fit through calls.

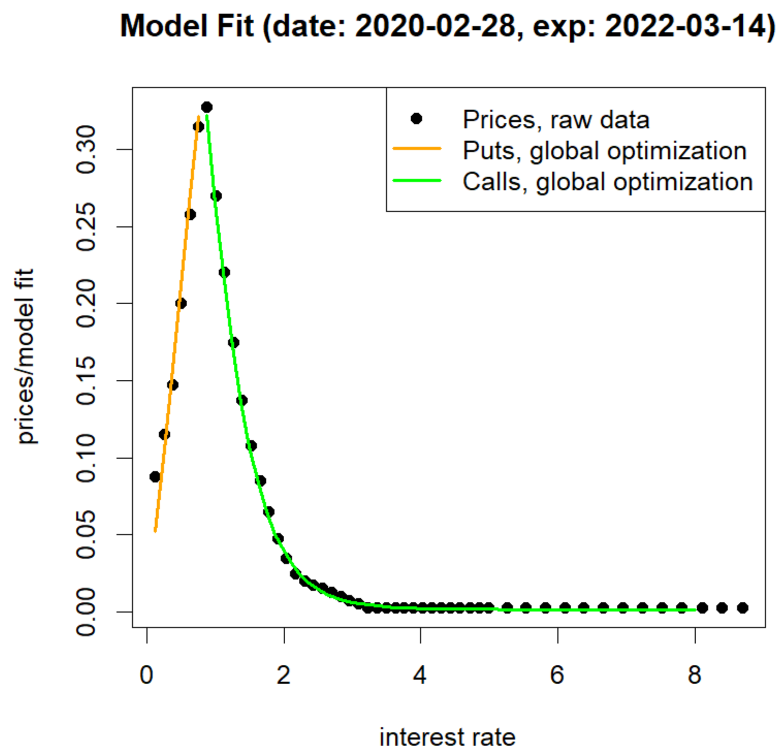

Online Appendix Figure 7: Fit of prices across strikes. The orange line fits through put prices while the green line fits through call prices. Both colored lines are derived from the same underlying mixture of lognormals distribution.