Asymmetric Monetary Policy Expectations*

Anthony M. Diercks[†]
Federal Reserve Board

Hiroatsu Tanaka[‡]
Federal Reserve Board

Paul Cordova[§]
Columbia

August 2025

Abstract

We document some novel empirical evidence of significant time-varying skewness in the aggregate forecast distribution of the monetary policy rate, using unique distributional information of the federal funds rate (FFR) forecasts in the Survey of Primary Dealers (SPD). We further show that a nonlinear New-Keynesian model with an effective lower bound (ELB) constraint and (symmetric) shocks to aggregate uncertainty can endogenously generate both positive and negative skewness patterns in the data, which cover periods at and away from the ELB. Our skewness measure is constructed from "physical" forecast distributions of the FFR, free of risk premia, in contrast to measures extracted from asset prices, and the time-variation of asymmetry in the aggregate distribution has implications for the measurement of monetary policy expectations more broadly. In particular, we show that mean and modal expectations can differ significantly, and the FFR forecasts from the Blue Chip Survey, a popular survey measure, track the mode more closely. In addition, the mean extracted from the SPD has distinct implications for forecast performance and the rationality of monetary policy expectations.

JEL: E43, E44, E52, G17

Keywords: Skewness, Federal Funds Rate Expectations, Survey of Primary Dealers, Blue Chip Survey, Monetary Policy, DSGE.

*We thank Michael Bauer, John Cochrane, Michiel De Pooter, Greg Duffee, Eric Engstrom, Andrey Ermolov (Discussant), Stefano Giglio, Haitham Jendoubi, Don Kim, Andrew Meldrum, Stefan Nagel, Andrew Patton, Zhaogang Song, Annette Vissing-Jorgensen, Benedict Wensley, Jonathan Wright as well as seminar participants at the Federal Reserve Bank of New York and the WFA Conference for helpful comments. All errors remain our sole responsibility. The views expressed herein are those of the authors and not necessarily those of the Board of Governors of the Federal Reserve System.

†Division of Monetary Affairs, Federal Reserve Board, 20th Street and Constitution Avenue N.W. Washington, D.C. 20551; http://www.anthonydiercks.com; Email: anthony.m.diercks@frb.gov

‡Division of Monetary Affairs, Federal Reserve Board, 20th Street and Constitution Avenue N.W. Washington, D.C. 20551; Email: hiroatsu.tanaka@frb.gov.

§Email: pac3nc@virginia.edu

1 Introduction

Gauging financial market expectations about the future course of monetary policy is vital for both market participants and policy makers. The expectations are typically measured by some (point) forecast of the federal funds rate, and various surveys of professional forecasters provide important information in this regard. However, for some surveys, it is not always clear what the responses truly indicate. In particular, it is unclear whether the responses indicate a statistical average (mean) outcome or some other statistic of central tendency, such as the most likely (modal) outcome of policy rates in the future. It is important to note that the mode is unlikely to capture increased likelihood of upside or downside outcomes that are reflected in the mean because the mode excludes information on the shape of the rest of the distribution. If the distribution of future policy rates were symmetric and unimodal, the distinction between means and modes would be irrelevant, but if the distribution was skewed, the mean-mode difference could be significant.

In this paper, we revisit the measurement of central tendencies by computing measures of mean and modal expectations from the aggregate forecast distribution of the fed funds rate. To this end, we rely on the Survey of Primary Dealers (SPD), which is compiled by the Federal Reserve Bank of New York (FRBNY) before each FOMC meeting, and provides important information on the primary dealers' outlook on monetary policy and the economy.[1] The survey results include the aggregate distribution of the fed funds rate forecast distribution, which is the average across each respondent's forecast distribution of the fed funds rate.[2] The SPD is publicly available since 2011 and covers forecast distributions for horizons up to about three years, but we primarily focus on the one-year ahead measure because it is among the most relevant horizons for policy makers.

The construction of the distributions is non-trivial since the results of each survey are summarized in individual pdf files and the format of the survey questions and responses is not consistent over time. Therefore, we hand-collect survey results necessary to compute measures of skewness, specifically the forecasts of the federal funds rate. These survey-based measures are unique in that, by construction, they do not leave ambiguity as to whether they are the mean, mode or any other measure of central tendency. Furthermore, the expectations derived from the survey distribution are distinct from measures extracted from interest rate options, which need additional assumptions to disentangle market expectations from risk premia. The survey-based measures directly provide central tendencies of the "physical distribution" of the fed funds rate.

We document some novel empirical evidence of significant time-varying skewness in the aggregate forecast distribution of the federal funds rate (based on the difference between the mean and mode), i.e. *asymmetric monetary policy expectations*. While the mode was generally lower than the mean when the policy rate was at the effective lower bound (ELB) from 2011 to 2015, the mode increased faster than the mean as the exit from the ELB drew closer. The mode generally stayed

[1] See Correia-Golay, Friedman, and McMorrow (2013) for an introduction to the SPD provided by the FRBNY staff.

[2] The Survey of Professional Forecasters provides forecast distributions for inflation and output growth, but do not provide forecast distributions for the fed funds rate.

significantly higher than the mean after liftoff. However, the skewness increased sharply once again after the pandemic hit, and the policy rate returned to the ELB. More recently, as expectations for liftoff has increased, the skewness of the distribution has dropped significantly. We show these dynamics are consistent across a broad range of skewness measures including those estimated from parametric and nonparametric continuous distributions that are fitted to the survey data.

We then provide a simple macroeconomic framework that can help explain the pattern of time-varying skewness of the policy rate distribution. We show that a stylized New-Keynesian model which takes into account an occasionally binding ELB constraint can endogenously generate not only the observed positive skew when the economy is at the ELB, but also the negative skew when the economy is above but still near the ELB, even under the assumption of *symmetric* shocks to the economy. The positive skew is a result of the policy rate distribution being truncated by the ELB constraint, as was the case from 2011 through 2014, as well as from 2020 through 2021, while the negative skew is due to the increased downside risks related to monetary policy being less able to effectively counter a recession at the ELB, consistent with the negative skew observed from 2015 through 2019. While we explain the key mechanism using a model that features only demand shocks, we also show that the result extends to a model with both demand and supply shocks.

We further extend the model to include *symmetric* second moment shocks ("uncertainty" shocks), as a novel mechanism for endogenously generating negative skewness in the forecast distribution of the policy rate even when the rate is sufficiently away from the ELB, which is consistent with our empirical evidence after the pandemic crisis. Intuitively, the consumer in the model has a greater precautionary savings motive when confronted with additional uncertainty at times when uncertainty is already high. The desire for even greater precautionary savings at times of high uncertainty leads to endogenously asymmetric responses of consumption, inflation and interest rates to symmetric shocks. Importantly, this mechanism relies solely on accurately capturing higher-order derivatives of a standard utility function and can be obtained as long as the model is solved globally or with an order of approximation that is above what is typically considered in the literature. This is in sharp contrast to other mechanisms in the literature that generate skew exogenously, such as those incorporating shocks to the third moment (Ferreira, 2024; De Polis, Melosi, and Petrella, 2024).

We further discuss some important implications of the time-varying skewness in the forecast distribution.

First, the fact that we have explicit measures of mean and modal policy expectations allows us to evaluate whether other measures of expectations are more likely to capture the mean or the mode of the average forecaster's distribution. We apply this idea to a popular measure of policy expectations extracted from the Blue Chip Survey (BCS).[3] Using some unpublished results from the SPD, we find that the BCS expectations tracked the average of the modal expectations in the SPD very closely over the past decade, based on a matched sample of respondents. In

[3] As discussed in the following sections, this measure specifically refers to the average of the respondents' forecasts from the Blue Chip Financial Forecasts Survey, otherwise known as the BCS "consensus" forecast.

contrast, we observe significant gaps between the BCS expectations and the mean of the aggregate distribution. This finding suggests some caution in treating the BCS forecasts as representative of mean expectations, which is not uncommon in the literature.

The discrepancy between monetary policy expectations measured by the mean or the mode also leads us to investigate the relative usefulness of these measures. In this regard, we argue that the measure of mean policy expectations from the aggregate distribution is a preferred measure of near-term monetary policy expectations for at least a few reasons. First, the relevance of the mean is deeply rooted in expected utility theory and is arguably the most widespread notion of expectations. For instance, most models of the yield curve decompose yields into mean expected rates and term premiums. It is hence useful to have a direct measure of mean expectations, which could also serve as a benchmark to various model-based measures of policy expectations. On that note, the mean expectations extracted from the SPD imply significantly less negative term premia on average, compared to term premia implied by the BCS and SPD modal forecasts since 2011.

We show the term premium computed from modal expectations includes a fraction of skewness in addition to the standard measure of term premium based on mean expectations. This is demonstrated by the following decomposition of the forward rate f_t^τ into modal expectation $\hat{\mathbb{E}}_t[r_{t+\tau}]$, skewness—as defined by mean minus mode normalized by standard deviation ("Pearson Skewness")—$Sk_t\left[r_{t+\tau}\right]$, and the term premium tp_t:

$$\begin{aligned}
f_t^\tau &= \mathbb{E}_t[r_{t+\tau}] + tp_t^{mean} \\
&= \rlap{\overbrace{\phantom{\hat{\mathbb{E}}_t[r_{t+\tau}] + \sigma_t[r_{t+\tau}]Sk_t[r_{t+\tau}]}}}\hat{\mathbb{E}}_t[r_{t+\tau}] + \underbrace{\sigma_t[r_{t+\tau}]Sk_t[r_{t+\tau}] + tp_t^{mean}} \\
&= \hat{\mathbb{E}}_t[r_{t+\tau}] + tp_t^{mode}
\end{aligned}$$

The second reason why a more precise measure of mean expectations could be useful is that it may lead to superior forecasts. Compared to the BCS forecasts which appear to be representative of the mode, we find that the mean expectations more accurately forecasted the one-year ahead policy rate based on the mean squared error (MSE) criterion. This result is consistent with the theory that a mean expectation is the optimal forecast based on the MSE loss function.

There is a large literature on federal funds rate forecasts, and the importance of considering the asymmetric nature of the forecast distribution is fairly recognized by staff in the Federal Reserve System and policy makers.[4] Yet, there are relatively few studies that explicitly analyze the time-series properties of the forecast distribution.[5] To our knowledge, Potter, Del Negro, Topa, and

[4]In terms of official policy discussions that are publicly available, see for example transcripts of the December 2014 (page 5) and the October 2015 FOMC (page 165) meetings, among others. The topic is also taken up by numerous pieces by Federal Reserve staff that discuss current policy issues, see e.g. Crump, Moench, O'Boyle, Raskin, Rosa, and Stowe (2014), Brodsky, Del Negro, Fiorica, LeSueur, Morse, Rodrigues, et al. (2016), Joergensen and Meldrum (2019). Wright (2017) offers a survey on density forecasts for interest rates and provides some discussion on the use of Desk surveys to extract information about the physical distribution of the fed funds rate.

[5]The distributional properties of GDP growth and inflation forecasts have been more studied, partly due to the availability of the distributions through the Survey of Professional Forecasters, see e.g. Giordani and Söderlind (2003), Engelberg, Manski, and Williams (2009), D'Amico and Orphanides (2014), Huang, Pilbeam, and Pouliot (2022).

Van der Klaauw (2017) were among the first detailed studies that show probability-weighted mean expectations of the fed funds rate do not necessarily equal modal forecasts, also using results from the Desk surveys. However, their discussion is limited to a few Desk surveys conducted around the beginning of 2016. Kwong, Leung, Wong, and Zhang (2020) similarly documents the difference in the mean and modal forecasts, but their sample period is shorter than ours, and do not discuss the relation with other measures of policy expectations, as we do. Importantly, the aforementioned studies do not provide an economic framework to explain the asymmetry in policy expectations. In this sense, our work is related to the literature analyzing business-cycle asymmetry with a macroeconomic model such as Berger, Dew-Becker, and Giglio (2020) and Jensen, Petrella, Ravn, and Santoro (2020). Mertens and Williams (2021) discuss interest rate distributions during a similar period and offer theoretical insights into how the probability of hitting the ELB interacts with the natural rate of interest and inflation expectations. Their focus is on the positive skewness of the long-term distribution when the economy is near or at the ELB. Moreover, the distributions they analyze are implied by option prices, which include risk premia. Contemporaneous work by Bauer and Chernov (2021) also analyzes interest rate skewness and its relation with excess returns and forecast errors. However, in contrast to our paper, their focus is on the forecast distribution of longer-term Treasury yields derived from derivative prices, similar to Mertens and Williams (2021). The economic mechanism they highlight is also different from ours. Bauer and Chernov (2021) argue that the effect of the ELB on skew at longer horizons may not be important, noting that if it were, skew and the level of yields should be negatively related. In contrast, with our focus on shorter horizons, we offer a general equilibrium mechanism where the ELB constraint can generate both a positive and negative relationship between skew and the level of yields, consistent with our new empirical evidence.

Since the asymmetry in policy rate expectations can have important implications on the measurement of term premia, our work builds on the vast literature on the decomposition of (default-free) yields into expected rates and term premia. While the number of papers in the field is too large to cite in full, our work is particularly related to recent studies that measure term premia using surveys, such as Piazzesi, Salomao, and Schneider (2015), Cieslak (2018), Schmeling, Schrimpf, and Steffensen (2020) and Nagel and Xu (2022). By offering a measure of expectations that directly represents the mean, we attempt to refine some of the common measures adopted in previous work. Our work also complements studies such as Bekaert, Engstrom, and Ermolov (2021a), which discusses the link between skew and term premia.

The structure of the paper is as follows. Section 2 discusses the survey data used in our analysis. Section 3 documents the time series of skewness in the aggregate forecast distribution of the federal funds rate and offers an economic interpretation using a stylized New-Keynesian model subject to a ELB constraint. Section 4 discusses some implications of the time-varying skew in the aggregate fed funds rate distribution. Section 5 concludes.

2 Construction of the Survey-Based Measures of Policy Rate Expectations

To construct measures of expectations from the aggregate forecast distribution of the fed funds rate, we rely on the Survey of Primary Dealers (SPD), which is compiled by the Federal Reserve Bank of New York (FRBNY) before each FOMC meeting, and provides important information on the Primary Dealers' outlook on monetary policy and the economy (see Correia-Golay, Friedman, and McMorrow (2013) for an introduction to the SPD provided by the FRBNY staff).[6] The survey results from January 2011 are publicly available on the FRBNY website.[7] Note that in contrast to other studies which focus on risk-neutral distributions, the expectations here are derived from "physical" distributions, arguably free of risk premia.[8]

However, the construction of the distributions is non-trivial since the results of each survey are summarized in individual pdf files and the format of the survey questions and responses is not consistent over time. Therefore, we go through the entire set of available pdf files and hand-collect the relevant survey results, specifically the forecasts of the federal funds rate. This is our starting point for the computation that follows.

The SPD typically only asks for distributions over the next few years, and in this study, we construct the one-year ahead and two-year ahead measures. Nevertheless, given that the two measures are fairly correlated, our analysis will mostly focus on the one-year ahead measure, which is often relevant to policy makers.[9] Each respondent's distribution is a set of probabilities assigned to specific bins, each of which represents a range for the federal funds rate at a certain point in the future. In line with the literature, we assume that the probabilities are concentrated on the mid-point of each bin when we calculate any statistic from the distribution.[10] Whenever a one-year ahead aggregate distribution (i.e., the average of each respondent's distribution) is available, we can simply compute its mean and mode (which are not explicitly provided). When only distributions for year-end dates are available, we first compute aggregate means/modes for the two consecutive year-end dates which bracket the date one year from the survey date, and then linearly interpolate the two aggregate means/modes for each year-end date to estimate the aggregate mean/modal expectations exactly one year from the survey date. This gives us a time series of the constant

[6]A similar survey is conducted for active investment managers as the Survey of Market Participants (SMP). We do not use the results from the SMP since it is available only for a shorter period since January 2014. In Appendix B, we show that at least in terms of the one-year ahead median of the modal fed funds rate forecasts, the SPD and SMP forecasts track each other very closely.

[7]`https://www.newyorkfed.org/markets/primarydealer_survey_questions`.

[8]See, for example, Mertens and Williams (2021) and Bauer and Chernov (2021), for studies that focus on risk-neutral distributions. The notion that survey expectations are more representative of physical distributions is standard in the literature. For example, Adam, Matveev, and Nagel (2021) reject the hypothesis that survey expectations of stock returns are risk-neutral.

[9]Another reason to focus on the one-year ahead measure is to allow for a more rigorous comparison with BCS forecasts, which are only available up to 6 quarters ahead on a monthly basis.

[10]See Appendix A for more details. Some papers make adjustments to the measures of standard deviation since the mid-point assumption may create bias (e.g. the "Sheppard correction" in D'Amico and Orphanides (2014)). We do not make such adjustments since they rely on additional assumptions about the true distribution while we are dealing with distributions that exhibit significant time-varying skew.

maturity one-year ahead aggregate mean/modal expectations of the fed funds rate.

Since the format of the survey responses is inconsistent over time, the construction of the distributions is non-trivial. Importantly, from the January 2016 to the October 2019 surveys, respondents were asked to report multiple distributions of the future fed funds rate based on certain scenarios about the future economy or monetary policy and the probabilities attached to each scenario, instead of one distribution for each forecast horizon. To construct the aggregate distributions for these surveys, the aggregate conditional distributions that are provided in the public pdf files are integrated.

The focus of this study is primarily the relation between the mean and mode of the aggregate distribution, or the discrepancy between the two (skewness). However, one can naturally compute statistics of the aggregate distribution beyond what we consider in the paper, which is left for future research.

To construct a measure of fed funds rate expectations from the Blue Chip Survey (BCS) that is directly comparable with the SPD measures, we use the consensus (average) estimates of the federal funds rate forecasts reported in the monthly Blue Chip Financial Forecasts. Since the BCS reports an average forecast over any given quarter, we assume the forecast for each quarter reflects the mid-quarter date, and we linearly interpolate two adjacent mid-quarter dates to derive a one-year ahead expectation from the survey date. To take into account the fact that the survey dates of the SPD and BCS do not coincide, we linearly interpolate the one-year ahead forecasts across the BCS dates to obtain forecasts that correspond to each survey date of the SPD.

We also use some unpublished results from the SPD for inputs in a regression analysis to make a clear comparison with the BCS expectations (See Section 4). Specifically, we use the mean and standard deviation across the respondents' expectations, including an anonymized subset.

Further details are deferred to Appendix A.

3 Time-Varying Skewness of the Aggregate Federal Funds Rate Distribution

3.1 Time Series of Skewness

The key empirical result of our paper is the time-varying skewness of the one-year ahead aggregate federal funds rate distribution (i.e., average of the respondents' distributions). In Figure 1, we show this based on a time series of two measures. One is the difference between the mean and the mode of the distribution for each available SPD since the beginning of 2011 (red), which is a simple and intuitive measure of skewness.

Despite its simplicity, the difference between the mean and mode is subject to several well-known issues; for instance, the mode can be sensitive to small changes in the distribution. Therefore, we also show a measure of skewness based on quantiles of the distribution (blue). Specifically, following Groeneveld and Meeden (1984), we integrate the interquantile measure of skewness for quantiles α in the range of $\alpha \in (0, 0.5)$ (henceforth, "GM skewness").[11] GM skewness does not rely on the mode, and is more robust than a measure using an arbitrary choice of the quantile. Although a widely preferred measure of skewness is based on the third moment of a standardized distribution, the measure has unbounded units which are hard to interpret, and can be sensitive to the standard deviation. In particular, we find that the skewness based on the third moment can look outsized especially when the fed funds rate distribution collapsed near the effective lower bound. In contrast, GM skewness appears more robust to large swings in the standard deviation as it is bounded between -1 and 1, and arguably provides a more balanced assessment of skewness over the sample period.[12] While the use of GM skewness appears to be limited in the economics and finance literature, papers such as Kim and White (2004) and Ghysels, Plazzi, and Valkanov (2016) highlight the advantages of this skewness measure in the context of financial asset returns.

We also discuss time series of alternative skewness measures in Appendix E, including the measure based on the standardized third moments, as well as the normalized (by the standard deviation of the distribution) difference between the mean and the mode/median (the latter two are the so-called "Pearson Skewness"). In addition, we analyze skewness from continuous distributions that fit the histograms of the survey distributions both parametrically and nonparametrically in Appendix F. The fitted distributions provide further robustness checks for the skewness measures.

Regardless of how we measure it, we observe significant time-varying skew in the aggregate distribution of the federal funds rate, resulting in sizable gaps between mean and modal expectations. The difference between the mean and mode can be as high as around 35 basis points and as low as about –50 basis points, equivalent to multiple 25 basis point changes in the policy rate.[13] For a sample period where the policy rate was close to zero for a majority of the time, the economic magnitudes of these differences seem large. As mentioned in the previous section, the expectations here are derived from "physical" distributions, free of risk premia, in contrast to other studies which focus on risk-neutral distributions. For reference, we compute corresponding skewness based on the implied risk-neutral distribution from Eurodollar options and find that although the broad pattern is similar to the skewness based on the SPD, there are also meaningful differences (see Appendix G for details).

[11] Mathematically, the measure is defined for a random variable X as:

$$\frac{\int_{0.5}^{1}(F^{-1}(1-\alpha)+F^{-1}(\alpha)-2F^{-1}(0.5))d\alpha}{\int_{0.5}^{1}(F^{-1}(1-\alpha)+F^{-1}(\alpha))d\alpha} = \frac{\mathbb{E}[X]-\mathbb{E}_{median}[X]}{\mathbb{E}|X-\mathbb{E}_{median}[X]|} \tag{1}$$

where F is the cumulative distribution function of X and $F^{-1}(\alpha)$ is the αth quantile of X (thus $F^{-1}(0.5)$ is the median).

We plot smoothed series by taking the centered moving average of three adjacent surveys. We do this merely for expositional purposes. All ensuing analysis is based on the raw series, unless otherwise stated. The raw series are shown in Appendix C, Figure C.1.

[12] For a time series of the standard deviation of the aggregate distribution, see Figure D.1 in Appendix D.

[13] Based on the raw survey data shown in Figure C.1 in Appendix C.

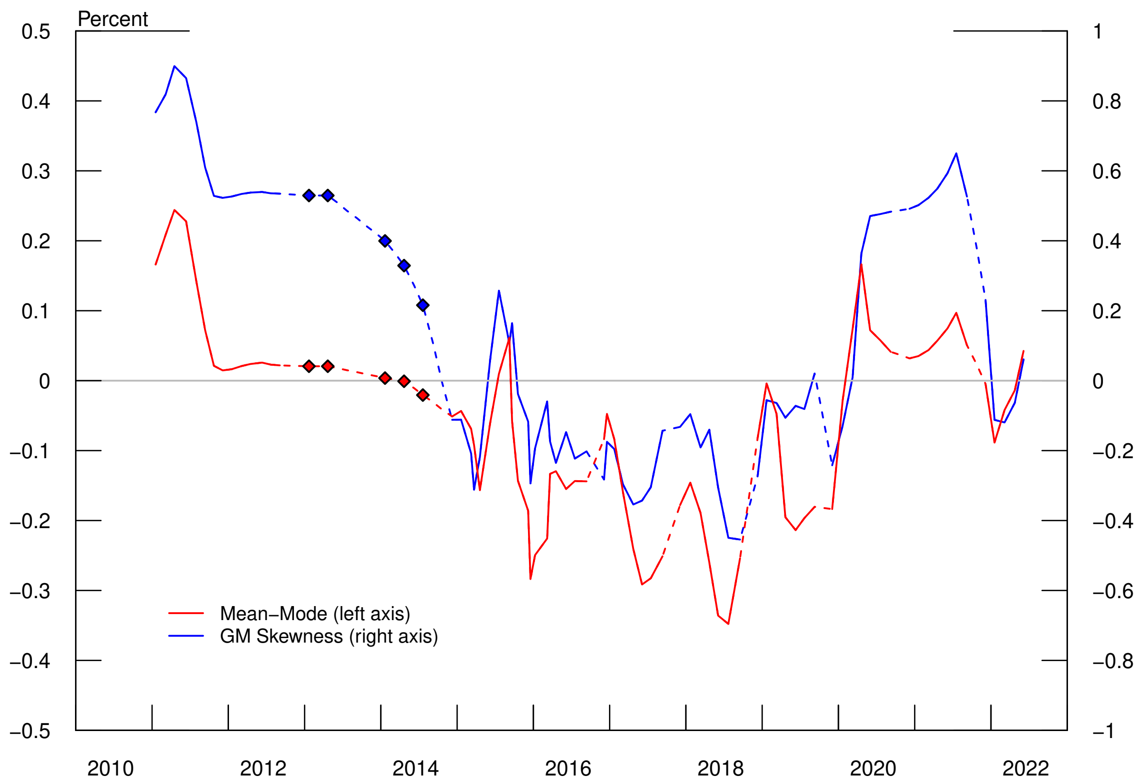

Figure 1: **Skewness of the Aggregate Federal Funds Rate Distribution from the SPD**

Note: Both series are smoothed by taking the centered moving average of three adjacent surveys. Whenever data from an adjacent survey is unavailable, we take the average using data up to two surveys apart. Dashed line segments are interpolations between survey dates where data is available. The diamonds indicate points that have no adjoining observations. Pearson Skewness has no units.

Following the financial crisis, skewness was significantly positive through mid-2011, but the mean – mode difference narrowed afterwards as it became increasingly likely that the fed funds rate would stay near the ELB for an extended period of time and the near-term fed funds rate forecasts trivially concentrated on near zero rates.

Skewness largely collapsed close to zero after the Fed announced date-based forward guidance in August 2011. The SPD mean forecast was generally somewhat higher than the mode even while the mean – mode difference was compressed till the end of 2014, and if the very small standard deviation of the expected rate distributions during this period is taken into account, the skewness still remained noticeably positive, as shown by the GM skewness. Positive skewness is consistent with the notion that rates are bounded from below; even a relatively small probability on higher rates would push the mean above the mode.

After liftoff at the end of 2015, the SPD modal forecast generally stayed higher than the mean, reflecting a shift towards a negative skew in the forecast distribution.[14] Skewness continued to

[14] Potter, Del Negro, Topa, and Van der Klaauw (2017) zoom in on the drop and rebound in the skew around the beginning of 2016, which can also be observed in Figure 1. The turn to negative skew in mid-2015 seems broadly consistent with the deterioration in the foreign economic outlook around that time (there was a large decline in equity

become more negative until mid-2019, after which it started to retrace. Apparently, survey participants perceived significant downside risk in policy rates during this period amid persistent trade tensions and a lack of pickup in inflation during the period. However, the mode is unlikely to capture increased likelihood of upside or downside outcomes that are reflected in the mean because the mode excludes information on the shape of the rest of the distribution. In March 2020, the policy rate returned to the ELB following the pandemic crisis, and the mean-mode difference moved sharply into positive territory, to a level observed back in 2011. More recently, as expectations for liftoff increased, the skewness of the distribution dropped significantly, similar to what we observed in 2015.

Figure 2 shows a more detailed snapshot of two recent data points, including the term structure of skew. We see that in January 2021, when the policy rate was at the ELB, the skew was very small due to a collapsed distribution in the near term, but increased with forecast horizon, with a very small probability on rates becoming negative, similar to the previous ELB period. However, in January 2022, as expectations for liftoff increased, the skew turned negative, with the absolute size of the skew largest for intermediate maturities. The skewness measures are generally well correlated across maturities. See Appendix H for a time series of the two-year ahead skewness measure.

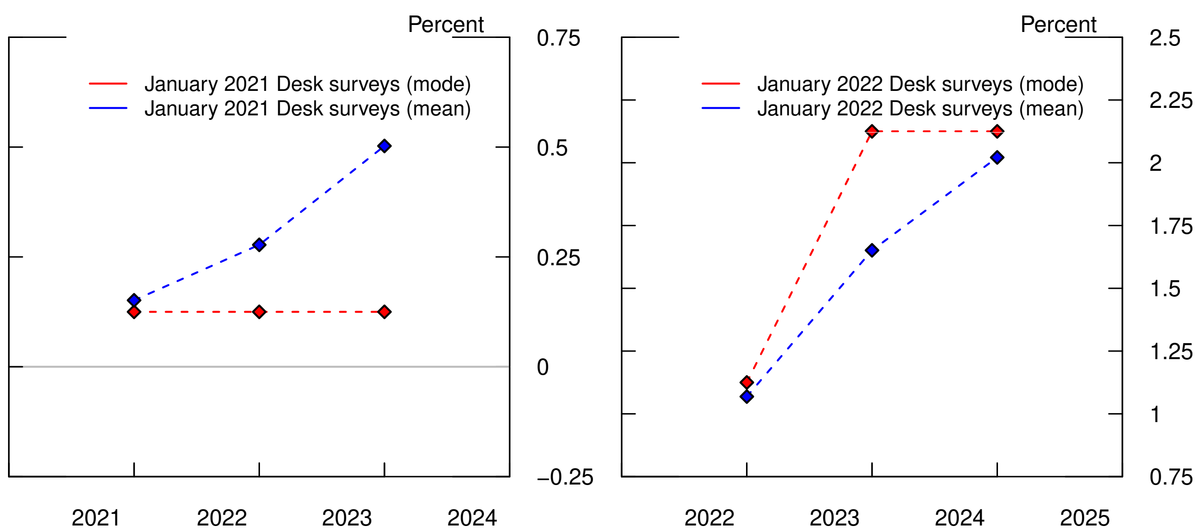

Figure 2: **Term Structure of Fed Funds Rate Expectations in 2021 and 2022**

Note: Means and modes are derived from the aggregate fed funds rate distribution for each year-end maturity.

3.2 Interpreting the Time-Varying Skewness with a Macroeconomic Model

In this section, we present a simple macroeconomic model to interpret more formally the time-varying skewness in the aggregate fed funds rate distribution documented above. The model is a

prices during the summer of 2015, before liftoff).

stylized New-Keynesian model in which the central bank implements monetary policy under the ELB constraint. We show that the observed pattern of time-varying skew is a natural consequence of the economy hitting the ELB, as well as the risk of hitting the ELB when the economy is near the ELB.

We keep the model simple so that the key intuition can be clearly understood. The New-Keynesian model features a representative consumer with log utility, firms facing nominal price rigidities in the form of Rotemberg adjustment costs, and a central bank that sets monetary policy based on a Taylor rule which only reacts to the deviation of inflation from the central bank's target. Monetary policy is subject to a ELB constraint that binds occasionally. The only exogenous shock to the economy is a shock to the time discount factor, which we interpret as a "demand" shock. All agents form expectations rationally. The model can then be summarized by the following equilibrium conditions, which have the standard "three-equation" form:

$$\mathbb{E}_t \left[\beta_t \frac{C_t}{C_{t+1}\Pi_{t+1}} R_t \right] = 1 \tag{2}$$

$$Y_t \left(\varphi \left(\Pi_t - 1 \right) \Pi_t - (1-\theta) - \theta w_t \right) \ = \mathbb{E}_t \left[\beta_t \frac{C_t}{C_{t+1}} Y_{t+1} \varphi \left(\Pi_{t+1} - 1 \right) \Pi_{t+1} \right] \tag{3}$$

$$R_t = \max \left[0, \Pi_t^{\phi_\Pi} \right] \tag{4}$$

with the aggregate resource constraint $Y_t = C_t + \frac{\varphi}{2} \left(\Pi_t - 1 \right)^2 Y_t$. $C_t, Y_t, \Pi_t, R_t, w_t$ are consumption, output, inflation, nominal one-period interest rate (i.e. policy rate) and (real) wage, all measured as deviations from their deterministic steady states. Equations (2) and (3) are the first order conditions for the consumer (Euler equation) and the firm (the Phillips curve), respectively, (4) is the monetary policy rule subject to the ELB constraint. The time discount factor β_t is exogenous and follows $\ln \beta_t = (1-\rho_\beta) \ln \bar{\beta} + \rho_\beta \ln \beta_{t-1} + \sigma_\beta \varepsilon_{\beta,t}$, where $\varepsilon_{\beta,t}$ is an i.i.d standard normal. The model is solved using a global solution method to take account of non-linearities properly. Further details of the model and calibration are discussed in Appendix I.1.

Figure 3 shows distributions of the time discount factor and the one-period nominal interest rate (i.e. a proxy for the policy rate) three periods ahead simulated from the model.[15] The left panel (Panel A) plots the three-period ahead distributions of the time discount factor based on three values of the current time discount factor. Since the shocks are normally distributed, the exogenous distributions are symmetric.[16] Within Panel A, the distributions to the left, in the middle and to the right respectively represent a state where the discount factor is small enough that the economy is far from the ELB and is likely to stay there in the near term (state I), a state where the discount factor takes an intermediate value such that the economy is away from the ELB,

[15] The model parameters use mostly standard values calibrated at a quarterly frequency, so the results presented here are only a rough approximation of the one-year ahead fed funds rate distribution. The model's main purpose is to capture the qualitative features of the distribution clearly and not necessarily to fit various aspects of the data.

[16] Although we do not explore the possibility in our study, we could alternatively assume that the exogenous shocks have asymmetric distributions. Studies such as Bekaert, Engstrom, and Ermolov (2021b) and Berger, Dew-Becker, and Giglio (2020) suggest this may be plausible. We view our work to be complementary to such an approach.

but still close enough that the likelihood of hitting the ELB in the near term is significant (state II), and a state where the discount factor is large enough that the economy is at the ELB and likely to stay there in the near term (state III).

The remaining three panels to the right (Panels B, C and D) show the three-period ahead distributions of the policy rate for the corresponding states of the economy. Although the model is simulated from an approximately continuous distribution of the discount factor shock, we collect observations in discrete bins each covering 25 basis point intervals to construct the distributions of the policy rate, in line with the design of the SPD.[17] The dark distribution in the top right panel is generated assuming the discount factor is at state I. In this scenario, the near-term policy rate distribution is essentially symmetric, showing little skewness, which can be seen from the roughly overlapping dashed-dotted lines, which represent the mean (grey) and mode (red) of the distribution, respectively.

The dark distribution in the middle right panel is generated assuming the discount factor is at state II. In this scenario, the discount factor takes an intermediate value where there is significant risk that the policy rate will hit the ELB over the near term. However, since the economy is not yet at the ELB, the fed funds rate distribution turns out to be skewed more towards lower rate outcomes. Indeed, we see that the mean (grey) is smaller than the mode (red) due to the large mass in the left tail. We explain this mechanism behind the negative skewness in further detail below. This scenario is consistent with the negative skew of the fed funds distribution observed for a few years after liftoff. The dark-colored distribution in the bottom right panel is generated assuming the discount factor is at state III. In this scenario, the policy rate is expected to stay at the ELB with a significant probability over the near term, and hence the distribution is positively skewed with a large mass at the effective lower bound. As shown again by the vertical lines, the mean (grey) is significantly higher than the mode (red). Indeed, such a scenario is consistent with the positive skew of the aggregate fed funds rate distribution observed during a majority of the ELB period, from 2011 through 2015 and after March 2020.[18] [19]

We emphasize that in our model the change in skewness is a result of the ELB constraint. To clarify this point, we also compute the policy rate distributions simulated from the same model without imposing the ELB constraint. The results are shown as light-colored distributions in Figure 3. We find that regardless of which state the economy is in, the distribution remains virtually symmetric.

It is worth noting that several well-known term structure models with the ELB constraint cannot generate short rate distributions with negative skewness.[20] In contrast, by accounting for

[17] Analogous to how we define the mode of the aggregate distribution in the SPD, the mode of the simulated distribution is defined as the mid-point of the bin with the largest probability.

[18] Even when the policy rate is (slightly) away from the ELB, the risk of reaching the ELB in the near future may be high enough that the simulated skew is still positive.

[19] If the time discount factor is large enough, the policy rate will almost certainly stay at zero within the forecast horizon. Then, the skew can become less positive as a result of the policy rate distribution collapsing trivially to a zero outcome. Such a case is consistent with the skewness decreasing around the time the Fed announced date-based forward guidance in August 2011.

[20] See Kim (2009) for a discussion. Although the paper focuses on risk-neutral distributions, the results also apply

richer general equilibrium dynamics, we show that the ELB constraint can be a *single* driver of *both* positive and negative skew in the policy rate distribution. Meanwhile, other papers offer mechanisms that lead to (endogenous) negative skew in distributions of various macroeconomic variables.[21] While these mechanisms are also likely in play, we show that the ELB constraint can be a driver of negative skew that is particularly relevant for the policy rate at or near the ELB.

The mechanism of time-varying skew due to the ELB constraint may be understood more clearly by examining the model's equilibrium decision rules, shown in Figure 4. The blue dashed lines in the two panels plot the decision rules for the policy rate and inflation, respectively, when the model is not subject to a ELB constraint, while the solid lines plot the corresponding decision rules when the model is subject to a ELB constraint.[22] The decision rules are essentially (locally) concave or convex transformations of the discount rate, which, by definition, has a symmetric conditional distribution. To map this transformation into a change in skewness, we can resort to a straightforward corollary of a theorem shown by van Zwet (1964):

Result (Corollary of Theorem 3.1 in van Zwet (1964)): *Suppose random variable x is symmetrically distributed. If ψ is a convex (concave) function on interval I, which is not constant on I, and if moments up to the third order exist, then* $\frac{\mu_3(\psi(x))}{\mu_2(\psi(x))^{3/2}} \geq (\leq) 0$ *on I where* $\mu_k(x) \equiv \mathbb{E}[(x - \mathbb{E}[x])^k]$ *(i.e., $\psi(x)$ is positively (negatively) skewed).*

In other words, a convex or concave function of a symmetric normal variable implies skewness that is not zero.[23] Consider the decision rules when the model is not subject to a ELB constraint (the dashed lines). In this case, the decision rules are essentially linear. A larger time discount factor leads the consumer to save more and consume less, lowering demand and causing disinflation. The near-linearity of the decision rules preserves the normality of the time discount factor distribution, and the policy rate distribution exhibits little skew regardless of the state of the economy, as shown in Figure 3.

In contrast, the solid lines show the decision rules when the model respects the ELB constraint. As shown in the left panel, the policy rate is locally convex at the point where the ELB constraint binds, as the policy rate cannot decrease further. This maps the symmetric distribution of the time discount factor into a positively skewed distribution of the policy rate with a significant mass at zero.

---

largely to physical distributions, since the risk-neutral and physical processes take similar forms in these models. For example, a common shadow rate process $ds = k^Q\left(\mu^Q - s\right)dt + \sigma dB^Q$ under the risk-neutral measure with a risk price process $\lambda = \lambda_a + \lambda_b s$ behaves as $ds = k^P\left(\mu^P - s\right)dt + \sigma dB^P$ under the physical measure where $k^P = k^Q - \sigma\lambda_b$, $\mu^P = k^P(\mu^Q k^Q + \sigma\lambda_a)$. Meanwhile, the positive skew of the short rate in shadow rate models is well-documented in the literature (e.g. Kim and Singleton (2012), Bauer and Rudebusch (2016)).

[21] Jensen, Petrella, Ravn, and Santoro (2020) show that a model with credit constraints can explain negative skewness in aggregate variables such as output growth.

[22] We do not show the decision rule for output (and consumption) since the discussion is similar to that of inflation.

[23] The result does not require continuity and hence differentiability of ψ, which is convenient for considering policy functions with kinks. Note the result is proven for the skewness measured by the centered third moment, but the result is still applicable to our discussion, which does not rely on the specific measure of skewness.

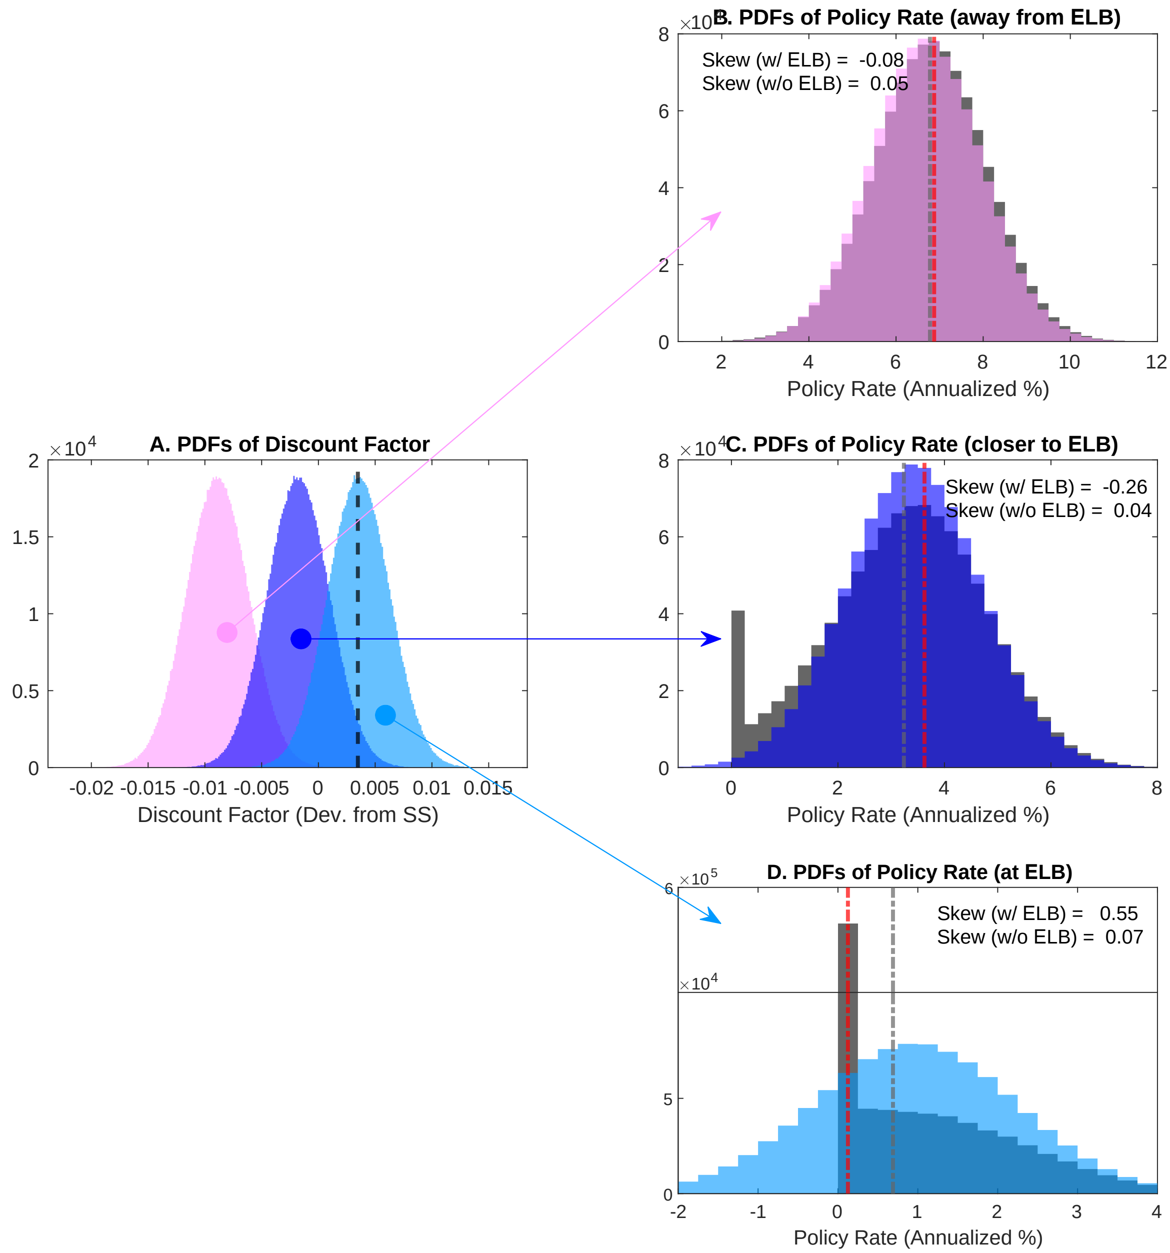

Figure 3: **Simulated Policy Rate Distributions from the New-Keynesian Model**

Note: The vertical dashed line in the left panel shows the value of the discount factor (its deviation from the steady state) when the ELB constraint binds. The dark-colored probability distributions in the right panels are generated from the model with the ELB constraint. The light-colored pdfs in the right panels are generated from the model without the ELB constraint. The dashed-dotted lines are the mean (grey) and mode (red) of the pdfs from the model with the ELB constraint.

However, as we move to a state of a lower discount factor, the policy rate rises above the ELB and the policy rate function becomes locally concave. This is because output and inflation are subject to significant downside risk near and at the ELB where monetary policy cannot lower interest rates further to effectively counter the recession. Indeed, the decision rule for inflation (solid line in the right panel) drops steeply as the economy reaches the ELB.[24] In turn, the central bank is expected to lower the interest rate aggressively to counter the downside risk until it reaches the ELB.

When the discount factor is small enough that the economy is far away from the ELB, there is a negligible chance that the economy will reach the ELB in the near term. Without a meaningful effect of the ELB constraint, the decision rules are essentially linear, and hence the interest rate distribution preserves the normality of the discount factor distribution, similar to the case when the ELB constraint is not imposed.

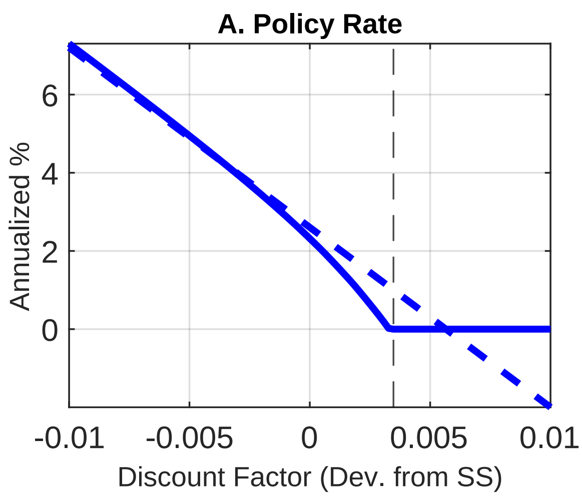

Figure 4: **Equilibrium Decision Rules for the Policy Rate and Inflation**

Note: The vertical dashed lines show the value of the discount factor (its deviation from the steady state) when the ELB constraint binds. The solid blue lines and the dashed blue lines show decision rules for the model with the ELB constraint and without the ELB constraint, respectively.

It is natural to focus on the demand shock because it is broadly considered to be the empirically relevant shock that drove the policy rate to the ELB, following both the financial crisis in 2008 and the pandemic crisis in 2020. Abstracting from other shocks also has the benefit of parsimony, which allows for a transparent explanation of how the ELB constraint can generate endogenous skewness. However, particularly in light of the pandemic crisis, during which supply shortages have

[24] The strong concavity of the decision rule for inflation/output implies a negatively skewed inflation/output distribution. The additional downside risk to inflation/output induced by the ELB constraint is well-known in the literature (e.g. Reifschneider and Williams (2000)). Bekaert, Engstrom, and Ermolov (2021b) provides empirical estimates of conditional skew of inflation and output growth, which show discrete drops in skew around 2015.

been widely reported, it may still be interesting to see how our analysis extends to a setup with "supply" shocks as well.

We introduce supply shocks in the standard fashion as a shock to the total factor productivity in the production function. It turns out that supply shocks alone can also generate similar endogenous skew when the ELB constraint is present. The result further extends to a setup which is buffeted with both demand and supply shocks. We find that a combination of a negative supply shock and a positive demand shock (or a negative shock to the discount factor) that lifts the policy rate from the ELB can generate a transition from a positive skew to negative skew in the forecast distribution of the policy rate, qualitatively consistent with the data. We defer details about this extension to Appendix I.

In sum, the ELB constraint can naturally produce time-varying skewness in the policy rate distribution which can be either negative or positive. When the economy is far away from the ELB, the model would predict that skewness is close to zero. When the economy moves closer to the ELB, we observe negative skewness due to the increased downside risks related to monetary policy being unable to effectively counter a recession. When the economy approaches the ELB, the model intuitively generates positive skewness as rates are bounded from below. If the economy has little chance of leaving the ELB within the forecast horizon, the distribution will trivially collapse, and skew will be smaller.

[Results of the extended model with symmetric second-moment shocks to be added.]

3.3 Survey Question Design and Reliability of the Responses

As described above, the format of the survey questions for the fed funds rate distribution has evolved across the sample period. In particular, while the format prior to 2016 was to ask for the distributions for certain future dates without further conditionalities (henceforth "unconditional distributions"), the format changed from January 2016 to a more complicated structure where each respondent was asked to report distributions conditional on specific economic scenarios such as whether or not the economy returned to the ELB by a certain date, or the direction and timing of the Committee's next policy action. This format continued up to the October 2019 survey, after which the format reverted to asking for unconditional distributions. Therefore for this period, we need to integrate the conditional distributions to recover the unconditional distributions. Due to this change in the question format, it may be natural to wonder how this affects the distributions in its reliability as well as its consistency with the other unconditional distributions in the sample.

However, while there may be some limits on how effectively the respondents can digest certain question formats, we believe the significant negative skew observed from 2016 through 2019 should be attributed mostly to factors other than the change in the question format, for several reasons.

First, we began to see negative skew in 2015, even before the survey format changed. Also, while we saw a drop in the mean-mode difference when the conditional distribution was first asked in the January 2016 survey, the decrease in the market rate (the one-year ahead OIS forward rate) was virtually the same magnitude. Moreover, when the survey format switched back to asking unconditional distributions in December 2019, we did not observe any discrete jump in skewness, which could have possibly occurred if the question format was creating downward bias in skew.

Second, as mentioned in Section 3.1, the decrease in skew towards lower rate outcomes through 2016-2019 seems broadly consistent with the overall sentiment during the period when risk toward the downside was perceived to be meaningful amid persistent trade tensions and lack of pickup in inflation. An extended period without a recession may also have caused some fear of the next recession looming, and given that rates were still less than 3 percent during the period, the experience of the Global Financial Crisis (and also in hindsight), may have fostered expectations for a rapid return to the ELB in the event of a recession. Indeed, a significant probability of returning to the ELB was persistently reported in the surveys at the time.

Third, as we show later, the mean forecast, which incorporates the ELB concern ended up being a better forecast of the realized FFR a year later compared to the mode, which may be another reason to believe the estimates were sensible. And fourth, our theoretical model-based exercises suggest that we should expect to see meaningful negative skew when we are above but near the ELB. For these reasons, we believe the measures of expectations from the SPD are largely reliable, including the period from 2016 to 2019.

4 Implications of Time-Varying Skewness in the Average Fed Funds Rate Distribution

4.1 What do the Expectations from the Blue Chip Survey (BCS) Represent?

4.1.1 BCS Expectations versus the Modes and Means from the SPD

If the forecast distribution of the fed funds rate has significant skew, it becomes important to distinguish whether certain measures of policy expectations are defined as means or modes.[25] However, unless the respondents are explicitly asked to provide certain statistical measures as their forecasts, it is unclear what type of central tendency a forecast represents. This appears to be the case with the fed funds rate expectations taken from the BCS—a widely used measure of policy expectations. In this section, we infer whether the BCS expectations (the "consensus" estimates) are more representative of mean or modal expectations by looking at how they relate to measures constructed from the SPD.

There are a couple important issues to consider when we compare the results from the BCS with those from the SPD. First, the BCS consensus forecast is the cross-sectional average of each respondent's expectation (i.e., mean, mode or other measures of central tendency). Therefore, we must look to the cross-section of the SPD rather than the aggregate distribution (which was the focus of our previous sections). It turns out the cross-sectional average of the means of each respondent's distribution is mathematically equivalent to the mean of the aggregate distribution,

[25] The expectations could also represent medians, but we believe it is less likely for respondents to consider them as a representation of their point forecasts, since medians are likely to be less intuitive than the mean or the mode, and in terms of the SPD, the median appears to be rarely asked for (in contrast to the mode).

but this is not true for the modes. In other words, we need to compare the BCS expectation with the average of the modes of each respondent's distribution and not the mode of the aggregate distribution.

Second, the pool of respondents differs substantially across the two surveys. Therefore, we should compare the measures based on information only from respondents who have provided data for both surveys in each given survey month. This will allow us to rule out any discrepancies arising from the difference in the set of respondents across surveys.[26] Since the cross-sectional information of the distribution is not publicly available, we use some unpublished survey results to conduct our analysis.

Table 1 shows the absolute average difference between the federal funds rate expectations computed from the BCS and the average across the modal expectations computed from the SPD (left panel, first column) as well as the absolute average difference between the BCS and the mean of the aggregate distribution (second column).[27] [28]

Average Absolute Difference betw. SPD and BCS (in bps)		
	Mode	Mean
Subset	6.86	17.52
Full Sample	10.48	17.92

Average Absolute Difference betw. σ(SPD) and σ(BCS) (in bps)		
	Mode	Mean
Subset	4.04	6.88
Full Sample	7.22	9.80

Table 1: **Summary Statistics of SPD versus BCS**

Note: This table shows the absolute average difference between the consensus Blue Chip Financial Forecast one-year ahead forecast of the federal funds rate and the average of the modal paths and mean coming from the aggregate forecast distribution of the Survey of Primary Dealers. The time period is 2011-2022 and the units are in basis points. The full sample includes all respondents, while the subset uses only the respondents that participated in both the Blue Chip Survey and the Survey of Primary Dealers.

We find that the average of the modes from the SPD tracks the BCS measure quite closely for most of the sample period. Since the beginning of 2011, the average absolute difference between BCS expectations and the average of the modes from the SPD is 6.86 basis points, while the average absolute difference between BCS expectations and the means from the SPD is notably higher at 17.52 basis points. In the last five years, the average gap between the BCS and the mean widens

[26] Note there could be other reasons that cause the results from the BCS consensus and those from the SPD to diverge. The linear interpolation of the BCS across survey dates to match the SPD survey date is one of them.

[27] As the measure of each respondent's mode, we use responses from the question about the modal path instead of the distribution, as the question format is better aligned with the BCS.

[28] It turns out that empirically, the cross-sectional mean of the modes track the mode of the aggregate distribution quite closely for most of the sample period.

to 25 basis points, while the gap between the BCS and the mode remains small (not shown). For reference, we also show the average absolute differences using the full set of respondents. We find that while the difference between the BCS and the SPD mode is smaller than the difference between the BCS and the SPD mean, the former difference is meaningfully larger compared to using the matched sample.

A similar point can be made with respect to the dispersion across each respondent's expectations. The right panel of Table 1 shows the average absolute difference between the cross-sectional standard deviation ("disagreement") of the respondents' forecasts from the BCS, SPD means, and SPD modes. Once again, we use the subset of firms that are in both the BCS and the SPD to control for the discrepancy in the set of respondents across the two surveys. Similar to what we saw for expectations, we find that the disagreement in BCS tracks the disagreement in the SPD mode more closely than the disagreement in the SPD mean. In addition, matching the set of respondents across surveys reinforces the result. This provides further evidence that the two measures coincide beyond their implications for the first moment.

Dependent Variable: BCS			
Mode	1.00^{***}		0.87^{***}
	(0.01)		(0.06)
Mean		1.18^{***}	0.16^{**}
		(0.04)	(0.07)
Constant	0.05^{**}	-0.03	0.04
	(0.02)	(0.04)	(0.03)
R-squared	0.99	0.98	0.99

Dependent Variable: Δ BCS			
Δ Mode	0.85^{***}		0.80^{***}
	(0.10)		(0.16)
Δ Mean		0.87^{***}	0.05
		(0.11)	(0.11)
Constant	0.01	0.02	0.01
	(0.01)	(0.01)	(0.01)
R-squared	0.72	0.66	0.72

Dependent Variable: σ(BCS)			
σ(Mode)	0.92^{***}		0.84^{***}
	(0.07)		(0.12)
σ(Mean)		1.33^{***}	0.17
		(0.12)	(0.12)
Constant	0.04^{***}	0.02	0.03^{**}
	(0.01)	(0.02)	(0.01)
R-squared	0.88	0.69	0.88

Dependent Variable: Δ σ(BCS)			
Δ σ(Mode)	0.75^{***}		0.75^{***}
	(0.13)		(0.16)
Δ σ(Mean)		0.67^{***}	0.00
		(0.20)	(0.14)
Constant	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)
R-squared	0.70	0.21	0.70

Table 2: **Regression Results**

Note: Numbers in parentheses denote Newey-West standard errors. Stars next to the regression coefficients indicate statistical significance levels. ***: 1%, **:5%, *:10%.

To further establish the relation between the BCS expectation and the mean and modal expectations from the SPD, we regress the levels of the BCS measure onto the levels of the mean and modes from the SPD, as shown in the upper left panel of Table 2. The level results show the mode has a significant coefficient close to 1 and when both the mode and mean are included, the coefficient for the mode is much closer to one than the coefficient for the mean. We see a similar dynamic hold for the second moments. In the bottom left panel, the standard deviation of the mode is highly significant while the standard deviation of the mean is insignificant and the opposite sign when both are included. The R-squared improves substantially when the standard deviation of the mode is included in the regression.

Due to concerns about the highly persistent nature of the time series, we also run similar exercises for the changes in these variables. The upper right panel shows results for the regression of changes in the BCS measure onto the changes in the mean and modes from the SPD.[29] Both the mean and mode have significant explanatory power when included in the regression individually, but the mean loses statistical significance when the regression also includes the mode. Meanwhile the mode retains its strong statistical significance with a coefficient closer to one compared to the mean. The R-squared is also higher when the mode is included in the regression.

As discussed above, part of these results are strengthened by the fact that we match respondents across the SPD and BCS. We discuss the regression results using the full set of respondents in Appendix J.1. As a robustness check, we also repeat our analysis excluding the period since the pandemic crisis, and find that the results continue to hold, with a somewhat larger benefit of matching respondents. Further details are discussed in Appendix J.2.

Potentially, the cross-sectional average of individual forecasts from the BCS may serve as a good proxy for the mean of the aggregate distribution, but we find this is not the case. There is a wide gap between the BCS consensus forecast and the SPD mean. Given that the difference in the mean and modal expectations from the SPD is sizeable, this suggests that the consensus forecasts from the BCS may be better interpreted as a measure of modal expectations, and some caution is warranted in treating such a measure as representative of mean expectations. For example, a term structure model that uses the BCS as a guide to measure interest rate expectations may provide estimates that are closer to modal expectations (see also Section 4.2).[30]

Given the evidence in this section on the proximity of BCS forecasts to the modes from the SPD, one may wonder why BCS respondents report modal outcomes instead of means. Although a thorough investigation is out of scope for this paper, we can point to a few possibilities. One is that the mode is simply more natural and intuitive to report. Another is that it could be a result of optimal reporting based on minimizing an "all-or-nothing" loss function, the solution to which

[29] Whenever data from the survey immediately before a particular survey is unavailable, we take the difference with respect to the closest observation. Since this results in changes which cover different periods of time, we check whether our results are affected by this issue by running the regressions excluding any observation that does not have an observation one survey before. We find that the results still hold.

[30] Table 2 in Engelberg, Manski, and Williams (2009) indicates that the point estimates for GDP growth and inflation in the US-SPF are closer to the mode than the mean, but the difference is very small and is not acknowledged in the text by the authors. Huang, Pilbeam, and Pouliot (2022) provide some suggestive evidence that the point estimates for the corresponding variables in the ECB-SPF are closer to the mode than the mean, although this distinction is not the focus of their study and the statistical significance of this distinction is not shown.

is known to be the modal forecast.[31] For example, a forecaster may care about obtaining a perfect forecast to gain publicity with occasional users.[32]

In any case, the fact that the mode is justified by a rather extreme form of the loss function appears to highlight the importance of considering the mean expectations explicitly, when the distribution is asymmetric.[33] The relevance of the mean is deeply rooted in expected utility theory and can be rationalized by a standard quadratic loss function. Indeed, in Section 4.3, we show the means from the SPD outperform the modes based on this criterion.

4.1.2 Comparison of other Moments

We can also compute a measure of skewness from the BCS by taking the difference between the cross-sectional average (consensus) and mode of the forecasts. Note this measure of skew is fundamentally distinct from the measure using the aggregate distribution, but one may be interested to know if an empirical pattern similar to what we document using the aggregate distribution from the SPD can be obtained from the BCS.

We show in Figure 5 the difference between the cross-sectional average (consensus) and the cross-sectional mode of each respondent's forecast in the BCS.

[31] The loss function can be defined using the Dirac Delta function (Gneiting (2011)).

[32] See, for example, Laster, Bennett, and Geoum (1999). The BCS selects the Lawrence Klein Best Forecaster, which is awarded to a single forecaster each year (the rest of the forecasters receive nothing). Such a prize may incentivize the forecaster to report the mode, assuming a (near-)perfect forecast is necessary to win the award.

[33] Such a view is shared by Svensson (2001) in the context of inflation-forecast targeting.

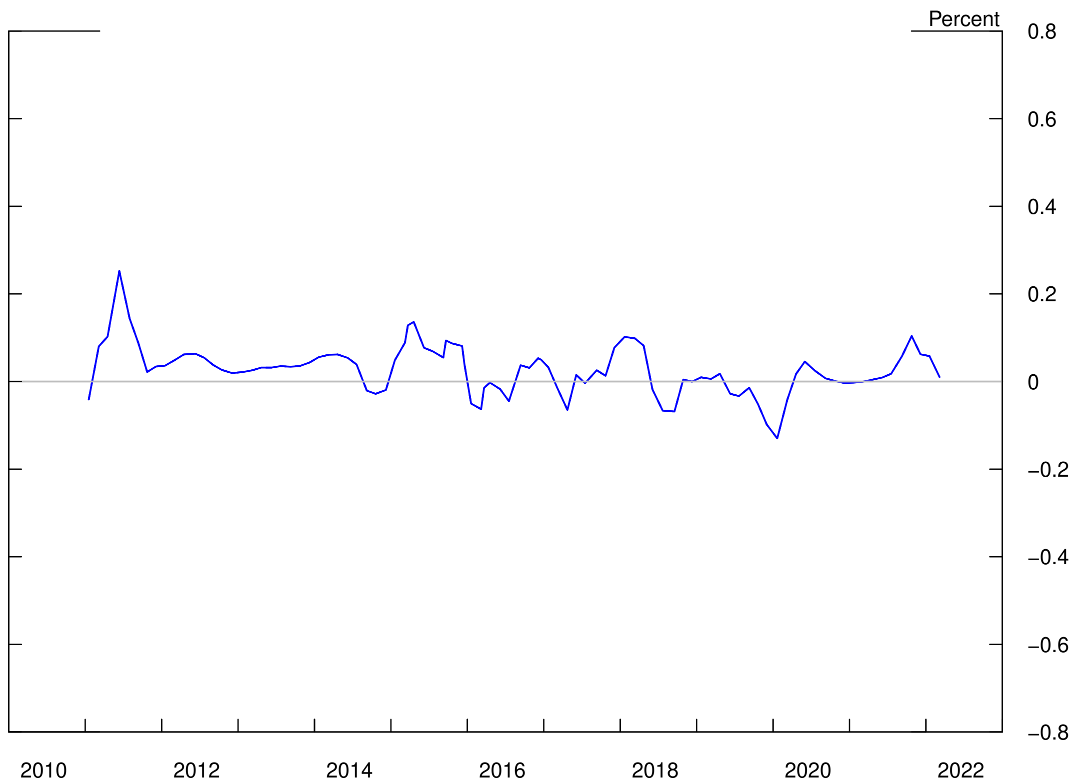

Figure 5: **Difference between the Cross-Sectional Mean and Mode of the BCS forecast**

Note: The series is smoothed by taking the centered moving average of three adjacent surveys.

We see little resemblance to the mean–mode difference coming from the SPD's aggregate distribution (see Figure 1). For example, from 2016 through 2020, the average gap between the mean and mode for the BCS is negligible at about 2 basis points. In contrast, the mean-mode gap averages -20 basis points from the SPD over the same period, a ten-fold increase in magnitude with the opposite sign. While at first it may seem odd that the mean-mode gap coming from the BCS is fairly different from the SPD, it is important to remember that each respondent's forecast in BCS is likely capturing most likely outcomes. Hence, it is not obvious that the average "most likely outcome" and its gap with respect to the most likely "most likely outcome" should proxy for the same values that one would derive from an aggregate probability distribution.[34] This measure of "skew" does not seem to adequately capture much downside/upside risks to the outlook.

[34] A similar point can be made about a related exercise using the Survey of Professional Forecasters (SPF). The SPF provides a time series of the mean and median of the 3-month T-bill "projections" across each respondent. We find that skewness measured by the mean minus the median follows a broadly similar pattern as our measure for the fed funds distribution, but the magnitude is significantly smaller. Needless to say, the forecast distribution of the 3-month T-bill yield is inherently different from that of the fed funds rate as evident from the variation in the spread between the T-bill and OIS rates as well as their forward counterparts.

4.2 Term Premia

The time-varying skewness of the average fed funds rate distribution implies that measures of policy expectations differ depending on whether they are measured based on the mean or the mode. In turn, term premium estimates may also differ. If we accept the standard definition of term premium as the residual of the forward rate after subtracting the expected rate measured by the mean, we can further decompose the forward rate into modal expectation, term premium and simple skewness defined as the mean minus the mode normalized by the volatility of rates as follows:

$$
\begin{aligned}
f_t^\tau &= \overbrace{\hat{\mathbb{E}}_t[r_{t+\tau}] + \sigma_t[r_{t+\tau}]Sk_t[r_{t+\tau}]}^{\mathbb{E}_t[r_{t+\tau}]} + tp_t^{mean} \\
&= \hat{\mathbb{E}}_t[r_{t+\tau}] + \underbrace{\sigma_t[r_{t+\tau}]Sk_t[r_{t+\tau}] + tp_t^{mean}}_{tp_t^{mode}}
\end{aligned}
$$

where $\hat{\mathbb{E}}_t\left[r_{t+\tau}\right]$ is the mode and $Sk_t\left[r_{t+\tau}\right] \equiv (\mathbb{E}_t\left[r_{t+\tau}\right] - \hat{\mathbb{E}}_t\left[r_{t+\tau}\right])/\sigma_t[r_{t+\tau}]$ is the simple skewness of the aggregate fed funds rate distribution. In other words, the "term premium" computed from modal expectations, tp_t^{mode}, includes a fraction of simple skewness in addition to the standard measure of term premium (correctly measured).[35]

Figure 6 shows the implied one-year ahead forward term premia computed from the one-year ahead (instantaneous) OIS forward rate and measures of expectations using the mean and mode of the aggregate distribution from the SPD, as well as the consensus estimate from the BCS.

[35] In fact, the same decomposition holds for GM skewness if we replace $\hat{\mathbb{E}}_t\left[r_{t+\tau}\right]$ with the median and $\sigma_t[r_{t+\tau}]$ with $\mathbb{E}_t|r_{t+\tau} - \hat{\mathbb{E}}_t\left[r_{t+\tau}\right]|$.

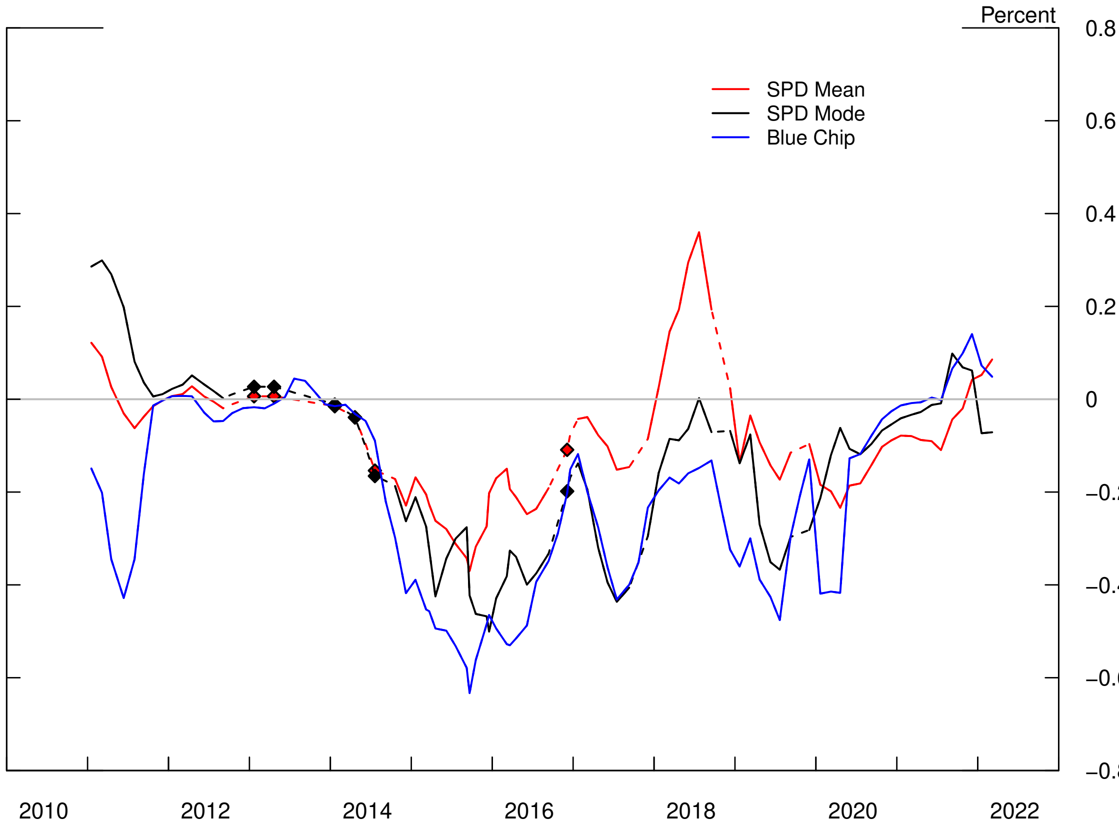

Figure 6: **One-year Forward Term Premia Implied by Various Policy Expectations**

Note: All series are smoothed by taking the centered moving average of three adjacent surveys. Whenever data from an adjacent survey is unavailable, we take the average using data up to two surveys apart. Dashed line segments are interpolations between survey dates where data is available. The diamonds indicate points that have no adjoining observations.

Our mean expectations extracted from the SPD generally imply significantly less negative term premia on average, compared to term premia implied by BCS forecasts since 2011, suggesting an alternative assessment of interest rate risk perceived over the period.[36] As discussed in section 4.1, the BCS forecasts are more representative of modal forecasts in recent years, and this is confirmed by the fact that term premia implied by the SPD mode tracks term premia implied by the BCS closely since 2012.[37] With the BCS forecast acting as an important input into several popular term structure models, it logically follows that the term premia generated from these settings may also imply more negative term premia than the mean expectations.[38]

[36]The averages of term premia based on the SPD mean and BCS were -8.6 and -21.2 basis points, respectively.

[37]A large part of the discrepancy between the BCS-implied term premia and the SPD mode-implied term premia before 2012 can be attributed to the fact that the BCS is more representative of the average of modes as opposed to the mode of the aggregate distribution.

[38]See, for example, Priebsch (2017).

4.3 Forecast Performance

4.3.1 A Visual Comparison

To further understand the usefulness of the newly computed mean expectations, we compare their forecast performance with alternative measures of policy expectations. First, we visualize how the measure of mean expectations performed as forecasts in hindsight, by plotting the mean and modal expectations from the aggregate distribution of the SPD as well as the consensus expectations from the BCS along with the realized effective fed funds rate (Figure 7).

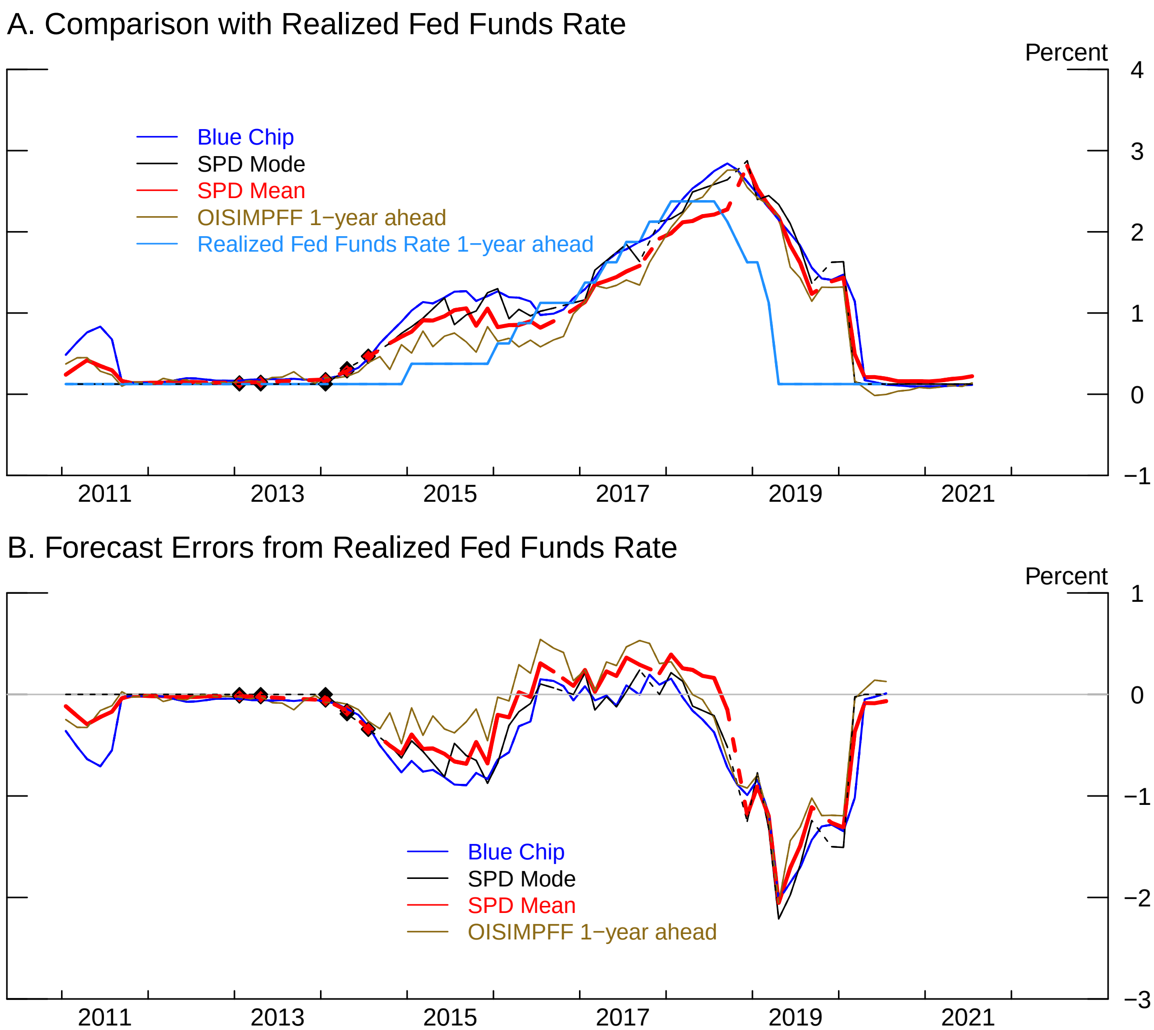

Figure 7: **Forecast Performance of Various Measures of Policy Expectations**

Note: Dashed line segments are interpolations between survey dates where data is available. The diamonds indicate points that have no adjoining observations.

For the first half of 2011, all measures of policy expectations consistently underestimated the pace of policy accommodation of the following year, except for the mode of the aggregate distribution, which correctly predicted rates to stay at the ELB. Surveys conducted from mid-2011 through 2013 forecasted the fed funds rate generally well, likely because the Fed's date-based forward guidance initiated in August 2011 helped anchor near-term policy expectations.

Survey measures in 2015 predicted the fed funds rate to increase significantly faster than the actual pace of tightening. In this regard, a straight read of the OIS-implied forward rates turned out to be the best predictor. However, the SPD mean generally ended up being the better predictor among the survey counterparts with the smallest (absolute) forecast errors. Meanwhile the BCS forecasts had the largest error, significantly over-predicting the actual federal funds rate. During 2017 and mid-2018, the SPD mode and the BCS forecasts still predicted a higher level of rates compared to the SPD mean, but this time, the realized fed funds rate a year later ended up relatively close to what the SPD mode and the BCS had predicted.

Beyond the second half of 2018, a general pattern of forecast errors seen during policy accommodation in response to the financial crisis reemerged, as the forecasts failed to predict the pace of accommodation again. BCS forecasts had predicted an increasing path of the policy rate with levels higher than the SPD mean through 2019, but this forecast ended up largely false as the fed held rates steady after the end of 2018 and started to lower its policy rate in July 2019. Although the SPD mean forecasts also predicted higher rates than was realized, it had put more weight on the possibility of lower rate outcomes, and the error was significantly smaller. More recently, after the policy rate returned to the ELB, all forecasts declined substantially, and the gap between the SPD mean and the BCS forecasts has narrowed accordingly.

4.3.2 Comparison based on the MSE Criterion

In this section, we formalize our discussion above on the forecast performance of various expectations by computing the root mean-squared forecast errors (RMSE) of each measure over the sample period. In terms of minimizing the RMSE, the (conditional) mean is theoretically the optimal forecast, and not the mode. We compute the RMSE for the full sample period as well as from 2011 to the period prior to the pandemic. The results are summarized in Table 3.

RMSE (Unconditional)				
	SPD Mean	SPD Mode	BCS	OIS-imp-FF
RMSE (2011-)	0.566	0.621***	0.699***	0.510***
RMSE (2011-2019.1)	0.372	0.391	0.509***	0.312*

RMSE (Conditional on SPD Mean - SPD Mode ≤ -12.5 bps)				
	SPD Mean	SPD Mode	BCS	OIS-imp-FF
RMSE (2011-)	0.801	0.915***	0.936***	0.747*
RMSE (2011-2019.1)	0.334	0.414*	0.450**	0.329

Table 3: **Forecast Performance of Various Expectation Measures**

Note: SPD Mode refers to the mode of the aggregate distribution. The stars next to the RMSE values show levels of statistical significance in the difference with the SPD mean forecast based on the Diebold-Mariano test. ***: 1%, **:5%, *:10%.

We find that for both sample periods, the mean expectations from the SPD outperform the modal expectations from the SPD and BCS expectations.[39] As discussed in Section 4.1, the BCS expectations are more representative of modal expectations, so this is consistent with what theory would predict. In addition, the RMSE of the mean expectations fares reasonably well against the top performer among our contestants, which is the one-year ahead OIS forward rate. In the second row, we show results using the sample period from January 2011 to January 2019, which excludes the sharp decline in rates following the onset of the pandemic. We find that the results generally hold with smaller RMSEs across all measures. Our results suggest that the common practice of evaluating forecast errors based on the assumption that the BCS forecasts represent means may lead to some overstatement of the magnitude of those errors.

We further investigate whether the differences in the forecasts are also statistically significant. Table 3 also indicates levels of statistical significance in the difference between the SPD mean forecast and the other measures of expectations, based on the Diebold-Mariano test (Diebold and Mariano (1995)). We find that the meaningful economic difference in the forecast errors largely translates to statistical significance.[40] In particular, evidence that the SPD mean outperforms the BCS expectations is statistically strong.

Part of our sample includes the period when the Fed's forward guidance anchored policy expectations strongly at the ELB, in which case we cannot expect to find significant differences in the

[39] We use the data from the full set of respondents in the BCS and SPD here. The results do not materially change if we used data matching the set of respondents in the BCS and SPD.

[40] To account for potential small-sample bias, we also checked statistical significance using the correction proposed by Harvey, Leybourne, and Newbold (1997) and found that the significance weakens only slightly.

forecast performance across measures. Also, it is interesting to see how forecast performance compares precisely when there is meaningful skew in the forecast distribution. Therefore, we also show the RMSE conditional on the mean and mode difference being at or below -12.5 basis points. This threshold corresponds to half of a 25 basis point rate cut, which seems to represent an economically meaningful skew in the respondents' outlook of the fed funds rate, but small enough such that we do not exclude too many observations.[41] We find that the results without conditioning is further magnified, in particular, the better forecast performance of the SPD mean compared to the SPD mode appears to be somewhat starker. We hence conclude that the mean expectation is relatively useful for forecasting purposes.

5 Conclusion

We document some novel empirical evidence of significant time-varying skewness in the aggregate forecast distribution of the federal funds rate, i.e. *asymmetric monetary policy expectations*. To this end, we construct measures of the one-year ahead expectations from responses to the Survey of Primary Dealers (SPD). The SPD provides a "physical" future distribution, in contrast to measures extracted from asset prices. Importantly, this allows us to explicitly compute mean and modal expectations and the discrepancy between the two measures, free of risk premia. We further show that a simple New-Keynesian model with the ELB constraint can endogenously generate both positive and negative skewness similar to patterns in the data.

The time-variation of asymmetry in the aggregate distribution highlights the importance of correctly measuring the mean when extracting FFR expectations from surveys. We argue that the fed funds rate forecasts from the Blue Chip Survey (BCS), a popular survey measure of monetary policy expectations, track the mode more closely than the mean since 2011, when the data becomes publicly available. This finding suggests some caution in treating the BCS forecasts as representative of mean expectations, which is not uncommon in the literature. As a result, the mean measure of policy expectations extracted from the SPD implies significantly less negative term premia on average, compared to term premia implied by BCS forecasts. The mean measure also outperforms the BCS forecasts based on the mean squared error loss, consistent with the theory of optimal forecasting.

References

Adam, K., D. Matveev, and S. Nagel (2021): "Do survey expectations of stock returns reflect risk adjustments?," *Journal of Monetary Economics*, 117, 723–740.

Bauer, M. D., and M. Chernov (2021): "Interest rate skewness and biased beliefs," Discussion paper, National Bureau of Economic Research.

[41] We do not consider conditioning on significant positive skew since the number of observations in which the mean – mode exceeded 12.5 basis points is too small (3 observations since September 2011) and decreasing the threshold below 12.5 basis points will defeat the purpose of conditioning on a meaningful amount of skew.

Bauer, M. D., and G. D. Rudebusch (2016): "Monetary policy expectations at the zero lower bound," *Journal of Money, Credit and Banking*, 48(7), 1439–1465.

Bekaert, G., E. Engstrom, and A. Ermolov (2021a): "Macro risks and the term structure of interest rates," *Journal of Financial Economics*, 141(2), 479–504.

——— (2021b): "Uncertainty and the Economy: The Evolving Distributions of Aggregate Supply and Demand Shocks," *Available at SSRN 3765164*.

Benaglia, T., D. Chauveau, D. R. Hunter, and D. S. Young (2010): "mixtools: an R package for analyzing mixture models," *Journal of statistical software*, 32, 1–29.

Berger, D., I. Dew-Becker, and S. Giglio (2020): "Uncertainty shocks as second-moment news shocks," *The Review of Economic Studies*, 87(1), 40–76.

Brodsky, B., M. Del Negro, J. Fiorica, E. LeSueur, A. Morse, A. P. Rodrigues, et al. (2016): "Reconciling survey-and market-based expectations for the policy rate," Discussion paper, Federal Reserve Bank of New York.

Cieslak, A. (2018): "Short-rate expectations and unexpected returns in treasury bonds," *The Review of Financial Studies*, 31(9), 3265–3306.

Correia-Golay, E., S. Friedman, and M. McMorrow (2013): "Understanding the New York Fed's Survey of Primary Dealers," *Current Issues in Economics and Finance*, 19(6).

Crump, R., E. Moench, W. O'Boyle, M. Raskin, C. Rosa, and L. Stowe (2014): "Interest rate derivatives and monetary policy expectations," *Federal Reserve Bank of New York Liberty Street Economics (blog), December*, 5, 2014.

D'Amico, S., and A. Orphanides (2014): "Inflation uncertainty and disagreement in bond risk premia," .

De Polis, A., L. Melosi, and I. Petrella (2024): "The taming of the skew: Asymmetric inflation risk and monetary policy," Discussion paper, University of Warwick, Department of Economics.

Diebold, F. X., and R. S. Mariano (1995): "Comparing predictive accuracy," *Journal of Business and Economic Statistics*, 13(3), 253–263.

Engelberg, J., C. F. Manski, and J. Williams (2009): "Comparing the point predictions and subjective probability distributions of professional forecasters," *Journal of Business & Economic Statistics*, 27(1), 30–41.

Ferreira, T. R. (2024): "Cross-sectional financial conditions, business cycles and the lending channel," *Journal of Monetary Economics*, p. 103597.

Figlewski, S. (2018): "Risk-Neutral Densities: A Review," *Annual Review of Financial Economics*, 10, 329–359.

Ghysels, E., A. Plazzi, and R. Valkanov (2016): "Why invest in emerging markets? The role of conditional return asymmetry," *The Journal of Finance*, 71(5), 2145–2192.

Giordani, P., and P. Söderlind (2003): "Inflation forecast uncertainty," *European Economic Review*, 47(6), 1037–1059.

Gneiting, T. (2011): "Making and evaluating point forecasts," *Journal of the American Statistical Association*, 106(494), 746–762.

Groeneveld, R. A., and G. Meeden (1984): "Measuring skewness and kurtosis," *Journal of the Royal Statistical Society: Series D (The Statistician)*, 33(4), 391–399.

Harvey, D., S. Leybourne, and P. Newbold (1997): "Testing the equality of prediction mean squared errors," *International Journal of forecasting*, 13(2), 281–291.

Huang, R., K. Pilbeam, and W. Pouliot (2022): "Are macroeconomic forecasters optimists or pessimists? A reassessment of survey based forecasts," *Journal of Economic Behavior and Organization*, 197, 706–724.

Jackwerth, J. (2020): "What do index options teach us about COVID-19?," *The Review of Asset Pricing Studies*, 10(4), 618–634.

Jensen, H., I. Petrella, S. H. Ravn, and E. Santoro (2020): "Leverage and deepening business-cycle skewness," *American Economic Journal: Macroeconomics*, 12(1), 245–81.

Joergensen, K., and A. Meldrum (2019): "Expectations about the Federal Funds Rate in the Long Run," *FEDS Notes*, (2019-10), 09.

Kim, D. H. (2009): "Zero bound, option-implied pdfs, and term structure models," .

Kim, D. H., and K. J. Singleton (2012): "Term structure models and the zero bound: an empirical investigation of Japanese yields," *Journal of Econometrics*, 170(1), 32–49.

Kim, T.-H., and H. White (2004): "On more robust estimation of skewness and kurtosis," *Finance Research Letters*, 1(1), 56–73.

Kwong, M., D. Leung, A. Wong, and J. Zhang (2020): "Interpreting survey-based federal funds rate forecasts: how accurate are they in reflecting market expectations?," *International Journal of Monetary Economics and Finance*, 13(1), 68–88.

Laster, D., P. Bennett, and I. S. Geoum (1999): "Rational bias in macroeconomic forecasts," *The Quarterly Journal of Economics*, 114(1), 293–318.

Melick, W. R., and C. P. Thomas (1997): "Recovering an asset's implied PDF from option prices: an application to crude oil during the Gulf crisis," *Journal of Financial and Quantitative Analysis*, 32(1), 91–115.

Mertens, T. M., and J. C. Williams (2021): "What to expect from the lower bound on interest rates: Evidence from derivatives prices," *American Economic Review*, 111(8), 2473–2505.

Nagel, S., and Z. Xu (2022): "Dynamics of subjective risk premia," Discussion paper, National Bureau of Economic Research.

Nakata, T. (2017): "Uncertainty at the zero lower bound," *American Economic Journal: Macroeconomics*, 9(3), 186–221.

Nakata, T., and H. Tanaka (2016): "Equilibrium yield curves and the interest rate lower bound," *FEDS Working Paper*.

Piazzesi, M., J. Salomao, and M. Schneider (2015): "Trend and Cycle in Bond Premia," .

Potter, S., M. Del Negro, G. Topa, and W. Van der Klaauw (2017): "The advantages of probabilistic survey questions," *Review of Economic Analysis*, 9(1), 1–32.

Prates, M. O., V. H. Lachos, and C. R. B. Cabral (2013): "mixsmsn: Fitting finite mixture of scale mixture of skew-normal distributions," *Journal of Statistical Software*, 54, 1–20.

Reifschneider, D., and J. C. Williams (2000): "Three lessons for monetary policy in a low-inflation era," *Journal of Money, Credit and Banking*, pp. 936–966.

Schmeling, M., A. Schrimpf, and S. Steffensen (2020): "Monetary Policy Expectation Errors," *Available at SSRN 3553496*.

Söderlind, P., and L. Svensson (1997): "New techniques to extract market expectations from financial instruments," *Journal of Monetary Economics*, 40(2), 383–429.

Svensson, L. E. (2001): "Price stability as a target for monetary policy: defining and maintaining price stability," in *The Monetary Transmission Process*, pp. 60–111. Springer.

van Zwet, W. (1964): "Convex transformations: A new approach to slcewness and kurtosis," *Statistica Neerlandica*, 18(4), 433–441.

Wright, J. H. (2017): "Forward-looking estimates of interest-rate distributions," *Annual Review of Financial Economics*, 9, 333–351.

Appendix

A Details on the Data Construction

Aggregate results from the SPD are publicly available from January 2011 (The SMP was launched in January 2014 and aggregate results are publicly available from then). While the aggregate distribution of the fed funds rate is a regular part of the surveys, it is not always available, and the question format can differ across surveys. The data are only available in separate Acrobat pdf files for each survey, and must be manually extracted for statistical analysis.

The majority of the surveys ask for unconditional distributions of the fed funds rate. An example of the question and average response made public by FRBNY is shown in the top panel of Figure A.1. From January 2011 through September 2012, the SPD includes only the 12 months ahead distribution. From April 2012, the SPD began asking for the end of 2014 distribution as well.

After suspending the question about the distribution for the October and December 2012 surveys, the SPD resumed the question from January 2013, but shifted the maturity of the distribution to year-end dates (up to three year-ends ahead). This means we need to interpolate across year-end distributions to obtain an estimate of the one-year ahead distribution. For December 2014 and December 2015, the year-end distribution starts from year-end 2015 and year-end 2016, respectively, not providing us the starting point of the interpolation. For these two dates, we simply assume that the closest year-end distribution reported is equivalent to the one-year ahead distribution.

We did not calculate the means from June 2013 through December 2013, as either the question about the fed funds rate distribution was not asked or the maturities started from end-2014, which prevented us from interpolating a one-year ahead estimate.

From January 2016 to October 2019, the SPD began asking various conditional distributions based on certain economic scenarios. The aggregate conditional distributions are disclosed to the public, but not the unconditional pdf-implied means. An example of the questions and responses made public by FRBNY that we use to construct the unconditional distribution is shown in the bottom panel of Figure A.1.

The October/November and December surveys from 2016 to 2018 only contain (conditional) distributions for the next year-end and not the current year-end. To avoid substantial interpolation error, we simply treat the one-year ahead distribution for the October/November data during this period as N/A, but use the next year-end distribution as a proxy for the one-year ahead distribution for the December surveys.

In order to compute moments from the discrete distributions, we assign the probabilities on each bin to the mid-point of its range. For the open-ended terminal bins we must make further assumptions. For the lower terminal bin, we assume it represents a value that is equal to one-half a bin width (of the adjacent bin) less than the end point of the lowest bin. For instance, if the lowest bin represents rates less than 0 percent and the adjacent bin represents rates between 0 and 0.25 percent, we assume the lower terminal bin represents a probability of the fed funds rate being at -0.125 percent. Where the bin width is not the same for all the bins, the tail bins assume a bin width of the adjacent bin. For the higher terminal bin, we distribute the probability of the tail bin across the federal funds rate beyond the upper bound using simple exponential decay. These assumptions about the tail bins do not materially affect our results except for a few surveys in 2022, where there has been significant probability mass at the higher terminal bin (Applying the same assumption for the lower terminal bin to the higher terminal bin tends to result in larger negative skew).

Unconditional Distribution (January 2022)

3d) Please indicate the percent chance that you attach to the target federal funds rate or range falling in each of the following ranges at the end of 2022, 2023, and 2024. If you expect a target range, please use the midpoint of that range in providing your response.
(21 responses)

Federal Funds Rate or Range at the End of 2022										
	< 0.00%	0.00 - 0.25%	0.26 - 0.50%	0.51 - 0.75%	0.76 - 1.00%	1.01 - 1.25%	1.26 - 1.50%	1.51 - 1.75%	1.76 - 2.00%	≥ 2.01%
Average	1%	1%	4%	13%	23%	33%	15%	7%	3%	1%

Federal Funds Rate or Range at the End of 2023										
	< 0.00%	0.00 - 0.25%	0.26 - 0.50%	0.51 - 0.75%	0.76 - 1.00%	1.01 - 1.25%	1.26 - 1.50%	1.51 - 1.75%	1.76 - 2.00%	≥ 2.01%
Average	1%	2%	1%	2%	3%	9%	11%	17%	20%	33%

Federal Funds Rate or Range at the End of 2024										
	≤ 0.75%	0.76 - 1.00%	1.01 - 1.25%	1.26 - 1.50%	1.51 - 1.75%	1.76 - 2.00%	2.01 - 2.25%	2.26 - 2.50%	2.51 - 2.75%	≥ 2.76%
Average	7%	3%	3%	4%	8%	16%	17%	15%	13%	13%

Conditional Distribution (January 2017)

3c) Please indicate the percent chance that you attach to the following possible outcomes for the Committee's next policy action in 2017.

	Next Change is Increase in Target Rate or Range	Next Change is Decrease in Target Rate or Range	No Change in Target Rate or Range in 2017
Average	85%	5%	10%

3d) Conditional on the Committee's next policy action in 2017 being an increase in the target federal funds rate or range, please indicate the percent chance that you attach to the following possible outcomes for the timing of such a change. Only fill out this conditional probability distribution if you assigned a non-zero probability to the Committee's next policy action in 2017 being an increase.

	Increase Occurs at January FOMC meeting	Increase Occurs at March FOMC meeting	Increase Occurs at May FOMC meeting or later
Average	3%	20%	78%

3e) Please indicate the percent chance that you attach to the target federal funds rate or range falling in each of the following ranges at the end of 2017, conditional on the following possible scenarios for the direction and timing of the Committee's next policy action in 2017. Only fill out the conditional probability distributions for which you assigned a non-zero probability to the conditioning event occurring. If you expect a target range, please use the midpoint of that range in providing your response.

Next change is an increase, occurs at Mar. meeting or earlier								
	≤ 0.50%	0.51 - 0.75%	0.76 - 1.00%	1.01 - 1.25%	1.26 - 1.50%	1.51 - 1.75%	1.76 - 2.00%	≥ 2.01%
Average	2%	8%	10%	27%	33%	16%	4%	1%

Next change is an increase, occurs at May meeting or later								
	≤ 0.50%	0.51 - 0.75%	0.76 - 1.00%	1.01 - 1.25%	1.26 - 1.50%	1.51 - 1.75%	1.76 - 2.00%	≥ 2.01%
Average	2%	4%	17%	41%	26%	9%	2%	0%

Next change is a decrease								
	< 0.0%	0.00 - 0.25%	0.26 - 0.50%	0.51 - 0.75%	0.76 - 1.00%	1.01 - 1.25%	1.26 - 1.50%	≥ 1.51%
Average	10%	53%	34%	3%	0%	0%	0%	0%

3f-ii) Please indicate the percent chance that you attach to moving to the ZLB at some point between now and the end of 2019.

	Probability of Moving to ZLB at Some Point between now and the end of 2019
25th Pctl	15%
Median	20%
75th Pctl	25%

3f-iii) Please indicate the percent chance that you attach to the target federal funds rate or range falling in each of the following ranges at the end of 2018 and 2019, conditional on moving to the ZLB at some point between now and the end of 2019. Only fill out these conditional probability distributions if you assigned a non-zero probability to moving to the ZLB at some point between now and the end of 2019. If you expect a target range, please use the midpoint of that range in providing your response.
(22 responses)

Year-end 2018								
	< 0.00%	0.00- 0.25%	0.26- 0.50%	0.51- 1.00%	1.01- 1.50%	1.51- 2.00%	2.01- 2.50%	≥ 2.51%
Average	6%	37%	16%	14%	11%	11%	4%	2%

Year-end 2019								
	< 0.00%	0.00- 0.25%	0.26- 0.50%	0.51- 1.00%	1.01- 1.50%	1.51- 2.00%	2.01- 2.50%	≥ 2.51%
Average	7%	47%	20%	12%	7%	4%	2%	1%

Figure A.1: **Examples of Survey Questions and Responses in the SPD**

B Comparison of the SPD and the SMP Modal Series

Figure B.1 compares the median of each respondent's modal forecast of the approximately one-year ahead fed funds rate from the SPD and the SMP since the SMP was launched in January 2014.[42] We find that although the two series diverge at times, we do not observe a systematic pattern, and they generally track each other very closely. The average difference between the two series is less than a basis point (with a standard deviation of 13 basis points), and the correlation of the two series is over 0.99.

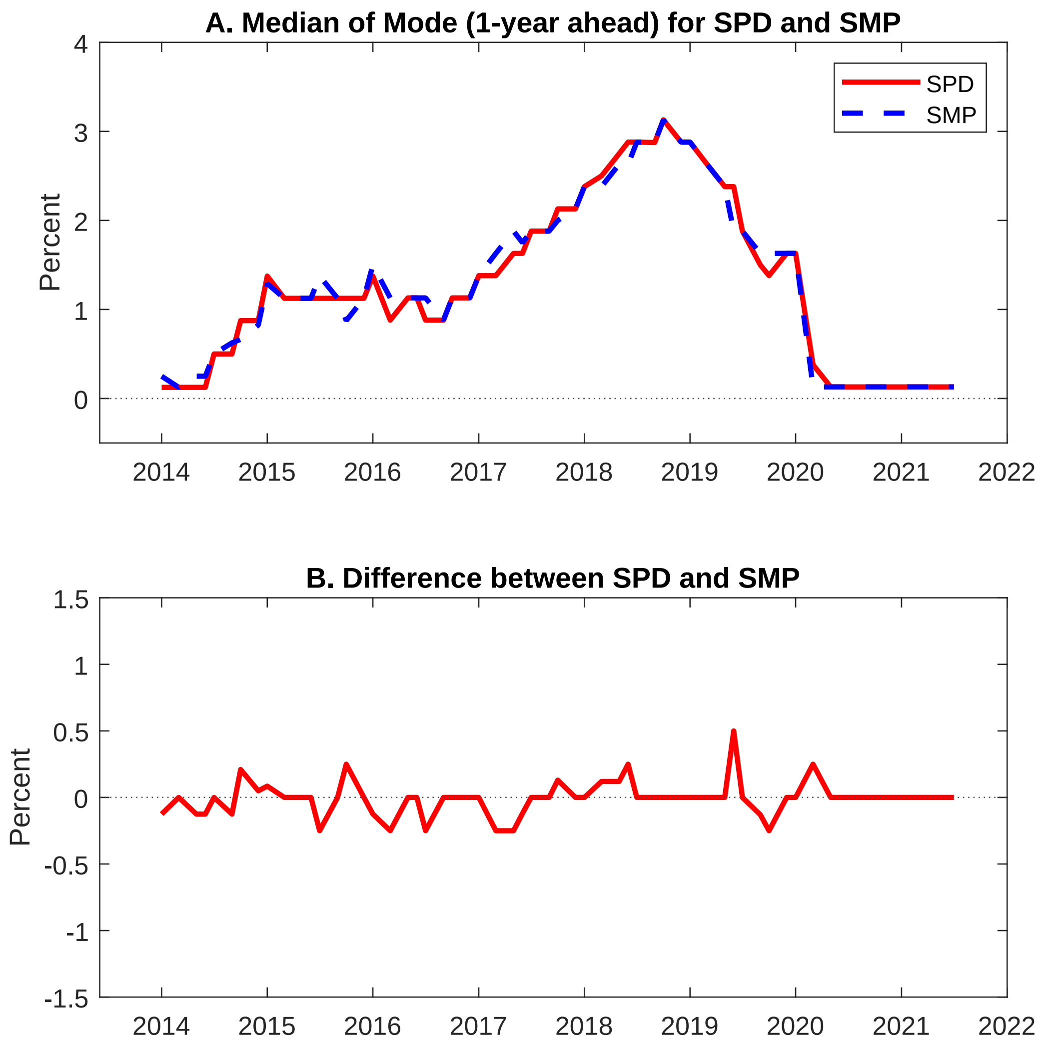

Figure B.1: **Comparison of the SPD and the SMP Modes**

Note: Panel A shows the median of each respondent's modal forecast of the approximately one-year ahead fed funds rate from the SPD and the SMP. Panel B shows the difference between the two series.

[42] The goal here is to compare the SPD and the SMP as precisely as possible and not necessarily to construct a constant maturity measure. Hence, for each survey, we take the forecast for the end of the quarter in which the date one-year from the survey date falls. This allows us to compare forecasts that are approximately one-year ahead without interpolation.

C Skewness based on Raw Survey Series

Figure C.1 shows skewness based on the raw survey series, which is the unsmoothed version of the series shown in Figure 1.

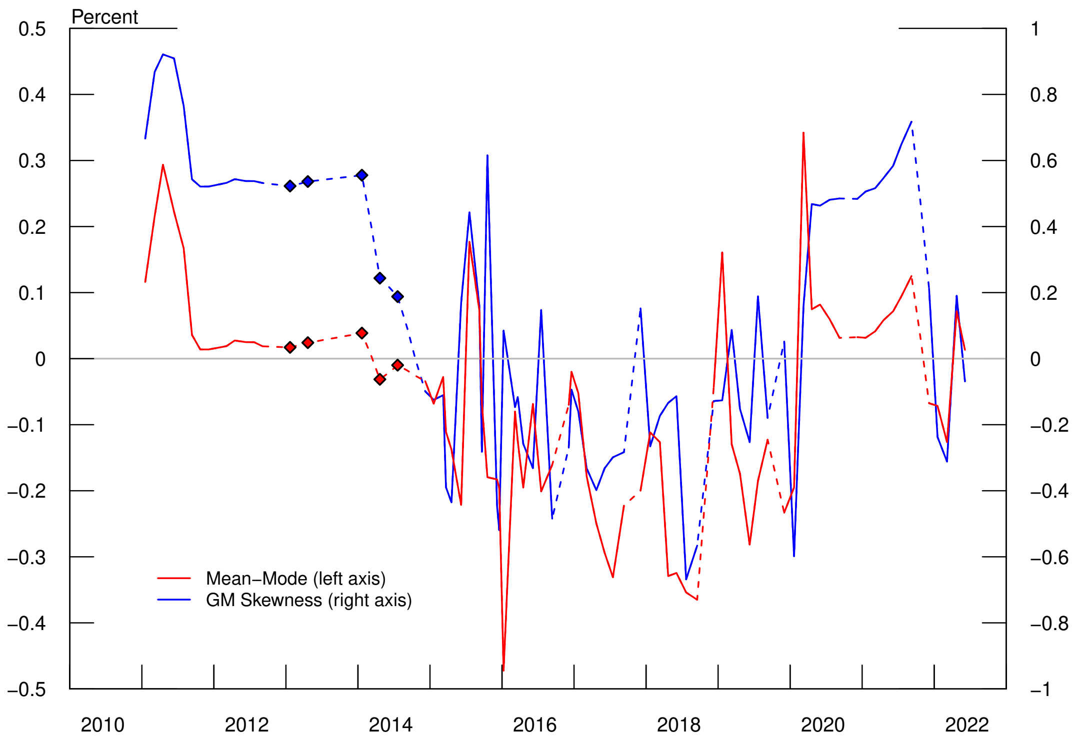

Figure C.1: **Unsmoothed Skewness of the Aggregate Federal Funds Rate Distribution from the SPD**

Note: Whenever data from an adjacent survey is unavailable, we take the average using data up to two surveys apart. Dashed line segments are interpolations between survey dates where data is available.

D Standard Deviation of the Aggregate Fed Funds Rate Distribution

Figure D.1 shows the time series of the standard deviation of the aggregate fed funds rate distribution. As described in the main text, the standard deviation collapsed in mid-2011 when the Fed introduced date-based forward guidance, contributing to a large increase in skewness at the time. The standard deviation collapsed once again at the onset of the pandemic crisis in March 2020 and has retraced some of the decline most recently.

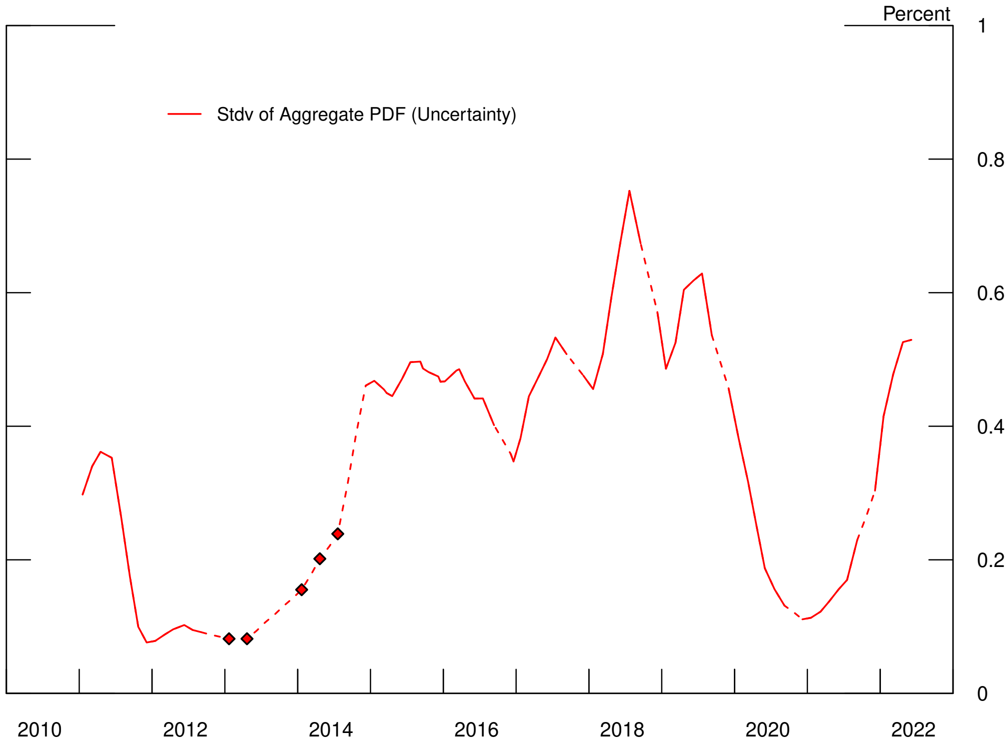

Figure D.1: **Standard Deviation of the Aggregate Fed Funds Rate Distribution**

Note: Whenever data from an adjacent survey is unavailable, we take the average using data up to two surveys apart. Dashed line segments are interpolations between survey dates where data is available. The diamonds indicate points that have no adjoining observations.

E Alternative Skewness Measures

There are various ways to compute measures of skewness and it is known that conventional measures are not necessarily robust (Kim and White (2004)). Hence, in this Appendix, we compute the skewness of the aggregate distribution of the fed funds rate using several measures as a robustness check. In particular, in addition to the two measures we presented in the main text, we compute three other measures of skewness that are well known: 1) Skewness using third moments: $E\left[\left(\frac{X-E[X]}{\sigma}\right)^3\right]$, 2) Normalized skewness using modes: $\frac{E[X]-E_{mode}[X]}{\sigma}$ and 3) Normalized skewness using medians: $\frac{E[X]-E_{median}[X]}{\sigma}$ (the latter two measures sometimes referred to as "Pearson skewness"), where $E[X]$, $E_{mode}[X]$, $E_{median}[X]$ and σ are the means, modes, medians and standard deviations of X. Figure E.1 plots the skewness measures. We find that the broad pattern of time-varying skew that we emphasized in the main text holds regardless of how we measure the skew—skewness was generally significantly positive while the fed funds rate was at the ELB, and turned negative after liftoff before moving back to positive territory following monetary policy accommodation towards the end of 2019 and 2020. However, evidence of skew before the financial crisis seems rather noisy and there appear to be significant quantitative differences among the measures particularly when the fed funds rate was close to the ELB. For instance, skewness measured by third moments rose rapidly after the volatility of the near-term fed funds rate collapsed following the announcement of date-based forward guidance in mid-2011, as well as after the sharp rate cut following the pandemic. This measure seems especially sensitive to volatility as the denominator raises volatility to its third power.

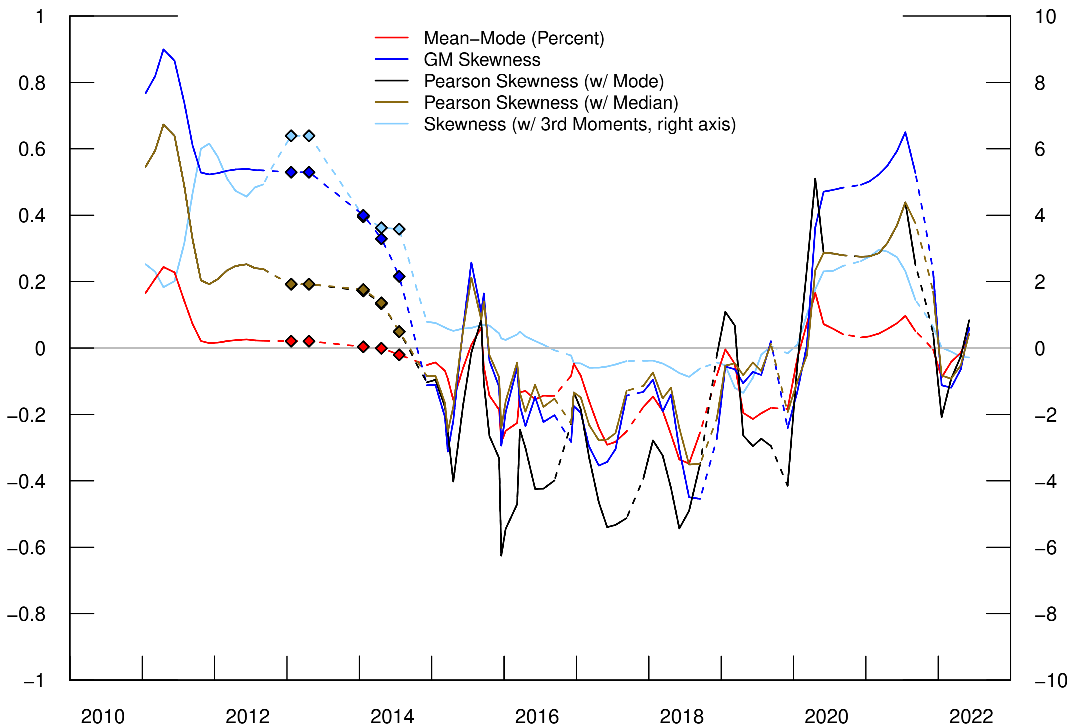

Figure E.1: **Alternative Skewness Measures of the Aggregate Fed Funds Rate Distribution**

Note: All series are smoothed by taking the centered moving average of three adjacent surveys. Whenever data from an adjacent survey is unavailable, we take the average using data up to two surveys apart. Dashed line segments are interpolations between survey dates where data is available. The diamonds indicate points that have no adjoining observations.

F Skewness from Fitted Distributions

In the main text, we showed skewness measures based on a distribution where we simply assume that the survey (average) probabilities are assigned to the mid-point of each bin. In this section, we show skewness measures based on alternative distributions where we fit the histograms with both parametric and nonparametric continuous distributions. The fitted distributions provide further robustness checks for the skewness measures.

We show results from two distinct fitting methods. One method is to fit the histogram parametrically using a mixture of two skew-normal distributions. A mixture distribution is crucial in capturing bimodality of the histogram observed in part of the sample. In addition, we choose the skew-normal distribution as a component of the mixture to account for skewness.[43] We apply the method described in Prates, Lachos, and Cabral (2013).

[43] We also fitted a mixture of two skew-t distribution (not shown) and found the results were broadly similar, with a somewhat worse fit.

The other method is to fit the histogram nonparametrically. We apply the method described in Benaglia, Chauveau, Hunter, and Young (2010) using a standard normal kernel and a choice of bandwidth according to the Silverman's rule. When implementing both methods, we first generate bootstrap data of the histogram assuming that draws within each bin are uniformly distributed across the width of the bin (as opposed to assuming the draws are concentrated on the mid-point of each bin).[44]

Figure F.1 shows examples of the histograms of the survey distributions along with the fitted distributions for four survey editions. The four editions are selected to capture the broad time-series pattern of skewness over the sample: 1) the positive skew following the financial crisis and the Federal Reserve's forward guidance at the ELB (April 2013, top left), 2) the negative skew during the tightening cycle since December 2015 (September 2018, top right), 3) the positive skew following the pandemic crisis and the return to the ELB (October 2020, bottom left) and 4) the negative skew observed more recently as the Federal Reserve tightened monetary policy (March 2022, bottom right). Note the forecast horizon is chosen to be as close as possible to one-year ahead, but not exactly, since the respondents typically report year-end forecasts and we only interpolate moments from the distributions.[45] The results show that both the parametric and nonparametric distributions capture the key features of the histograms (both positive and negative skew, as well as occasional bimodality) in each period reasonably well.

[44]Assuming the draws are concentrated on the mid-point effectively assumes that the true distribution assigns zero probability on the range of a bin except for the mid-point, which results in a nonsensical fit.

[45]The forecast horizon for each panel is end-2013 (for April 2013), end-2019 (for September 2018), end-2021 (for October 2020), and end-2022 (for March 2022).

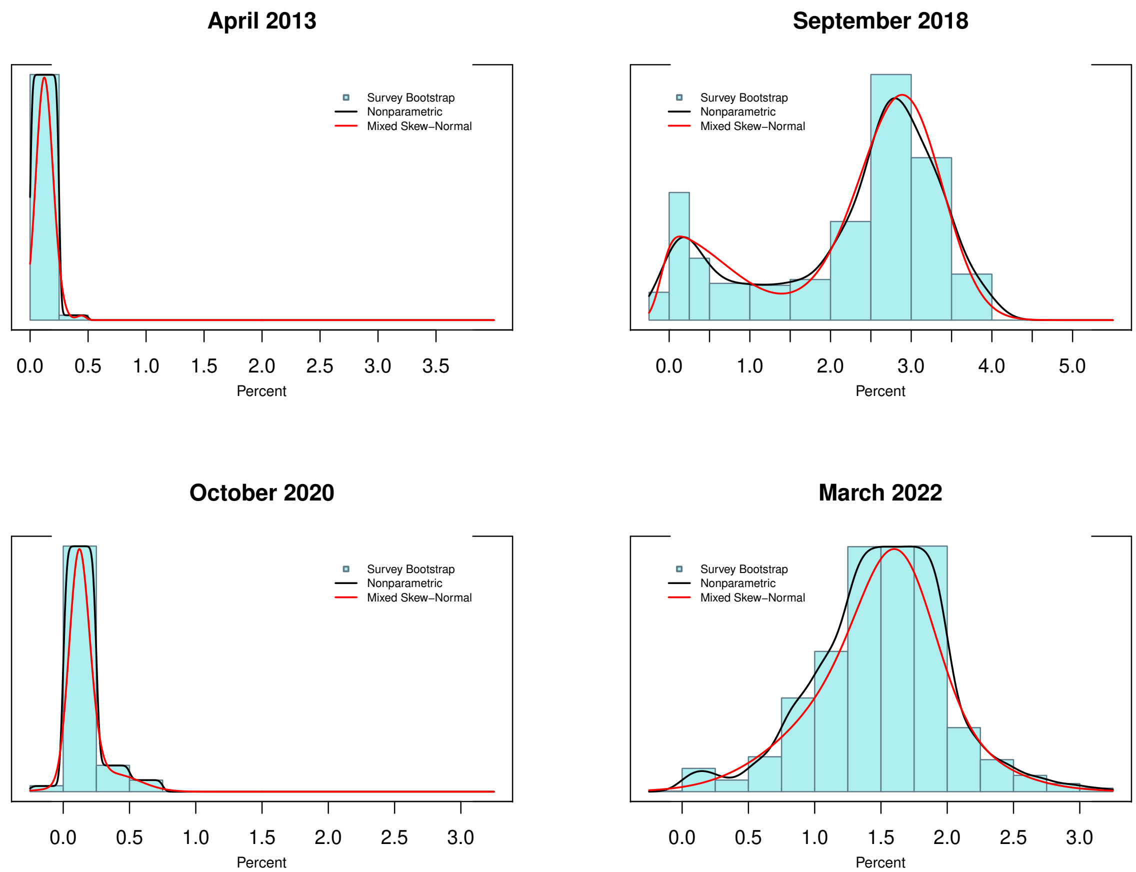

Figure F.1: **Fitted Aggregate Fed Funds Rate Forecast Distributions**

Note: The mixed skew-normal distribution is scaled to fit the panel for April 2013, October 2020 and March 2022. The forecast horizon is approximately one-year ahead.

Figure F.2 shows the skewness measures based on the histograms of the survey distributions along with those based on the fitted distributions. We find that for both the GM skewness and the skewness based on the standardized third moment, the time-series of skew follow a similar pattern across distributions that are fitted using different methods.

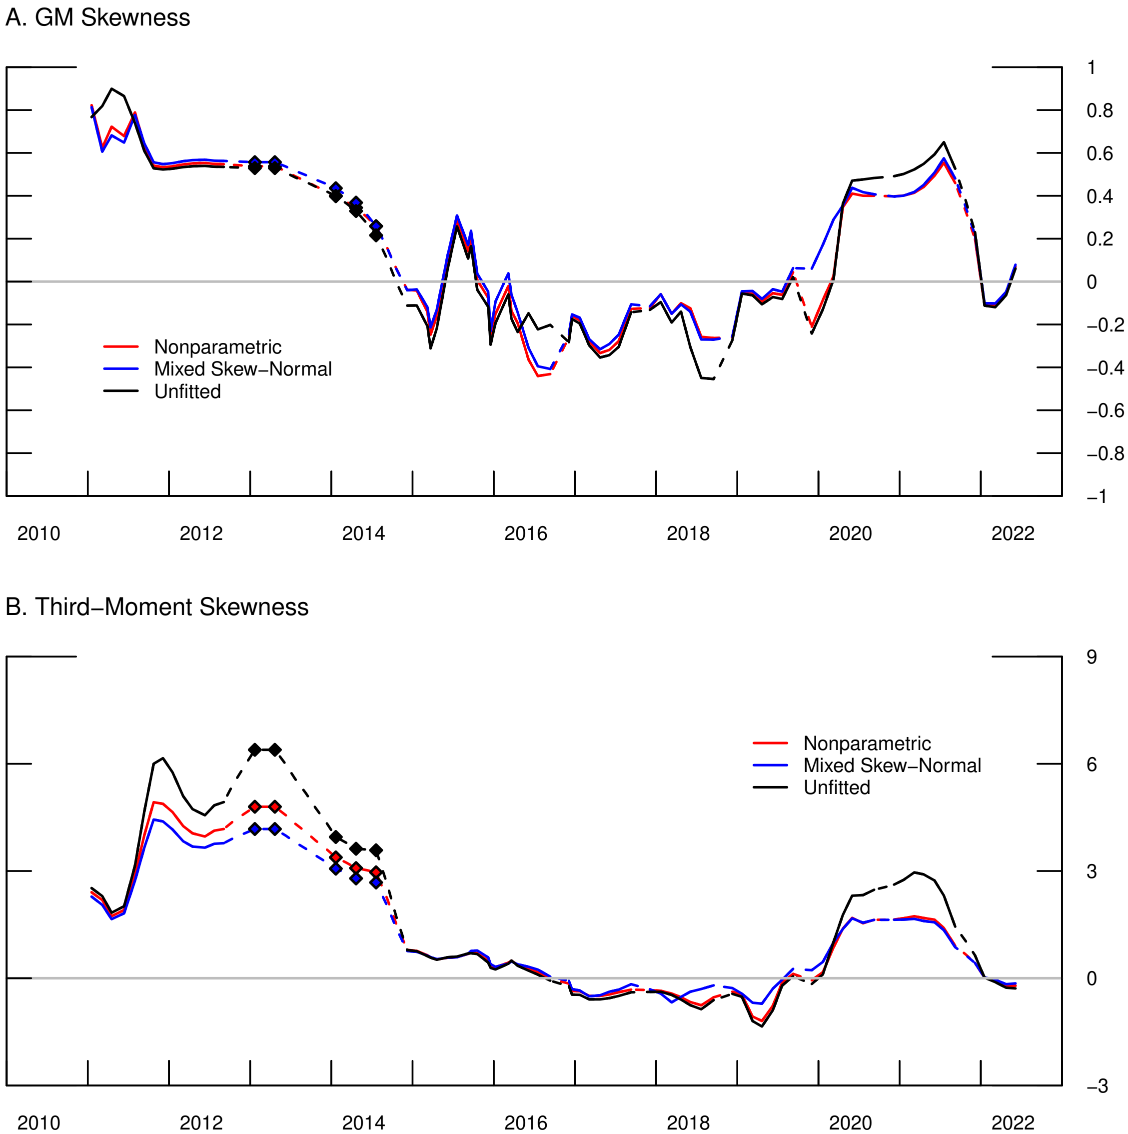

Figure F.2: **Skewness of the Fitted Aggregate Fed Funds Rate Distribution**

Note: All series are smoothed by taking the centered moving average of three adjacent surveys. Whenever data from an adjacent survey is unavailable, we take the average using data up to two surveys apart. Dashed line segments are interpolations between survey dates where data is available. The diamonds indicate points that have no adjoining observations.

G Eurodollar Skewness

Figure G.1 shows several one-year ahead skewness series computed from distributions implied by Eurodollar futures options. In terms of the construction of these particular measures, there are a number of differences compared to the approach used in the main text. First, the underlying asset for these securities is the 3-month London Interbank Offered Rate (LIBOR), which is in contrast to our measures based on the expected federal funds rate. Second, the associated distributions that are extracted are risk-neutral and may contain liquidity and risk premiums, which is in contrast to our physical measures derived from the surveys. Nonetheless, this measure is available on a higher frequency basis and can act as a robustness check.

To extract the probability distribution function of the 3-month LIBOR, we use a flexible parametric form with a mixture of distributions and conduct a nonlinear optimization to obtain the parameters that minimize the average distance between the pdf-implied option prices and the observed options prices at various option strikes. To account for the lower bound, we sum up the weights below zero. This particular approach is just one of many that are available in the literature: see Figlewski (2018) for an extensive review of risk-neutral densities. We use the parametric mixture because alternative methods such as the Breeden-Litzenberger method can lead to somewhat noisy results, while the "Shimko method" based on implied volatility can have difficulty accommodating unusual shapes of the pdf. Previous studies such as Melick and Thomas (1997), Söderlind and Svensson (1997) and more recent studies such as Jackwerth (2020) use parametric mixtures similar to our approach.

The broad pattern of skew resembles the skewness series based on the SPD (Figure 1), but the option-implied skew is also different in several ways. In particular, the option-implied skew is somewhat volatile from 2012 through 2014, and the period of negative skewness is shorter and its magnitude is smaller.

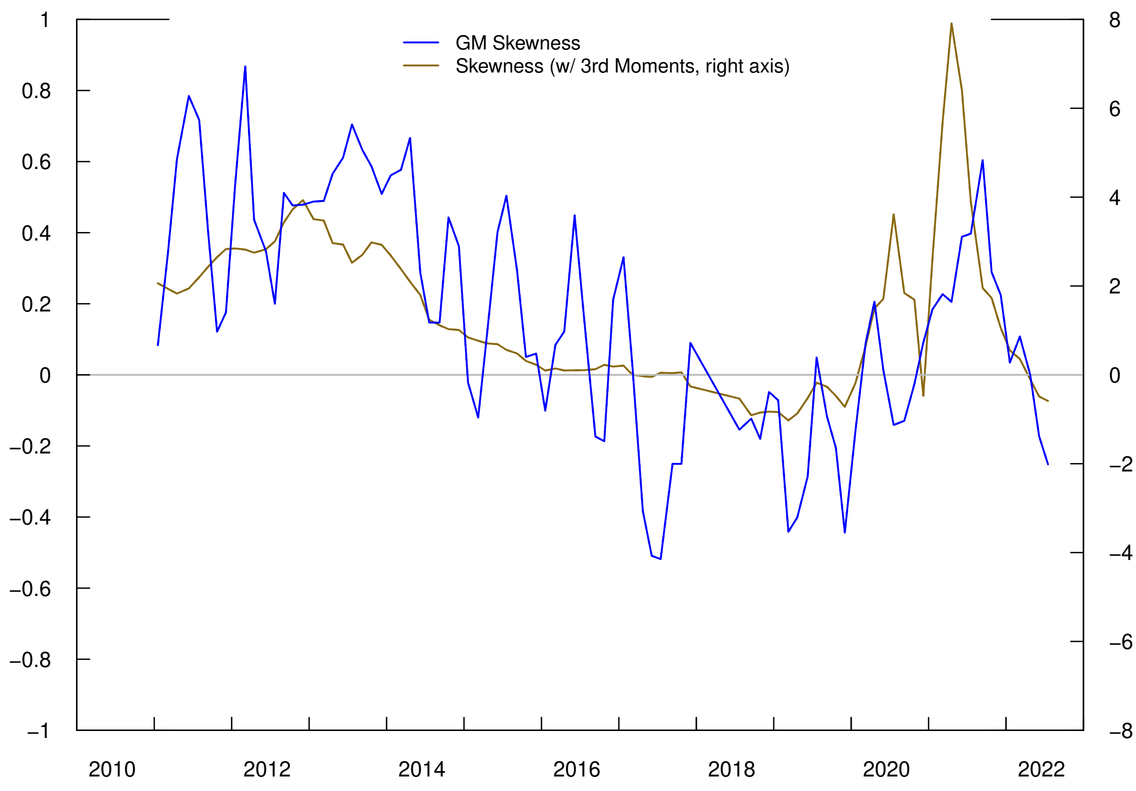

Figure G.1: **Eurodollar Skewness**

Note: All series are smoothed by taking the centered moving average of three adjacent surveys.

H Two-year Ahead Skewness

Figure H.1 compares the one-year ahead skewness series with the two-year ahead skewness series, based on the GM measure. The two-year skew generally correlates well with the one-year series. Figure H.2 plots alternative skewness measures for the two-year ahead maturity, analogous to Figure E.1. As discussed in Appendix E, we find that the broad pattern of time-varying skew holds regardless of how we measure skew.

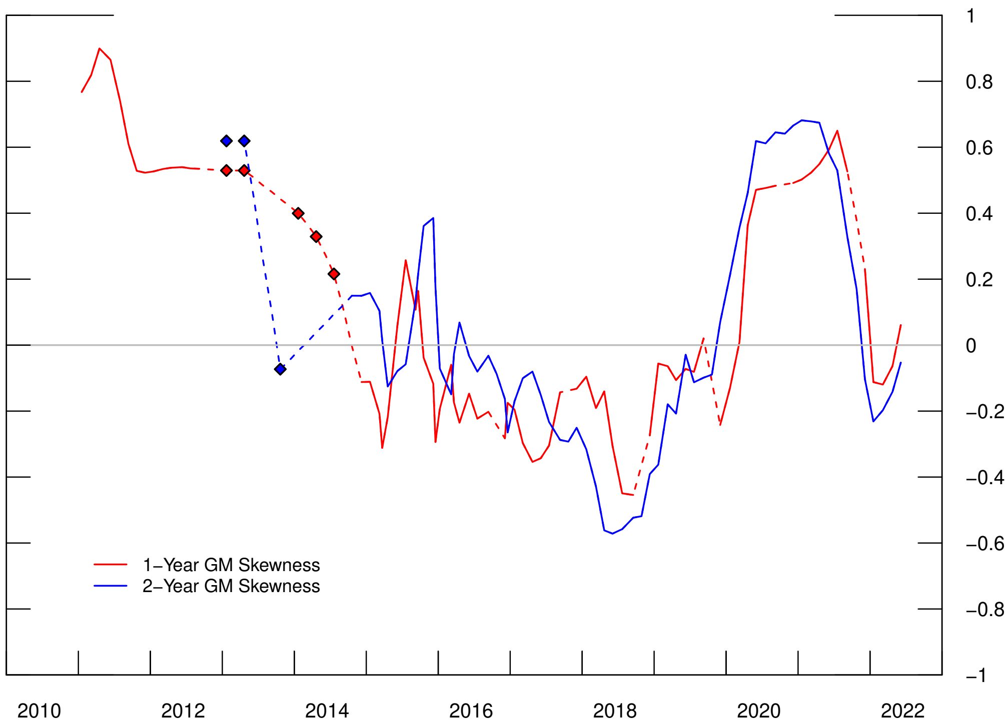

Figure H.1: **One-year Ahead v.s. Two-year Ahead Skewness of the Aggregate Fed Funds Rate Distribution**

Note: All series are smoothed by taking the centered moving average of three adjacent surveys. Whenever data from an adjacent survey is unavailable, we take the average using data up to two surveys apart. Dashed line segments are interpolations between survey dates where data is available. The diamonds indicate points that have no adjoining observations.

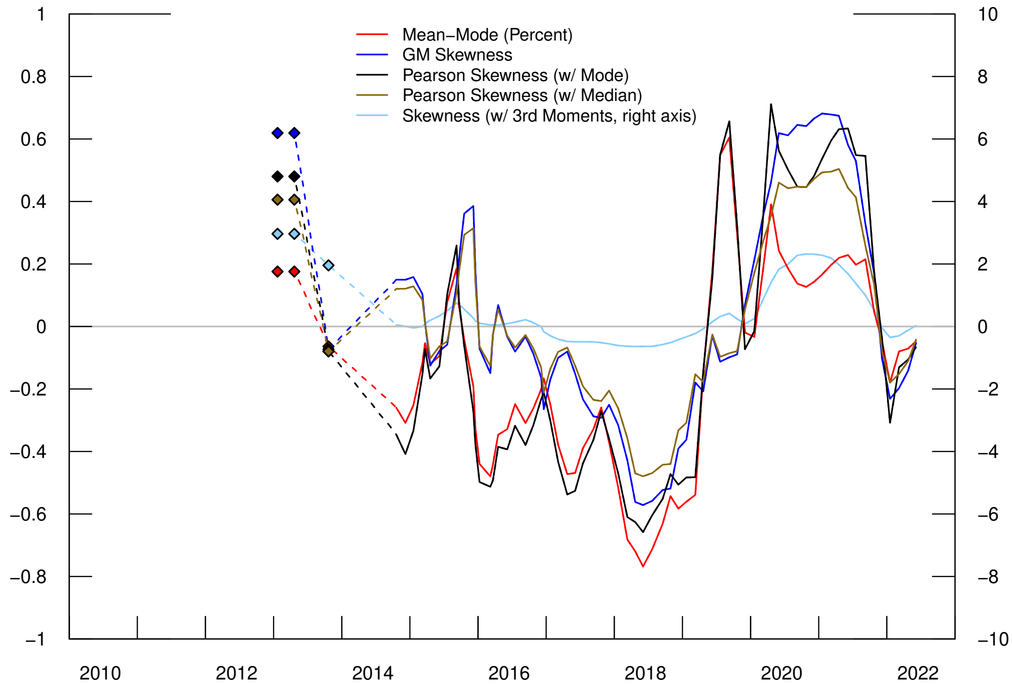

Figure H.2: **Alternative Skewness Measures of the Aggregate Fed Funds Rate Distribution (Two-Year Ahead)**

Note: All series are smoothed by taking the centered moving average of three adjacent surveys. Whenever data from an adjacent survey is unavailable, we take the average using data up to two surveys apart. Dashed line segments are interpolations between survey dates where data is available. The diamonds indicate points that have no adjoining observations.

I Further Details on the Macroeconomic Model and its Implications

I.1 Model Specification

In this section, we provide details of the New-Keynesian model used to simulate the aggregate fed funds rate distribution in section 3.2. The model is a stylized New-Keynesian model with an occasionally binding ELB constraint, fairly similar to the models used by Nakata (2017) or Nakata and Tanaka (2016), among others.

Representative Household: The representative household maximizes the value function:

$$V_t = \ln(C_t) + \chi \ln(1 - N_t) + \beta_t \mathbb{E}_t [V_{t+1}] \tag{I.1}$$

subject to its budget constraint:

$$P_t C_t + \mathbb{E}_t [M_{t+1} \mathcal{W}_{t+1}] = W_t N_t + \mathcal{W}_t + \Xi_t + T_t$$

where $\mathcal{W}_t$ is the consumer's wealth portfolio of state contingent claims (complete markets). Ξ_t is firms' profit rebated back to the consumer. β_t is a time-varying time discount factor, the process of which is described below.

The Euler equation for the nominal 1-period risk-free rate is:

$$\mathbb{E}_t[M_{t+1}R_t] = 1, \qquad M_{t+1} = \beta_t \frac{C_t}{C_{t+1}} \frac{1}{\Pi_{t+1}} \tag{I.2}$$

Intermediate Goods Producers: Monopolistically competitive intermediate goods producers $i \in [0, 1]$ maximize profits:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} M_{t-1,t} \left[P_t(i)Y_t(i) - W_t N_t(i) - \frac{\varphi}{2} \left(\frac{P_t(i)}{P_{t-1}(i)\bar{\Pi}} - 1 \right)^2 P_t Y_t \right]$$

with nominal rigidities based on Rotemberg adjustment costs. Producers are subject to the demand and production functions:

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\theta} Y_t$$

$$Y_t(i) = N_t(i) \tag{I.3}$$

Monetary Policy: Monetary policy is an interest rate rule with an occasionally binding ELB constraint:

$$R_t = \max\left[1, R_t^*\right]$$

where the shadow rate R_t^* follows:

$$R_t^* = \bar{R} \left[\frac{\Pi_t}{\bar{\Pi}} \right]^{\phi_\Pi}$$

Market Clearing/Exogenous Process: Aggregation and market clearing imply:

$$Y_t = N_t$$

$$Y_t = C_t + \frac{\varphi}{2} \left(\frac{\Pi_t}{\bar{\Pi}} - 1 \right)^2 Y_t$$

The exogenous discount rate process follows:

$$\ln \beta_t = (1 - \rho_\beta) \ln \bar{\beta} + \rho_\beta \ln \beta_{t-1} + \sigma_\beta \varepsilon_{\beta,t} \qquad \varepsilon_{\beta,t} \sim i.i.d\, \mathcal{N}(0, 1)$$

Calibration: The model is calibrated to quarterly data and is largely standard. Table I.1 summarizes the parameter values:

Parameter	Description	Parameter Value
β	Time discount rate at steady state	$\frac{1}{1.0015}$
χ	Preference over consumption vs leisure	0.25
θ	Elasticity of substitution among intermediate goods	6
φ	Price adjustment cost	55
$400(\bar{\Pi}-1)$	(Annualized) inflation target	2
ϕ_π	Coefficient on inflation in the Taylor rule	2.5
ρ_β	Persistence of discount factor shock	0.77
σ_β	Standard deviation of discount factor shock	$\frac{0.2174}{100}$

Table I.1: **Parameter Values**

The model is solved using a global solution method (time-iteration) to take account of non-linearities properly. To compute each forecast distribution in Figure 3, we simulate one million paths of the policy rate from the model.

I.2 Extension with Supply Shocks

In this section, we discuss details of the extended model with supply shocks. Supply shocks are introduced in a standard way as shocks to the total factor productivity (TFP). In other words, we simply augment the production function (I.3) with an exogenous process for TFP A_t:

$$Y_t(i) = A_t N_t(i)$$

where A_t follows:

$$\ln A_t = \rho_a \ln A_{t-1} + \sigma_a \varepsilon_{a,t} \qquad \varepsilon_{a,t} \sim i.i.d\,\mathcal{N}(0,1)$$

We first turn off demand shocks to clarify the role of supply shocks in our model. Figure I.1 plots the model's equilibrium decision rules, analogous to Figure 4 in the main text. The blue dashed lines in the two panels are the decision rules for the policy rate and inflation, respectively, when the model is not subject to a ELB constraint, while the solid lines are the corresponding decision rules when the model is subject to a ELB constraint. The decision rules shows a concave/convex shape with respect to TFP, similar to the decision rules with respect to the discount factor in Figure 4. The main difference (which is documented in various studies that use similar models) is the decision rule for output, which is generally increasing in TFP, while output decreases as the discount factor increases (not shown). Intuitively, higher TFP decreases the cost of production, leading to a deceleration of inflation. Consequently, monetary policy reacts by lowering the policy rate. However, in the presence of the ELB constraint, the central bank ends up accommodating more aggressively near the ELB to counter the significant downside risk to inflation caused by the limited effect of monetary policy at the ELB. Given the form of the decision rules and an analogous mechanism highlighted for the model with demand shocks, we conclude that the presence of the ELB constraint can generate both positive and negative skew endogenously from TFP shocks, which, by assumption, are drawn from a symmetric distribution.

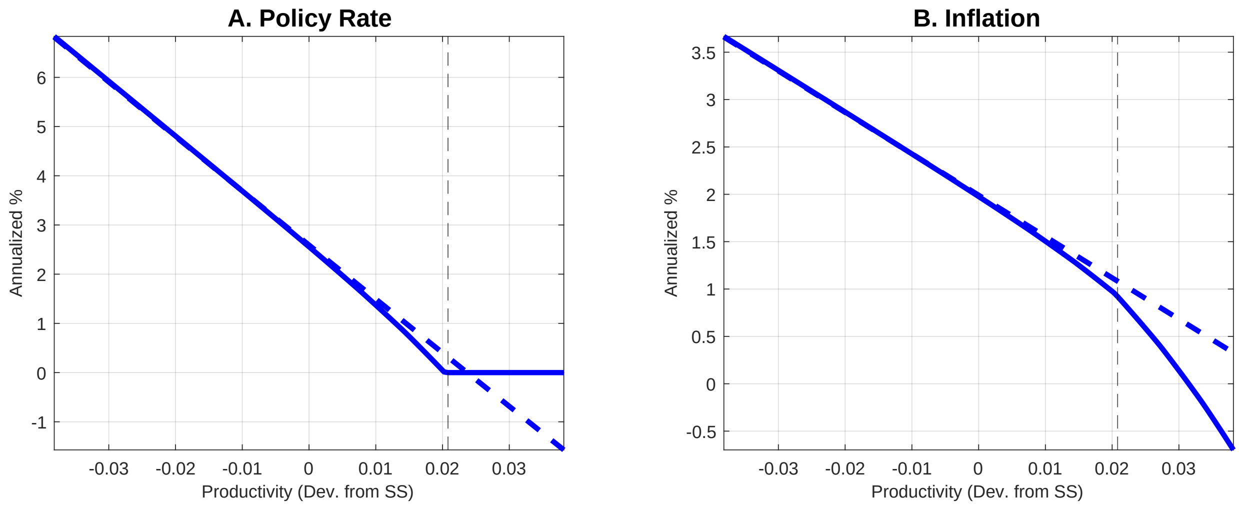

Figure I.1: **Equilibrium Decision Rules for the Policy Rate and Inflation**

Note: The vertical dashed lines show the value of TFP (its deviation from the steady state) when the ELB constraint binds. The solid blue lines and the dashed blue lines show decision rules for the model with the ELB constraint and without the ELB constraint, respectively.

We next study the implications of the extended model with both demand and supply shocks. The model parameters are fixed at those reported in Table I.1 except for the exogenous processes of the discount factor and TFP, which are specified as in Table I.2:

Parameter	Description	Parameter Value
ρ_β	Persistence of discount factor shock	0.68
σ_β	Standard deviation of discount factor shock	$\frac{0.2}{100}$
ρ_a	Persistence of TFP shock	0.75
σ_a	Standard deviation of TFP shock	$\frac{0.85}{100}$

Table I.2: **Parameter Values for the Extended Model**

While the parameter values are chosen so that the model is consistent with some basic aspects of the macroeconomy, the purpose of the model is to highlight the non-linear pattern of skew that is qualitatively consistent with the data. Figure I.2 plots the extended model's equilibrium decision rules, analogous to Figure 4 in the main text. Since the extended model includes TFP as an additional state variable, we plot decision rules as a function of the discount factor, and for different levels of TFP. As shown in the left panel, an increase (decrease) in TFP shifts the policy rate closer to (away from) the ELB for any level of the discount factor as long as the policy rate is not subject to the ELB constraint. This shift also changes the degree of local concavity around a specific discount factor, and hence skew. However, the skew can either increase or decrease, depending on both the discount factor and TFP.

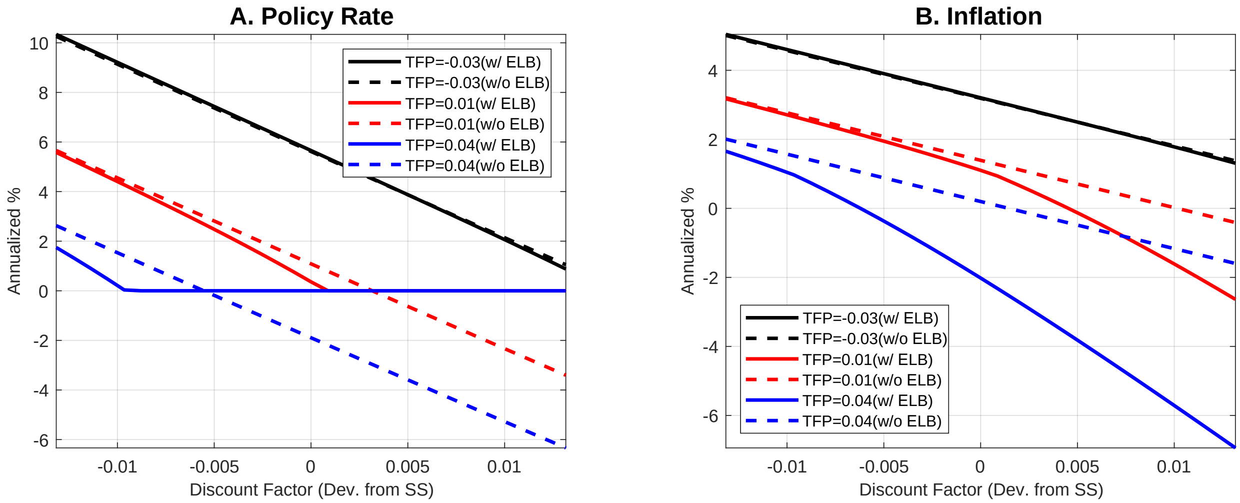

Figure I.2: **Equilibrium Decision Rules for the Policy Rate and Inflation**

Note: The solid lines and the dashed lines show decision rules for the model with the ELB constraint and without the ELB constraint, respectively.

To further understand the relation between the discount factor, TFP and the skewness of the policy rate forecast distribution, Figure I.3 plots the policy rate skew as a function of the discount factor and TFP.[46] The left panel plots the skew for the model with the ELB constraint. The colored contours indicate the level of skew, where white indicates zero skew, darker red indicates more negative skew, and darker blue indicates more positive skew. As implied by the decision rules for the policy rate in Figures 4 and I.1, the policy rate is decreasing in both the discount factor and TFP, hence the policy rate is highest at the bottom left corner of the panel. The dashed light blue line indicates the pair of TFP and the smallest discount factor that is associated with a zero policy rate (ELB threshold). Thus, the policy rate is zero to the right of the threshold and positive to the left.

Policy rate skew is generally positive to the right of the threshold where the policy rate is zero and expected to stay there within the forecast horizon. However, the magnitude of the skew is smaller for pairs of the discount factor and TFP that are close to the top right corner, as the forecast distribution collapses towards zero. Meanwhile, when the discount factor/TFP pair crosses the ELB threshold from the upper right and moves closer to the bottom left, skew decreases towards zero and eventually turns negative. Similar to the intuition gained from the model with either a demand or a supply shock, this reflects the possibility of a more aggressive rate cut before the policy rate reaches the ELB, due to the limited power of monetary policy at the ELB. For pairs of very low discount factor/TFP, the skew moves back towards zero, as the policy rate rises significantly above the ELB constraint and the decision rule becomes essentially linear. Since the policy rate is decreasing in both the discount factor and TFP, it follows that policy rate skew is not a monotonic function of the level of the policy rate, and exhibits an intuitive non-linearity.

The right panel is a similar plot of the policy rate skew for the model without the ELB constraint. In this case, the panel is simply blank (white), indicating skew is negligible across the entire state space. This is because the decision rule of the policy rate is essentially linear globally.

[46]We plot the Pearson skewness based on medians, since it is computationally less burdensome. Results are qualitatively similar using other skewness measures such as those based on modes or the third moment.

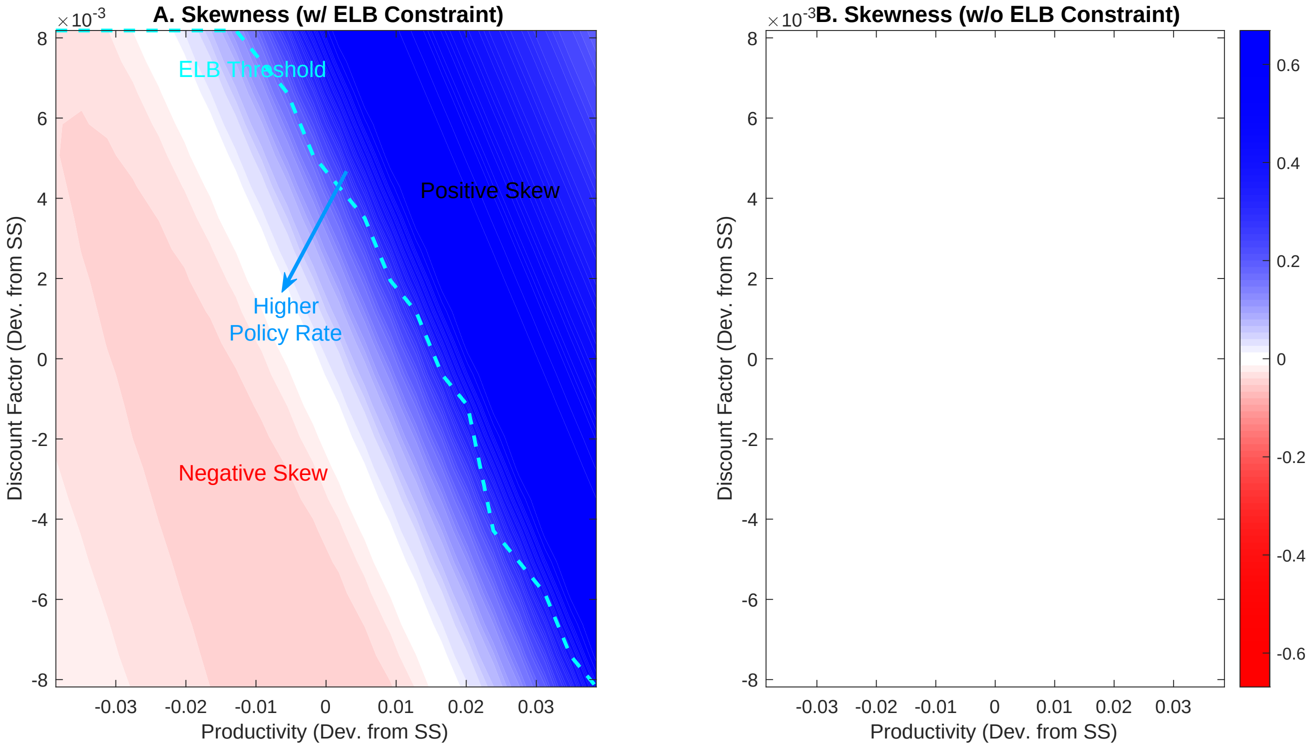

Figure I.3: **Skewness of the Policy Rate Forecast Distribution**

Note: The colors indicate the size of skewness. The light blue dashed line shows the values of the discount factor and productivity (both deviation from their steady states) when the ELB constraint starts to bind.

Incorporating supply shocks in the model appears particularly relevant for understanding the dynamics of policy rate skew since the onset of the pandemic crisis, during which supply shortages have been widely reported and have been regarded as a crucial contributor to the rise in inflation since early 2021. High inflation readings and a tight labor market prompted the Federal Reserve to lift the federal funds rate from the ELB in March 2022. Our model indeed suggests that a combination of a negative supply shock and a positive demand shock (or a negative shock to the discount factor) that lifts the policy rate from the ELB leads to a transition from positive skew to negative skew in the forecast distribution of the policy rate, qualitatively consistent with the data.[47]

In sum, we find that the ELB constraint can naturally produce time-varying skewness in the policy rate distribution which can be either negative or positive, even when we extend our setup to allow for both demand and supply shocks.

[47]Whereas the skewness measures dropped notably towards the end of 2021, the (approximately) one-year ahead aggregate distributions of the federal funds rate since January 2022 assign only a small probability of the federal funds rate being at the ELB. However, it is worth noting that our model also requires only a small probability of the policy rate returning to the ELB for the policy rate distribution to have negative skew. While we do not rule out other channels that lead to negative skew, we believe the mechanism based on the ELB can at least in part explain the changes in skew observed since the start of the pandemic crisis.

J Additional Regression Results

J.1 Regression Results with the Full Set of Respondents

In the main text, we provided evidence from regression analysis that the levels and changes in the BCS measure were better explained by the modes from the SPD rather than the means. This analysis was conducted on a subset of respondents that fill out both the BCS and the SPD. In this section, we provide further robustness by conducting similar analysis with the full sample of respondents. Using the full sample of respondents from both surveys is less ideal because there are many respondents in the BCS that are not included in the SPD. Despite this concern, we continue to find strong evidence that the modes from the SPD are statistically associated with the BCS measure in a significant way, as we discuss below.

Similar to the results in the main text, Table J.1 shows that for levels, the mode has a significant coefficient close to 1 and when both the mode and mean are included, the mode's coefficient is much closer to one compared to the mean's coefficient. We see a similar dynamic hold for the second moments. In the bottom left panel, the standard deviation of the mode is highly significant while the mean is insignificant and has the opposite sign when both are included. The R-squared also increases noticeably when the mode is in the regression. These results generally hold for regressions using first differences (right panels). This suggests that despite concerns associated with using the full set of respondents in both measures, the results in the main text continue to hold. The BCS measure appears to reflect the mode, rather than the mean.

There are advantages in focusing on the results only using dealers who answered both the SPD and BCS. This can be seen by comparing Table 2 and Table J.1. According to the top left panels, the coefficient of the modes are notably higher and closer to 1 for the subset regression. In the bottom two panels, while the R-squared improve significantly when the mode is included in the regression, regardless of the sample of respondents, the subset results show an even larger improvement.

Meanwhile, the regression with the first differences of BCS as the dependent variable appears to show somewhat stronger evidence that the BCS is closer to the mode using the full set of respondents (top right panels). However, as we show below (Appendix J.2) this result turns out to be not robust to the sample period.

Dependent Variable: BCS			
Mode	0.95^{***}		0.82^{***}
	(0.03)		(0.31)
Mean		1.12^{***}	0.16
		(0.07)	(0.42)
Constant	0.12^{**}	0.04	0.11
	(0.05)	(0.07)	(0.10)
R-squared	0.98	0.97	0.98

Dependent Variable: Δ BCS			
Δ Mode	0.88^{***}		0.98^{***}
	(0.08)		(0.15)
Δ Mean		0.89^{***}	−0.11
		(0.11)	(0.23)
Constant	0.00	0.01	0.00
	(0.01)	(0.01)	(0.01)
R-squared	0.77	0.69	0.77

Dependent Variable: σ(BCS)			
σ(Mode)	0.98^{***}		0.74^{***}
	(0.10)		(0.29)
σ(Mean)		1.43^{***}	0.42
		(0.19)	(0.36)
Constant	0.07^{***}	0.04	0.05^{**}
	(0.02)	(0.03)	(0.02)
R-squared	0.74	0.66	0.75

Dependent Variable: Δ σ(BCS)			
Δ σ(Mode)	0.75^{***}		0.65^{*}
	(0.27)		(0.34)
Δ σ(Mean)		1.038^{**}	0.26
		(0.480)	(0.28)
Constant	−0.00	-0.018^{**}	−0.01
	(0.01)	(0.008)	(0.01)
R-squared	0.41	0.24	0.42

Table J.1: **Full Sample Regression Results**

Note: Numbers in parentheses denote Newey-West standard errors. Stars next to the regression coefficients indicate statistical significance levels. ***: 1%, **:5%, *:10%.

J.2 Results Excluding the Covid Period (2020.3 ∼)

In this section, we briefly discuss regression results when we exclude observations starting in March 2020, when the pandemic crisis intensified in the U.S. Overall, the results we described in Section 4.1 and Appendix J.1 continue to hold, with some improvements in fit. Although we do not report the full analysis here (results are available upon request), we show the regression results that were most impacted by removing the post-Covid sample in Table J.2.

Dependent Variable: Δ BCS (Subset)			
Δ Mode	0.89^{***}		0.97^{***}
	(0.07)		(0.14)
Δ Mean		0.79^{***}	−0.10
		(0.13)	(0.14)
Constant	−0.00	0.01	−0.00
	(0.01)	(0.02)	(0.01)
R-squared	0.73	0.46	0.73

Dependent Variable: Δ BCS (Full Sample)			
Δ Mode	0.92^{***}		0.88^{***}
	(0.08)		(0.10)
Δ Mean		0.85^{***}	0.05
		(0.14)	(0.12)
Constant	−0.00	0.01	−0.00
	(0.01)	(0.02)	(0.01)
R-squared	0.745	0.52	0.75

Table J.2: **Regression Results (2011.1 ∼ 2020.1)**

Note: Numbers in parentheses denote Newey-West standard errors. Stars next to the regression coefficients indicate statistical significance levels. ***: 1%, **:5%, *:10%.

The table shows the results when the dependent variable is the first difference of the BCS, both for the subset of respondents (left) and the full set of respondents (right). Excluding the Covid period, we do not observe the stronger evidence for the BCS being closer to the mode using the full set of respondents. In contrast, especially in the regression that includes both the mean and the mode, the subset results indicate that the mode coefficient is very close to 1, while the performance of the mean regression is somewhat reduced.