Learning and subjective beliefs about good and bad inflation ranges*

Mohammad Ghaderi † Sang Byung Seo ‡ Ivan Shaliastovich §

July 14, 2025

Abstract

We identify desirable/undesirable inflation outcomes under subjective beliefs by comparing survey-based and risk-adjusted distributions of inflation. Intuitively, investors dislike inflation at both extremes, preferring a range in the middle. This "good inflation" region, which investors associate with lower-than-average marginal utility, varies substantially over time in position and width, revealing time-varying preferences across inflation ranges. Different ranges contribute to the inflation risk premium with mixed signs, offsetting each other and often masking important insights into the pricing of inflation risk. We rationalize these empirical patterns using a model where investors learn and update beliefs about hidden deflationary and inflationary recession states.

*We thank Hengjie Ai, Ravi Bansal, Will Bazely, Hui Chen, Olivier Darmouni, Anthony Diercks, Bjørn Eraker, Amit Goyal, Mohammad Jahan-Parvar, Christopher Jones, Lubomir Litov, Daniel Neuhann, David Schreindorfer, Yu Shan, Hiroatsu Tanaka, David Weinbaum, Lai Xu, Ram Yamarthy, Tony Zhang, and seminar participants at the Columbia Workshop in New Empirical Finance, Derivatives and Asset Pricing Conference, DePaul University, Federal Reserve Board, Michigan State University, Syracuse University, University of Kansas, University of Wisconsin-Madison, Virtual Household Finance Seminar, and World Symposium on Investment Research.

†School of Business, University of Kansas; mghaderi@ku.edu

‡Wisconsin School of Business, University of Wisconsin-Madison; sang.seo@wisc.edu

§Wisconsin School of Business, University of Wisconsin-Madison; ivan.shaliastovich@wisc.edu

1 Introduction

Inflation risk has long been recognized as an important risk factor in the asset pricing literature, especially for valuing nominal securities. Unexpected changes in inflation not only erode the real value of nominal cash flows but also introduce an additional pricing channel – the so-called inflation risk premium. This premium arises because investors expect future inflation to be correlated with future growth and, hence, marginal utility. Depending on their outlook, investors may perceive future inflation as a positive or negative outcome, leading to different signs and magnitudes of the inflation risk premium. With the drastic shift to the post-pandemic inflationary environment, understanding this premium and its implications is now at the center of discussion.

Despite its importance, properly capturing inflation risk in an asset pricing framework is challenging due to its dual nature. Both high inflation and excessively low inflation/deflation are typically viewed as bad economic outcomes. Unlike other risk factors, where more exposure is either unequivocally good (e.g., higher growth) or bad (e.g., higher uncertainty), inflation risk requires a more nuanced approach. Models treating inflation risk as a monotonic factor, either good or bad, inevitably miss the full picture. Even assuming that perceptions of inflation risk alternate between good and bad over time may not be sufficient, if investors have different preferences for various ranges of inflation outcomes at a given point in time.

To illustrate, consider a scenario where investors fear a deflationary recession, as in the Great Depression of the 1930s, as well as an inflationary recession, like the stagflation of the 1970s. Since investors dislike extreme future inflation outcomes and associate them with low growth, both ends of the inflation spectrum are priced negatively. However, this also implies that as we move toward the center of the spectrum, there should exist a moderate range of inflation that investors prefer, seeing it as correlated with higher growth. In this middle range, inflation risk is priced positively. This example highlights why accounting for investors' preferences across different inflation ranges is critical for understanding the inflation risk premium, which reflects the sum of premiums generated over the entire range of inflation outcomes: (i) the high

inflation range is disliked by investors and generates a positive premium; (ii) the moderate range is favored by investors and generates a negative premium; (iii) the deflationary range generates a negative premium, despite being disliked by investors, because the inflation outcomes in this range themselves are negative. As such, the premiums across different inflation ranges often offset one another. Even a small overall premium does not necessarily indicate that inflation risk is weakly correlated with future growth or that investors assign a low price of risk. By overlooking investors' preferences for inflation ranges, we lose much of the underlying dynamics behind the pricing of inflation risk.

In this paper, we extract investors' preferences for different inflation ranges and examine how they vary over time. To this end, we compare the distribution of future inflation under two different probability measures: the survey-based distribution, reflecting survey respondents' subjective beliefs, and the risk-adjusted distribution, estimated from inflation caps and floors. We then show that the ratio between the two distributions, also known as the Radon-Nikodym derivative, has a direct connection with the pricing kernel (equivalently, marginal utility). Specifically, if the risk-adjusted probability is higher than the survey-based probability for a certain inflation realization, it implies that investors perceive it as a bad economic outcome associated with higher-than-average future marginal utility. Conversely, if the risk-adjusted probability is lower, investors regard it as a good economic outcome associated with lower-than-average future marginal utility.

This approach implicitly assumes that the beliefs of survey respondents are meaningful and align closely with those of marginal investors in the inflation derivatives market. While this is a strong assumption, it is not an uncommon one; previous studies using survey forecasts to gauge investors' expectations make a similar assumption. In our analysis, we focus on the density forecasts provided by two key surveys: the Survey of Professional Forecasters (SPF) and the Survey of Primary Dealers (SPD). The SPF, conducted by the Federal Reserve Bank of Philadelphia, targets professional forecasters, many of whom work in the financial sector. The SPD surveys primary dealers, mostly large banks or securities companies, authorized to trade with the Federal Reserve Bank of New York. We exclude household surveys to ensure that the

forecasts we rely on better reflect those of market participants.

The density forecasts reported by the SPF and SPD represent the probabilities of the average annualized inflation rate falling into various ranges. From the SPF, we obtain the distribution of inflation, based on the GDP price index, for the next calendar year. The data frequency is quarterly, and we take the time series from the fourth quarter of 2009 when our sample period starts. From the SPD, we obtain the distribution of longer-term inflation, in terms of the CPI, over a 5-year horizon. Since the survey is conducted right ahead of each FOMC meeting, we have eight monthly observations per year, beginning in 2014.

These survey-based distributions of inflation can directly be compared with the risk-adjusted ones, thanks to inflation caps and floors – essentially call and put options written on future inflation rates. As demonstrated by Breeden and Litzenberger (1978), option prices at various strikes allow us to extract the risk-adjusted probability density, calculated as the second derivative of the option price with respect to the strike. In our sample period, which begins in October 2009, inflation options are traded with various maturities over a wide range of strike inflation rates, from -3% to 7%. This allows us to apply the nonparametric estimation approach of Aït-Sahalia and Duarte (2003) to obtain the risk-adjusted distribution of inflation on each trading day. By averaging the estimated distributions within each quarter or month, we generate measures that align with the timing, frequency, and horizon of the SPF and SPD.

We find that the probability ratio (or the probability distortion) between the risk-adjusted and subjected probability measures shows a clear U-shaped pattern. This confirms that investors indeed dislike both high and low inflation environments in the future while favoring moderate inflation outcomes in between. During our sample period, the probability ratio goes through significant time series variation, leading to different good and bad inflation regions over time. For example, in the second quarter of 2018, the good inflation region for the next calendar year was relatively narrow and tightly centered around the Fed's 2% target. However, in the second quarter of 2020, when deflation concerns emerged due to the pandemic, the good inflation region expanded noticeably to the downside. By the second quarter of 2022, this pattern was reversed: the good inflation region shifted upward, reflecting heightened fears of

inflationary pressures. We document qualitatively similar but much less pronounced patterns for the 5-year horizon, consistent with the view that long-term inflation expectations remain well anchored, even in light of the disruptions in 2020 and 2022.

We emphasize that this pattern is not just driven by the level of inflation but also shaped by the economic news inflation conveys. To demonstrate, let us revisit June 2022, when year-on-year inflation exceeded 9%. At that time, investors perhaps recognized that a low inflation outcome in the following year could materialize under two different scenarios with mixed economic implications. On the one hand, it could indicate a *soft landing* scenario where we see a smooth return to a low inflation environment without triggering a recession. On the other hand, it could indicate a potential *hard landing* scenario marked by severe economic contraction. Given the highly persistent nature of inflation, the latter scenario carried more weight when considering a sharp decline in inflation within a year. As a result, inflation ranges around 1%, and even 2% in some quarters, were perceived to be bad and fell outside the good region. Interestingly, we find that the 2% inflation rate is always located within the good region for the 5-year horizon; in the long run, good economic outcomes (characterized by lower-than-expected marginal utility) come with inflation leveling out to a moderate range around 2%, regardless of whether near-term inflation is high or low.

Investors' preferences toward inflation can be more effectively analyzed with the premiums attached to different ranges of inflation. The inflation risk premium, as a whole, is defined as the difference between the risk-adjusted and survey-based expected inflation rates. Equipped with the conditional distributions of inflation under the two measures, we not only calculate the inflation risk premium but also decompose it over different ranges, using the methodology proposed by Beason and Schreindorfer (2022). Our results show that the inflation risk premium and its components fluctuate significantly over time. Comparing 2018, 2020, and 2022 again, we argue that focusing solely on the total inflation risk premium can overlook valuable insights, as it often masks the underlying dynamics of investors' preferences and expectations over different inflation ranges. By breaking down the premium, we uncover which inflation ranges contribute disproportionately to the overall premium and how these contributions shift over time, providing

a deeper and more nuanced understanding of inflation risk.

Finally, we show that the empirical patterns we document can be rationalized using a model with learning. We assume that investors have imperfect information about the true state of the economy, which switches among three regimes: (i) normal/favorable regime with relatively high consumption growth and moderate inflation; (ii) deflationary recession regime marked by low growth and very low inflation, or even deflation; and (iii) inflationary recession regime characterized by low growth paired with a sharp increase in inflation. Since the true state is not directly observable, investors form their subjective beliefs based on historical consumption and inflation realizations. We show that the fear of the two recessionary regimes leads to a U-shaped pricing kernel, projected on future inflation. In line with empirical evidence, investors' preferences across different inflation ranges fluctuate over time, as they learn and update their beliefs.

Literature review

Our paper is built on prior studies that exploit inflation caps and floors to gauge market-based inflation expectations. Kitsul and Wright (2013) estimate risk-adjusted inflation densities from inflation options and examine their responses to macroeconomic announcements through event-study regressions. Their work is closely related to ours in that they also document a U-shaped pattern in the pricing kernel with respect to future inflation. The key distinction is that they construct the pricing kernel based on physical inflation densities, which are obtained from estimating econometric models using historical data. In contrast, we build the pricing kernel based on survey-based inflation densities, directly aiming to study investors' inflation preferences under subjective beliefs. In our approach, the pricing kernel is estimated without relying on past inflation time series: both risk-adjusted and survey-based distributions are forward-looking, allowing us to estimate the conditional pricing kernel and characterize its time series variation without imposing any econometric model. Equipped with this conditional information, our focus is to identify the good and bad inflation ranges perceived by investors and explore their implications for the inflation risk premium.

In addition to Kitsul and Wright (2013), a few more studies highlight the usefulness of the information embedded in inflation options in addressing macroeconomic questions. Fleckenstein, Longstaff, and Lustig (2017) document that deflation risk is significantly priced with a high market price of risk, revealing that the inflation option market places substantial weight on deflationary scenarios. Using inflation options, Mertens and Williams (2021) tackle the issue of multiple equilibria in a New Keynesian model with the zero lower bound by finding empirical support for the target equilibrium where the central bank largely succeeds in stabilizing the economy. Hilscher, Raviv, and Reis (2022) develop a copula estimator for the option-implied joint distribution of future inflation rates and study how much public nominal debt might be inflated away. In ongoing work, Hilscher, Raviv, and Reis (2024) extract physical probabilities of inflation disasters from inflation option prices, adjusting for the effect of inflation on inflation option payoffs, horizons, and risk premiums.

Our paper also relates to the literature on survey expectations of inflation. Various surveys have been extensively studied in both research and policy analysis.[1] These surveys have demonstrated significant value in forecasting future inflation (Ang, Bekaert, and Wei, 2007). While the point forecasts of expected inflation have received considerable attention so far, relatively little focus has been placed on the density forecasts provided by some of the surveys. A few studies utilize these distribution surveys, but their primary emphasis is rather on their adequacy (Diebold, Tay, and Wallis, 1997) or their potential to provide extra information about the level and uncertainty of inflation (Rich and Tracy, 2010; Kenny, Kostka, and Masera, 2014; Clements, 2018). Unlike prior work, we explore an understudied aspect of distributional forecasts: the insights they offer into individuals' preferences over different inflation outcomes, particularly when combined with data from inflation derivatives.

Recently, there has been growing interest in understanding the subjective beliefs of various economic agents, particularly with regard to the role of information rigidity in belief formation. Mankiw and Reis (2002) propose a model where agents update their beliefs infrequently due

[1]These include, but are not limited to, the Livingston Survey, the Survey of Professional Forecasters, the University of Michigan Surveys of Consumers, and the New York Fed Survey of Consumer Expectations.

to the costs of acquiring information. Analyzing survey data from consumers, firms, central bankers, and professional forecasters, Coibion and Gorodnichenko (2012) show that the patterns of forecast errors indicate the presence of information rigidity. Follow-up work by Coibion and Gorodnichenko (2015) further demonstrate that forecast errors can be predicted by forecast revisions and quantify the degree of information rigidity. Intriguingly, Bordalo, Gennaioli, Ma, and Shleifer (2020) uncover overreaction rather than underreaction when the methodology of Coibion and Gorodnichenko (2015) is applied to individual forecasts. They reconcile this overreaction with the underreaction observed in consensus forecasts through a diagnostic expectations model.[2] In line with these empirical and theoretical findings, we allow our model to depart from the Bayesian benchmark by incorporating information rigidity, which we calibrate to match the empirical relation between forecast errors and forecast revisions.

Lastly, our paper contributes to the asset pricing literature on investors' learning about inflation risk. David and Veronesi (2013) rationalize the time-varying comovement and volatility of stock and Treasury bond prices based on investors' learning about unobservable economic regimes governing consumption, earnings, and inflation. Their model shows that inflation news can signal either positive or negative future growth, depending on the prevailing regime, altering the signs of stock-bond correlations. Bianchi, Lettau, and Ludvigson (2022) discuss how monetary policy shocks can have long-lasting effects on real variables within a learning model. In their framework, learning about the duration of monetary policy regimes, coupled with the fading memory of past regimes, leads to persistent changes in asset valuations and real interest rates. Andrei and Hasler (2023) examine how investors' learning about the Fed's ability to manage inflation directly affects equity market dynamics; when the Fed's credibility wanes, investors begin to view inflation as more persistent, increasing both the risk premium and volatility in the stock market.

[2]Attempts to measure the degree of information rigidity and to understand its role in belief formation have been made across various contexts. These include, but are not limited to, studies on firm expectations (Coibion, Gorodnichenko, and Kumar, 2018; Coibion, Gorodnichenko, and Ropele, 2020), equity market expectations (Bordalo, Gennaioli, and Shleifer, 2018; Bordalo, Gennaioli, Porta, and Shleifer, 2019; Bouchaud, Krueger, Landier, and Thesmar, 2019), the influence of Fed communications on household beliefs (Coibion, Gorodnichenko, and Weber, 2022), and the role of data-generating processes and forecasters' information sets in large-scale randomized experiments (Afrouzi, Kwon, Landier, Ma, and Thesmar, 2023).

The rest of the paper is organized as follows. Section 2 reviews the survey-based distribution of inflation and its use in our study. Section 3 introduces inflation options and explains how the risk-adjusted distribution of inflation can be estimated. Section 4 analyzes what these distributions jointly reveal about investors' preferences for inflation ranges. Section 5 presents an economic model with learning that can explain the data. Section 6 concludes.

2 Survey-based distribution of inflation

This section provides an overview of survey-based density forecasts of price indices and explains how we leverage them in our study.

2.1 Survey of Professional Forecasters

The Survey of Professional Forecasters (SPF) is a quarterly survey of experts affiliated with financial and non-financial institutions, covering a wide range of macroeconomic and financial variables. The survey began in the first quarter of 1968 and was initially conducted by the American Statistical Association and the National Bureau of Economic Research. Since 1990, the Federal Reserve Bank of Philadelphia has administered the survey. Each quarter, the SPF sends survey questionnaires to participants after the release of the Bureau of Economic Analysis (BEA)'s advance report on the national income and product accounts, ensuring that the panelists' information sets include the latest data. Beginning with the 2005 survey, participants have been required to submit their projections by the second week of the middle month of each quarter. The results are then released to the public before the BEA's second report on the national income and product accounts.

In our analysis, we rely on a special section of the SPF called "Mean Probability Forecasts," which provides probability forecasts for various price indices, both at the individual respondent level and averaged across respondents. One of the longest-standing forecasts in the survey is the PRPGDP (Probability of Changes in GDP Price Index). This variable reports the average probabilities that individual panelists assign to the *annual-average over annual-average* percent

change in the chain-weighted GDP price index falling into various ranges. The term "annual-average over annual-average change" refers to the percent change in the average level of GDP prices from one year to the next, with the annual average being calculated as the mean of the quarterly levels across all four quarters of a calendar year. The underlying index for PRPGDP has changed over time. Before the adoption of the chain-weighted GDP price index, from 1992 to 1995, the survey asked respondents about changes in the implicit deflator for GDP with fixed weights. Prior to 1992, the implicit deflator for GNP with fixed weights was used.

The probability ranges for PRPGDP have undergone several modifications since the survey's inception. Focusing on the subsample relevant to our study, which begins in the last quarter of 2009, there was only one change.[3] From 1992 to 2013, respondents assigned probabilities to ten buckets of outcomes: below 0%, from 0% to 8% in 1% increments, and above 8%. Starting in the first quarter of 2014, the upper limit was reduced to 4%, but the buckets were refined to narrower intervals: below 0%, from 0% to 4% in 0.5% increments, and above 4%. All inflation buckets are right-open intervals. It is important to note that the probability estimates pertain to "fixed-event forecasts." In each quarterly survey of a given calendar year, respondents provide their probability estimates for changes in the current calendar year (i.e., the year in which the survey is conducted) and the following calendar year. This means that the forecast horizon varies across different quarters of the year. For example, the survey conducted in the fourth quarter of each calendar year includes a *nowcast* for the current year and a forecast for the next year. In our analysis, we focus on the survey responses for the subsequent calendar year to ensure that the reported probabilities represent valid forecasts.

Table 1 presents the time series mean, standard deviation, and the 5th, 50th, and 95th percentiles of the PRPGDP forecasts.[4] Panel A shows the descriptive statistics for the period

[3]Our sample period of interest begins in October 2009 due to the availability of inflation option data.

[4]Following Bansal and Shaliastovich (2010), we exclude a small number of outlier responses from the SPF survey data to ensure the reliability of our analysis. Outliers are identified based on each respondent's probability assessments of the tails – specifically, the probability assigned to inflation being below 1% or above 3% in each survey month. We compute the first quartile (Q1) and the third quartile (Q3) of these tail probabilities across all respondents and apply a standard interquartile range (IQR) rule: responses are flagged as outliers if their tail probabilities fall below $Q1 - 1.5 \times IQR$ or above $Q3 + 1.5 \times IQR$. Any respondent exceeding these bounds is excluded from that month's analysis. Our main findings remain robust when these outliers are retained or identified using alternative criteria.

Panel A: PRPGDP with narrower buckets (2014Q1 - 2024Q2)										
	$(-\infty, 0)$	$[0, 0.5)$	$[0.5, 1)$	$[1, 1.5)$	$[1.5, 2)$	$[2, 2.5)$	$[2.5, 3)$	$[3, 3.5)$	$[3.5, 4)$	$[4, \infty)$
Mean	0.21	0.84	3.34	11.06	26.74	29.19	14.66	6.59	4.05	3.33
SD	0.30	1.19	3.39	8.04	13.22	10.45	8.04	7.84	8.08	8.42
5th	0.00	0.00	0.00	0.17	2.74	6.45	4.48	0.16	0.00	0.00
50th	0.09	0.44	2.73	9.79	29.45	30.24	13.74	2.75	0.40	0.10
95th	0.72	2.37	9.75	25.26	43.21	42.56	28.23	21.51	25.07	23.17
Panel B: PRPGDP with wider buckets (2009Q4 - 2024Q2)										
	$(-\infty, 0)$	$[0, 1)$	$[1, 2)$	$[2, 3)$	$[3, 4)$	$[4, \infty)$				
Mean	0.67	6.37	39.48	40.78	9.83	2.88				
SD	0.95	5.89	17.39	14.06	12.96	7.13				
5th	0.00	0.00	4.69	22.19	0.28	0.00				
50th	0.17	5.01	42.39	38.85	5.06	0.35				
95th	2.53	18.82	63.06	63.68	43.49	19.60				

Table 1: Descriptive statistics for PRPGDP. This table presents the mean, standard deviation, and the 5th, 50th, and 95th percentiles of PRPGDP (Probability of Changes in GDP Price Index) forecasts from the Survey of Professional Forecasters. We focus on the survey responses for the next calendar year. Panel A covers 2014Q1 to 2024Q2 with narrower probability buckets: $(-\infty, 0)$, $[0, 0.5)$, $[0.5, 1)$, $[1, 1.5)$, $[1.5, 2)$, $[2, 2.5)$, $[2.5, 3)$, $[3, 3.5)$, $[3.5, 4)$, and $[4, \infty)$. Panel B provides the same statistics for 2009Q4 to 2024Q2 with wider buckets: $(-\infty, 0)$, $[0, 1)$, $[1, 2)$, $[2, 3)$, $[3, 4)$, and $[4, \infty)$. All values are in percentages.

beginning in 2014, during which the survey used narrower buckets with finer intervals. During this period, inflation was predominantly expected to hover around the Fed's 2% target, with the average probabilities concentrated in the $[1.5\%, 2\%)$ and $[2\%, 2.5\%)$ buckets. The tail portions of the distribution show much smaller probabilities. The average probability of the lowest bucket (below 0%) is 0.21%, reflecting the consensus that deflation was considered an unlikely outcome. Similarly, the highest bucket (above 4%) retains a modest average probability of 3.33%. However, the standard deviations of these bucket probabilities are quite sizable relative to their means. Related, we observe substantial differences between the 5th and 95th percentiles, particularly in the high inflation regions. The probability of over 4% inflation is 23.17% at the 95th percentile, indicating that forecasters did occasionally assign significant probabilities to these extreme outcomes.

Panel B extends the data to cover the entire sample period of our interest, from the last quarter of 2009 to the second quarter of 2024. The descriptive statistics shown in this panel are based on the wider buckets that were in use before 2014. To achieve this, we aggregate the finer intervals from 2014 onward into broader 1% increments that correspond to the earlier survey

structure. For instance, the probabilities for the $[1\%, 2\%)$ bucket in Panel B were computed by summing the probabilities assigned to the $[1\%, 1.5\%)$ and $[1.5\%, 2\%)$ intervals in the case of the post-2014 data. The results remain similar in the extended sample. Forecasters largely anticipated that inflation would fall within the $[1\%, 3\%)$ range, assigning it a total probability of more than 80% on average. Compared to Panel A, we notice that the distribution in Panel B is slightly shifted to the left, putting higher probabilities on lower inflation outcomes. This is because the full sample period in Panel B includes the low-inflation period following the global financial crisis, unlike the one in Panel A.

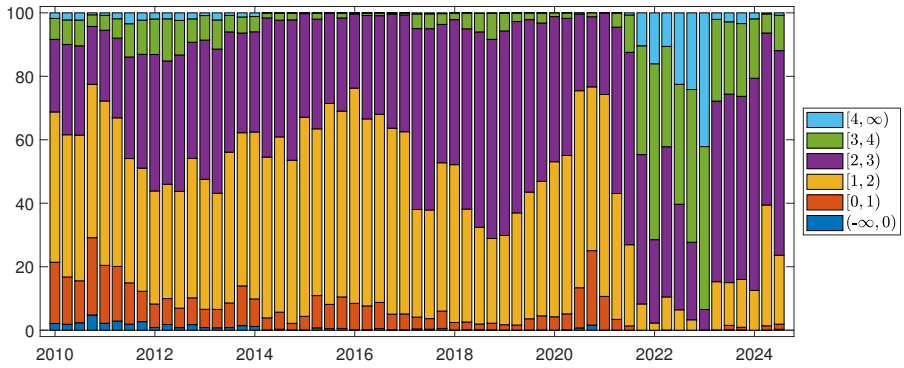


Figure 1: Time series of PRPGDP. This figure illustrates the time series of PRPGDP (Probability of Changes in GDP Price Index) forecasts from the Survey of Professional Forecasters, with each color representing one of the six probability ranges: $(-\infty, 0)$, $[0, 1)$, $[1, 2)$, $[2, 3)$, $[3, 4)$, and $[4, \infty)$. The sample period is from 2009Q4 to 2024Q2. All values are in percentages.

Figure 1 visualizes the variation in the forecasts for PRPGDP over time, with each color representing one of the six probability buckets in Panel B of Table 1. The figure highlights significant shifts in inflation expectations across different periods. We first notice a non-negligible probability of deflation following the financial crisis around 2008-2009 and again during the COVID-19 crisis in 2020. During these periods, the probabilities of the GDP price index falling into lower buckets (below 0% or 1%) were significantly elevated, reflecting concerns about deflationary recessions. In contrast, the inflationary environment emerged from mid-2021 to 2023, with a substantial increase in the probabilities for higher inflation ranges. Accordingly, forecasters assigned a greater likelihood to inflation rates in the $[2\%, 3\%)$, $[3\%, 4\%)$, and $[4\%, \infty)$ buckets, signaling much higher expected inflation. Overall, the figure encapsulates the significant

time series variation in inflation expectations, in response to major economic events/conditions.

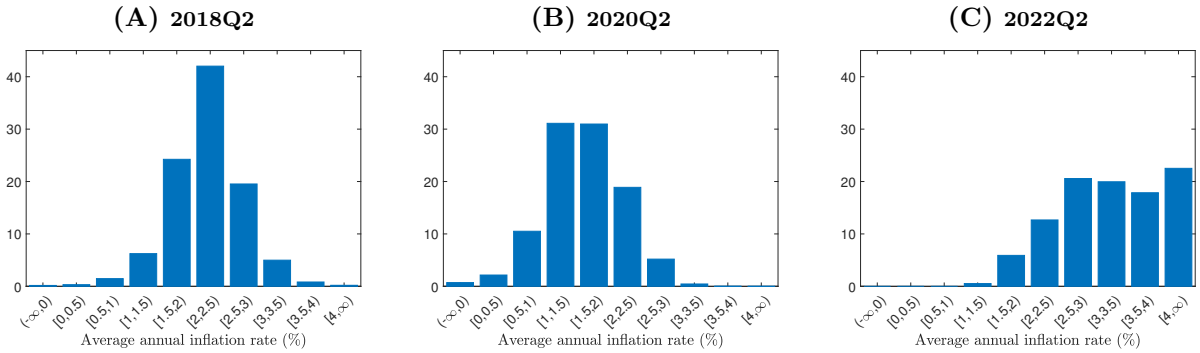


Figure 2: PRPGDP densities. This figure presents the survey results for PRPGDP (Probability of Changes in GDP Price Index) forecasts from the Survey of Professional Forecasters for three different quarters: the second quarter of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). The x-axis represents the average annual inflation rate buckets: $(-\infty, 0)$, $[0, 0.5)$, $[0.5, 1)$, $[1, 1.5)$, $[1.5, 2)$, $[2, 2.5)$, $[2.5, 3)$, $[3, 3.5)$, $[3.5, 4)$, and $[4, \infty)$. All values are in percentages.

To highlight the distinct patterns in inflation expectations over time, Figure 2 juxtaposes the distributions of PRPGDP for three different quarters: the second quarter of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). Inflation expectations were centered around the $[2\%, 2.5\%)$ bucket in 2008, close to the Fed's target. However, the distribution became more dispersed and substantially shifted to the left two years later, due to fears of deflationary contractions at the peak of the COVID-19 crisis. The pattern saw a dramatic shift to the right in 2022, with significant probabilities in the $[3\%, 4\%)$ and above 4% ranges, as forecasters were expecting much higher inflation coming their way.

While our main variable is PRPGDP, the SPF also reports panelists' average probabilities for other price indices. For instance, the PRCCPI (Probability of Core CPI Inflation) and the PRCPCE (Probability of Core PCE Inflation) represent the probabilities of *fourth-quarter over fourth-quarter* changes in the core CPI and the core PCE falling into the same ten post-2014 buckets as PRPGDP. The fourth-quarter level of each index is defined as the average of the monthly levels over the three months of the fourth quarter. We provide descriptive statistics for PRCCPI and PRCPCE in the Internet Appendix. The distributional forecasts for these indices show patterns similar to those of PRPGDP in general. Yet, they exhibit lower time series volatility, resulting in tighter distributions. This is expected, as these price measures

exclude food and energy prices, which tend to be more volatile over time.

2.2 Survey of Primary Dealers

Beginning in 2011, the New York Fed's Open Market Trading Desk has conducted the Survey of Primary Dealers (SPD) ahead of each FOMC meeting. This survey targets primary dealers to gather their expectations on key variables such as the federal funds rate, the future size of the Federal Reserve's balance sheet, and inflation. The survey results are among the inputs used by Federal Reserve staff to evaluate market expectations regarding the economic outlook, monetary policy, and financial markets. Survey questions are published on the Federal Reserve Bank of New York's website approximately two weeks before each FOMC meeting. Summaries of the survey results are released around three weeks after each FOMC meeting, following the publication of the meeting minutes. Unlike the SPF, which provides individual-level responses, the SPD only reports average responses aggregated across all respondents.

	$(-\infty, 1)$	$(1, 1.5]$	$(1.5, 2]$	$(2, 2.5]$	$(2.5, 3]$	$(3, \infty)$
Mean	3.63	9.87	24.83	33.55	18.13	10.17
SD	1.06	3.72	7.88	5.39	6.86	9.07
5th	2.30	4.30	12.00	23.00	11.30	3.00
50th	3.00	11.00	28.50	35.00	14.00	5.00
95th	6.00	15.70	33.70	40.70	29.70	29.70

Table 2: Descriptive statistics for SPD inflation density forecasts. This table presents the mean, standard deviation, and the 5th, 50th, and 95th percentiles of the Survey of Primary Dealers inflation density forecasts. The survey is conducted ahead of eight regularly scheduled FOMC meetings every year and provides the likelihood that the average annualized CPI inflation rate over the next five years will fall into various ranges: $(-\infty, 1]$, $[1, 1.5)$, $[1.5, 2)$, $[2, 2.5)$, $[2.5, 3)$, and $[3, \infty)$. The sample period is from December 2014 to June 2024. All values are in percentages.

We use this survey to better understand longer-term inflation expectations – namely over the 5-year horizon. We focus on a specific survey question that asks respondents to assign probabilities to the likelihood that the average annualized CPI inflation rate over the next five years will fall into various ranges. Since its introduction in December 2014, this question has been consistently included in the survey, with the exception of April 2020. Initially, respondents assigned probabilities to six ranges: below 1%, from 1% to 3% in 0.5% increments, and above

3%. Since June 2022, two additional 0.5% ranges have been added at both ends of the spectrum. The inflation buckets for the SPD are left-open intervals.

Table 2 presents the descriptive statistics across the six ranges included in the survey since December 2014. Respondents consistently expected the 5-year average inflation to remain very close to the Fed's 2% target, with the total average probability of 58% for the $(1.5\%, 2\%]$ and $(2\%, 2.5\%]$ ranges. Compared to the 1-year horizon forecasts reported in Panel A of Table 1, the standard deviations of these central buckets are lower, consistent with the notion that long-term inflation expectations are anchored.

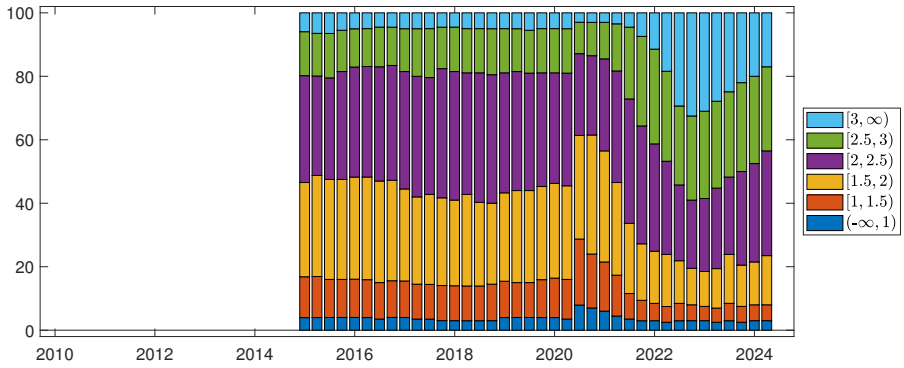


Figure 3: Time series of SPD inflation density forecasts. This figure illustrates the time series of the 5-year density forecasts for the Survey of Primary Dealers, with each color representing one of the six probability ranges: $(-\infty, 1]$, $[1, 1.5)$, $[1.5, 2)$, $[2, 2.5)$, $[2.5, 3)$, and $[3, \infty)$. The sample period is from December 2014 to June 2024. All values are in percentages.

Figure 3 depicts the changes in the SPD inflation density forecasts over time, with each color corresponding to one of the six probability ranges shown in Table 2. To facilitate comparison with Figure 1, the original time series is converted to a quarterly frequency. Note that there are typically eight pre-scheduled FOMC meetings per year, two per quarter. Each bar represents the average survey results within the quarter. The probability of being in the $(1.5\%, 2\%]$ range or below has steadily decreased since the COVID-19 pandemic, while the probability of being in the $(2.5\%, 3\%]$ range or above has steadily increased. This trend is also evident in Figure 4, which presents survey responses collected before the FOMC meeting in June of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). From 2018 to 2020, inflation expectations transitioned from a stable outlook centered around 2% to a more dispersed distribution tilted toward the

lower range. By 2022, the outlook was reversed, with a substantial shift to the right: the largest probability mass was allocated to over 3% average inflation, as inflation continued to rise.

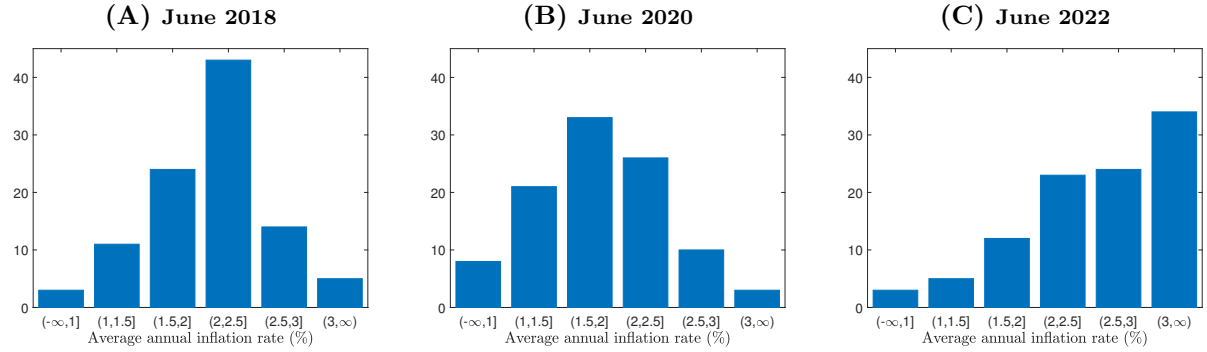


Figure 4: SPD inflation densities. This figure presents the distribution of 5-year inflation expectations from the Survey of Primary Dealers collected before the fourth FOMC meeting for three different years: June 2018 (Panel A), June 2020 (Panel B), and June 2022 (Panel C). The x-axis represents the average annual inflation rate buckets: $(-\infty, 1]$, $[1, 1.5)$, $[1.5, 2)$, $[2, 2.5)$, $[2.5, 3)$, and $[3, \infty)$. All values are in percentages.

Still, what stands out in Figure 3 is that the probability for the $(2\%, 2.5\%]$ bucket has remained similar since 2020, even during the high inflation episode in 2022; the lengths of the purple bars show relatively small variation in the figure. Overall, the bucket probabilities are much more stable over time compared to Figure 1, suggesting that survey respondents expect inflation to level out in the long run.

3 Risk-adjusted distribution of inflation

As discussed in the previous section, the survey-based distributions reflect survey respondents' beliefs about the future evolution of inflation. In this sense, such distributions can be viewed as the ones under the respondents' subjective probability measure. This section now examines the distribution of inflation under a risk-adjusted probability measure, which can be estimated nonparametrically from the prices of inflation options.

3.1 Inflation caps and floors

Inflation caps and floors are financial derivatives that are used to hedge or speculate on inflation. These instruments are essentially options written on future inflation rates. The most commonly traded are zero-coupon and year-on-year contracts. While both types of contracts critically depend on future realized inflation, they differ in their payoff structure and the way they are exercised.

As the name suggests, a zero-coupon inflation option is a single-payment instrument that can be exercised only once at the contract's maturity. A zero-coupon inflation cap (floor), with a strike rate of k maturing in T years, is a call (put) option whose payoff depends on the difference between realized cumulative inflation over the option's lifetime and the annually compounded strike price $(1+k)^T$. Formally, the cap and floor payoffs can be expressed as:

$$\text{cpay}^{\text{ZC}}_{t+T} = \Big[\Pi_{t\to t+T} - (1+k)^T\Big]^+ \quad \text{and} \quad \text{fpay}^{\text{ZC}}_{t+T} = \Big[(1+k)^T - \Pi_{t\to t+T}\Big]^+, \tag{1}$$

where $\Pi_{t\to t+T} = (\text{CPI}_{t+T}/\text{CPI}_t)$ represents the gross inflation rate between time t and time $t+T$. Here, CPI_t and CPI_{t+T} represent the price levels, in terms of CPI, at inception and maturity, respectively. In simpler terms, if the average annualized inflation rate over the next T years, given by $\sqrt[T]{\Pi_{t\to t+T}} - 1$, exceeds the strike inflation rate k, the cap is exercised. Conversely, if the average annualized inflation rate is below the strike rate k, the floor is exercised.

On the other hand, a year-on-year inflation option, with the same strike and maturity, consists of T yearly caplets (floorlets) maturing consecutively every year. Each caplet/floorlet $j \in \{1, 2, \cdots, T\}$ can be seen as a 1-year zero-coupon inflation option starting at time $t+j-1$ and expiring at $t+j$. Hence, at the end of each year $t+j$, the contract pays the difference between year-on-year gross inflation realized over the year $\Pi_{t+j-1\to t+j} = (\text{CPI}_{t+j}/\text{CPI}_{t+j-1})$ and the strike value $(1+k)$, if exercised. This leads to the following T yearly payoffs, each of which is expressed by

$$\text{cpay}^{\text{YoY}}_{t+j} = \Big[\Pi_{t+j-1\to t+j} - (1+k)\Big]^+ \quad \text{and} \quad \text{fpay}^{\text{YoY}}_{t+j} = \Big[(1+k) - \Pi_{t+j-1\to t+j}\Big]^+, \tag{2}$$

where $j = 1, 2, \cdots, T$. For a 1-year maturity ($T = 1$), there is only one period and one payoff to consider. Thus, year-on-year options are identical to zero-coupon options for this maturity, and the two terms can be used interchangeably.

Comparing equations (1) and (2) makes clear the difference between zero-coupon and year-on-year inflation options. Let $\pi_{t+j} = \log \Pi_{t+j-1 \to t+j}$ denote the log inflation rate. In the case of a zero-coupon contract, the exercise decision is based on cumulative inflation $\Pi_{t \to t+T} = \exp\left(\sum_{j=1}^{T} \pi_{t+j}\right)$. In the case of a year-on-year contract, however, the exercise decision is made for each $\Pi_{t+j-1 \to t+j} = \exp(\pi_{t+j})$ separately every year $t+j$. Put differently, the former contract can be seen as an *option on a portfolio* of future log inflation rates (i.e., $\sum_{j=1}^{T} \pi_{t+j}$), whereas the latter contract can be seen as a *portfolio of options* on future log inflation rates (π_{t+j} individually).[5]

3.2 Option-implied density

The seminal work by Breeden and Litzenberger (1978) demonstrates that by analyzing European options with different strike prices, one can infer the risk-neutral distribution of future asset prices. Specifically, the risk-neutral probability density is derived as the second derivative of the option price with respect to the strike price. Applying a similar method, we extract the distribution of inflation implied by caps and floors.

First, consider a zero-coupon inflation cap with strike rate k and time to maturity T.[6] For notational convenience, let $K = (1+k)^T$ denote the strike value. The pricing relation implies that the cap price $C_t^{\mathrm{ZC}(T)}$ can be expressed as

$$C_t^{\mathrm{ZC}(T)} = D_t^{(T)}\, \mathbb{E}_t^{\mathbb{Q}_{t+T}}\left[\mathrm{cpay}_{t+T}^{\mathrm{ZC}}\right] = D_t^{(T)} \int_K^{\infty} (x - K)\, \tilde{q}_{t,t \to t+T}(x) dx, \tag{3}$$

[5] Exploiting this relation, Hilscher, Raviv, and Reis (2022) estimate the joint distribution of π_{t+j} by introducing a copula function.

[6] While we illustrate the method using a cap (i.e., call), the second derivative of a floor (i.e., put) leads to the same mathematical expression. In our empirical implementation in Section 3.3, we follow the literature and work with cap prices. We still make use of floor prices by converting them into equivalent cap prices based on put-call parity.

where $D_t^{(T)}$ denotes the time-t price of a T-year risk-free nominal zero-coupon bond and $\tilde{q}_{t,t\to t+T}$ represents the time-t probability density of $\Pi_{t\to t+T}$. Unlike Breeden and Litzenberger (1978), the pricing measure we choose is not the risk-neutral measure $\mathbb{Q}_*$, which takes the bank account as the numéraire: $C_t^{\mathrm{ZC}(T)} = \mathbb{E}_t^{\mathbb{Q}_*}\left[e^{-\int_t^{t+T} r_s ds} \times \mathrm{cpay}_{t+T}^{\mathrm{ZC}}\right]$. As shown in equation (3), $\mathbb{Q}_{t+T}$ is the forward measure where the price of a zero-coupon bond maturing at time $t+T$ is used as the numéraire (see also Kitsul and Wright, 2013). The distinction between the two martingale measures carries different importance depending on the context. For instance, when future interest rates $\{r_s | t \leqslant s \leqslant t+T\}$ are constant or deterministic, $e^{-\int_t^{t+T} r_s ds}$ simply collapses to the zero-coupon bond price $D_t^{(T)}$ and can be pulled out of the expectation. More generally, these two measures become identical if future interest rates and the underlying asset determining the option payoff are independent. This is why, for practical purposes, the distinction between the two is often overlooked for short-maturity equity options. In our context, however, it is difficult to assume such independence, as future interest rates and the option payoff are both affected by the future inflation path over a long horizon.[7]

The integral expression in equation (3) reveals that the zero-coupon cap price can be seen as a function of the strike price: $C_t^{\mathrm{ZC}(T)} = C_t^{\mathrm{ZC}(T)}(K)$. By the second fundamental theorem of calculus, twice differentiating the equation with respect to K leads to the following relation:

$$\frac{\partial^2 C_t^{\mathrm{ZC}(T)}(K)}{\partial K^2} = D_t^{(T)}\, \tilde{q}_{t,t\to t+T}(K),$$

which implies that the risk-adjusted probability density of the cumulative inflation rate $\Pi_{t\to t+T}$ can be obtained as the second derivative of the option price, scaled by the zero-coupon bond price. Note that the surveys discussed in Section 2 provide the distribution of future inflation in terms of the average annualized inflation rate $\sqrt[T]{\Pi_{t\to t+T}} - 1$, not the cumulative one. For a direct comparison, we define $q_{t,t\to t+T}$ to be the density of the average annualized inflation rate

[7]For a similar reason, the forward measure is frequently employed as a risk-adjusted probability measure for interest rate derivatives. See, for example, Li and Zhao (2009).

and derive it from $\tilde{q}_{t,t\to t+T}$ as follows:

$$q_{t,t\to t+T}(x) = \tilde{q}_{t,t\to t+T}\left((1+x)^T\right) \times T(1+x)^{T-1},$$

where the term $T(1+x)^{T-1}$ comes from the Jacobian determinant of the mapping from $\sqrt[T]{\Pi_{t\to t+T}} - 1$ to $\Pi_{t\to t+T}$.

While we obtain $q_{t,t\to t+T}$ from zero-coupon contracts, year-on-year contracts allow us to extract $q_{t,t+j-1\to t+j}$, the density of inflation over a yearly period starting $j-1$ years and ending j years from today. Recall that each caplet constituting the year-on-year cap can be seen as a forward-starting 1-year zero-coupon cap. The price of caplet j, say $C_{j,t}$, should equal the price difference between the two year-on-year caps with adjacent maturities, one with j years and the other with $j-1$ years:

$$C_{j,t} = C_t^{\mathrm{YoY}(j)} - C_t^{\mathrm{YoY}(j-1)},$$

because buying the former and selling the latter exactly replicates the caplet. By applying the same method to the caplet price, we can show that

$$\frac{\partial^2 C_{j,t}(K)}{\partial K^2} = D_t^{(j)}\, \tilde{q}_{t,t+j-1\to t+j}(K),$$

where $\tilde{q}_{t,t+j-1\to t+j}$ is the risk-adjusted probability density of future 1-year gross inflation $\Pi_{t+j-1\to t+j}$ between year $t+j-1$ and year $t+j$. The density of net inflation $\Pi_{t+j-1\to t+j} - 1$ is then given by $q_{t,t+j-1\to t+j}(x) = \tilde{q}_{t,t+j-1\to t+j}(1+x)$.

3.3 Estimation

We download the historical prices of zero-coupon and year-on-year inflation options from Bloomberg. The pricing data, collected at the daily frequency, span approximately 15 years from October 5, 2009 to August 30, 2024. Bloomberg offers various data sources/contributors. For the majority of our sample period (nearly 12 years starting in January 2013), we rely on

"BVOL," which provides option prices derived from smooth, arbitrage-free volatility surfaces constructed by Bloomberg. For the period prior to 2013 – when BVOL data are unavailable – we use the "CMP," Bloomberg's composite prices aggregated from multiple sources. There are three versions of CMP, corresponding to different time zones: CMPN (New York), CMPL (London), and CMPT (Tokyo). We use CMPN, although the prices are very similar across the three.[8] All prices are quoted in cents per $100 of notional value.

Inflation options are typically traded for a wide range of fixed maturities ranging from 1 to 30 years. Among them, we select contracts with the following three maturities: (i) 1-year zero-coupon options, (ii) forward-starting 1-year options maturing in two years (i.e., 2-year caplets/floorlets, which are replicated by buying 2-year and selling 1-year year-on-year contracts), and (iii) 5-year zero-coupon options. As described below, these contracts are chosen so as to construct the risk-adjusted distributions that are directly comparable to the survey-based distributions introduced in Section 2.

The strike rates observed in the data range from -3% to 7%, with a central value of 2%. At the lower and higher ends, the strike rates are spaced in 1% increments, specifically from -3% to -1% and from 5% to 7%. The middle range, between -1% and 5%, uses finer increments of 0.5%. Not surprisingly, not all of these strikes are actively traded. Options with the two most extreme strikes (-3% and 7%) are rarely observed, and their prices are virtually nonexistent. Additionally, prices for extremely deep-in-the-money options are generally unavailable for both caps and floors. Consequently, caps are typically observed with higher strike rates and floors with lower strike rates. For instance, BVOL consistently provides prices for 10 floors with strikes ranging from -2% to 3% and 10 caps with strikes ranging from 1% to 6%. The number of options available per day is similar but slightly higher for CMPN, with an average of around 26 (median of about 23) until 2013. Based on put-call parity, we convert floor prices into equivalent cap prices so we acquire the prices of caps for all available strikes, not just in the

[8]Unlike BVOL prices, Bloomberg composite prices occasionally become stale and remain unchanged for extended periods. To exclude these stale prices from our analysis, we drop observations if the price has not been updated over the past five trading days.

out-of-the-money or near-the-money region, but also in the in-the-money region.[9]

Equipped with a rich cross-section across the strike dimension, we estimate the implied density of inflation using the methodology developed by Aït-Sahalia and Duarte (2003). This process involves two main steps. First, we apply the algorithm of Dykstra (1983) to the observed prices, retaining arbitrage-free prices that satisfy the shape restrictions related to the slope and convexity of the option price curve. This filtering ensures, in particular, that butterfly spread strategies, whose payoffs mimic state prices, do not yield negative values. Second, we estimate the second derivative of the option price with respect to the strike price using locally polynomial kernel smoothing. Specifically, we employ the locally linear estimator (order of 1), which, according to Aït-Sahalia and Duarte (2003), produces a better fit than the traditional Nadaraya-Watson kernel estimator (order of 0).

To illustrate, Figure 5 presents the estimation results for three dates: June 13, 2018 (Panel A), June 10, 2020 (Panel B), and June 15, 2022 (Panel C), all of which were the fourth scheduled FOMC announcement days that year. For each date, we extract the risk-adjusted distribution of 1-year inflation $q_{t,t\to t+1}$ from $C_t^{\mathrm{ZC}(1)}$ (blue dashed line, left graph) and that of 1-year forward-starting 1-year inflation $q_{t,t+1\to t+2}$ from $C_{2,t} = C_t^{\mathrm{YoY}(2)} - C_t^{\mathrm{YoY}(1)}$ (yellow dotted line, left graph). To directly compare with the survey-based distribution from the SPF, we need the risk-adjusted distribution of inflation over the next calendar year, denoted by q_t^{ncy} (black solid line, left graph). Intuitively, q_t^{ncy} should lie somewhere between $q_{t,t\to t+1}$ and $q_{t,t+1\to t+2}$, considering that

$$q_t^{\mathrm{ncy}} = \begin{cases} q_{t,t+1\to t+2} & \text{if } t \text{ is the beginning of the current calendar year,} \\ q_{t,t\to t+1} & \text{if } t \text{ is the end of the current calendar year.} \end{cases}$$

For an arbitrary t, we approximate q_t^{ncy} by the weighted average of $q_{t,t\to t+1}$ and $q_{t,t+1\to t+2}$, with weights depending on the number of days included in each yearly period. Lastly, we calculate

[9]In the Internet Appendix, we compare inflation swap rates observed in the market with those implied by inflation option prices. The option-implied rates are derived from the strike at which the put and call price curves intersect; put-call parity implies that the inflation swap rate corresponds to the strike where the two prices are equal. We find that the option-implied rates closely match the market rates, demonstrating consistency in both levels and time series variation.

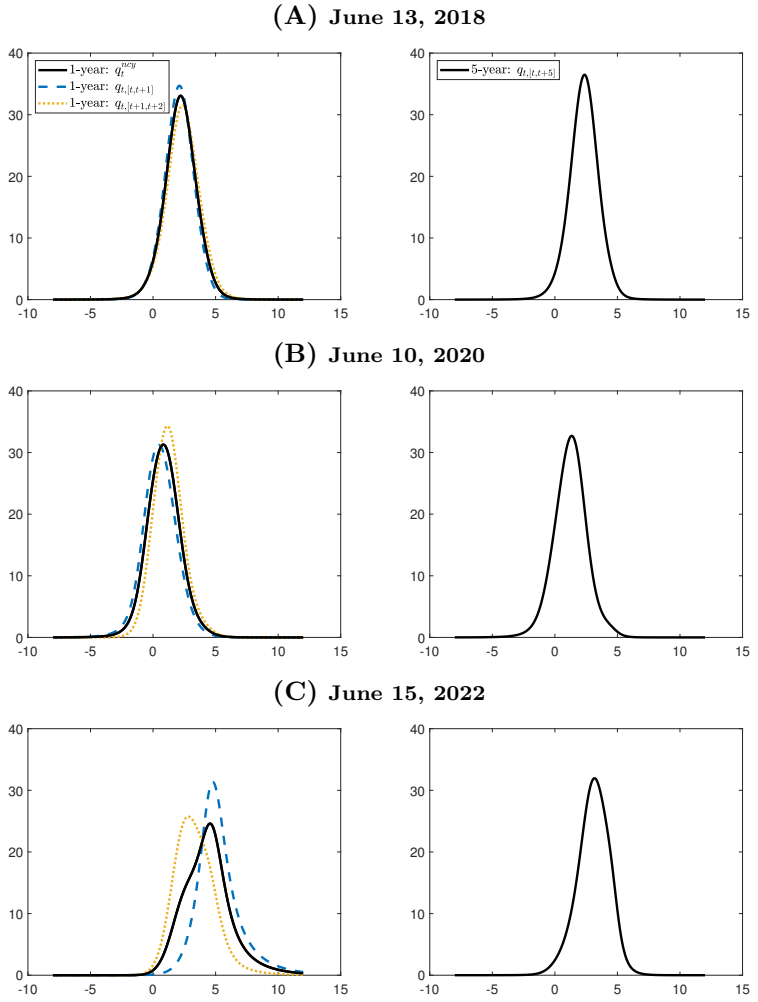


Figure 5: Risk-adjusted densities. This figure presents the risk-adjusted distributions of inflation for three dates: June 13, 2018 (Panel A), June 10, 2020 (Panel B), and June 15, 2022 (Panel C), all corresponding to the 4th scheduled FOMC announcement days of those years. For each date, the left graph shows the 1-year inflation distribution $q_{t,t\to t+1}$ (blue dashed line), the 1-year forward-starting 1-year inflation distribution $q_{t,t+1\to t+2}$ (yellow dotted line), and the inflation distribution over the next calendar year q_t^{ncy} (black solid line). The right graph shows the 5-year average inflation distribution $q_{t,t\to t+5}$ (black solid line).

the risk-adjusted distribution of 5-year average inflation $q_{t,t\to t+5}$ (black solid line, right graph), which serves as the counterpart to the survey-based distribution from the SPD.

The figure shows drastically different patterns across the different dates, reflecting the evolving economic landscape and inflation expectations over time. In June 2018 (Panel A), the 1-year inflation distribution $q_{t,t\to t+1}$ and the forward-starting 1-year inflation distribution $q_{t,t+1\to t+2}$, both of which are under the risk-adjusted measure, are closely aligned. This indicates a stable short-term outlook for inflation as perceived by investors. Naturally, the risk-adjusted distri-

bution of inflation over the next calendar year q_t^{ncy} is also positioned very close to these two distributions. Notably, both the 1-year inflation distributions and the 5-year average inflation distribution $q_{t,t\to t+5}$ remain symmetric with little skew.

As we move to June 2020 (Panel B), the inflation distributions shift substantially to the left, reflecting deflationary expectations in the wake of the COVID-19 pandemic. As investors grappled with an unpredictable economic environment, the divergence between $q_{t,t\to t+1}$ and $q_{t,t+1\to t+2}$ becomes pronounced. Reflecting heightened uncertainty and fears of deflation during this period, q_t^{ncy} also displays a wider spread and a larger probability mass in the deflationary region below 0%, relative to June 2018. Despite this, the 5-year average inflation distribution $q_{t,t\to t+5}$ remains fairly symmetric, albeit slightly shifting to the left, suggesting that long-term inflation expectations were still anchored.

However, we observe the completely opposite pattern in June 2022 (Panel C). The inflation distributions have shifted to the right, reflecting the steady rise in inflation as economic conditions changed dramatically post-2020; the year-on-year inflation rate peaked in June 2022 at about 9%. The 1-year inflation distribution $q_{t,t\to t+1}$ is now located further to the right than the forward-starting 1-year inflation distribution $q_{t,t+1\to t+2}$, as short-term inflation expectations surged alongside realized inflation. Similar to June 2020, q_t^{ncy} has a large spread but now assigns the majority of its probability mass to regions with high inflation. While investors' expectations about short-term inflation had significantly increased, the 5-year average inflation distribution $q_{t,t\to t+5}$ shows a moderate shift and remains relatively symmetric. This is consistent with the belief that, despite the immediate inflationary pressures, inflation would eventually stabilize over a longer horizon.

4 Preferences for different inflation ranges

In this section, we combine the survey-based probability distribution from Section 2 with the risk-adjusted probability distribution from Section 3 to explore what they reveal about investors' preferences for different ranges of inflation.

To this end, we assume that marginal investors in the inflation option market hold beliefs similar to those of survey respondents. While strong, this assumption is not uncommon in the literature. Prior studies that use survey forecasts as proxies for physical expectations and incorporate them into model estimation implicitly make a similar assumption (see, e.g., Haubrich, Pennacchi, and Ritchken, 2012). To minimize potential frictions, we intentionally exclude household surveys and focus on surveys of professional forecasters and primary dealers. Respondents to the SPF are professional economists, analysts, and forecasters, many of whom work at major banks and investment firms. The SPD also targets primary dealers – mostly large banks and their affiliates. Hence, it is not unreasonable to assume that the beliefs of these survey respondents are closely aligned with those of marginal investors trading inflation options. In fact, we show in the Internet Appendix that the average survey-based distribution from the SPD is fairly similar to the historical distribution constructed from realized inflation outcomes over the same sample period. The assumption can always be revisited if estimation results appear unreasonable, but, as we demonstrate later, that is not the case.

Building on this premise, we can establish a direct connection between the two measures $\mathbb{P}$ and $\mathbb{Q}_{t+T}$ through investors' marginal utility. Since $\mathbb{Q}_{t+T}$ uses a zero-coupon bond $D_t^{(T)} = \mathbb{E}_t\left[M^{\$}_{t\to t+T} \times 1\right]$ as the numéraire, the Radon-Nikodym derivative of $\mathbb{Q}_{t+T}$ with respect to $\mathbb{P}$ is expressed by

$$\left.\frac{d\mathbb{Q}_{t+T}}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \frac{M^{\$}_{t\to t+T}}{D_t^{(T)}} = \frac{M^{\$}_{t\to t+T}}{\mathbb{E}_t\left[M^{\$}_{t\to t+T}\right]}, \tag{4}$$

where $M^{\$}_{t\to t+T}$ denotes the (nominal) pricing kernel, or equivalently, marginal utility. The key intuition from equation (4) is that $\mathbb{Q}_{t+T}$ is obtained by re-weighting $\mathbb{P}$ based on future marginal utility values. Let ω represent an arbitrary event realized at time $t+T$, i.e., $\omega \in \mathcal{F}_{t+T}$. If the event constitutes a "good" economic state, in terms of lower-than-expected marginal utility ($M^{\$}_{t\to t+T}(\omega) < \mathbb{E}_t\left[M^{\$}_{t\to t+T}\right]$), the Radon-Nikodym derivative in equation (4) assigns a lower probability to this event under $\mathbb{Q}_{t+T}$ than $\mathbb{P}$. In contrast, if this event is perceived as a "bad" economic state with higher-than-expected marginal utility ($M^{\$}_{t\to t+T}(\omega) > \mathbb{E}_t\left[M^{\$}_{t\to t+T}\right]$), the

event probability under $\mathbb{Q}_{t+T}$ is determined to be higher than that under $\mathbb{P}$. In other words, a measure change from $\mathbb{P}$ to $\mathbb{Q}_{t+T}$ tilts the probabilities such that good events are under-weighted and bad events are over-weighted. This implies that the probability ratio between the two measures contains crucial information about how investors perceive a certain event, whether it be good or bad.

This interpretation still holds when we project both sides of equation (4) on future realized inflation. If the risk-adjusted probability is lower (higher) than the subjective probability for a certain realization of inflation, it indicates that investors regard it as a good (bad) economic outcome that is associated with lower (higher) marginal utility on average. As such, we can define and determine *good inflation* versus *bad inflation*, as perceived by investors, based on the probability ratio between the two measures.

4.1 Implied conditional pricing kernel

We begin by examining the implied pricing kernel, projected on inflation over the next calendar year. To take the ratio between the risk-adjusted and physical probability distributions, we first need to ensure they are directly comparable. In the case of the risk-adjusted distribution, inflation options allow us to nonparametrically estimate the entire probability density function over a continuous domain, as outlined in Section 3. Using this density, we calculate the probabilities of inflation falling into 20 distinct inflation ranges, from -3% to 7%, with 0.5% intervals. By averaging these risk-adjusted probabilities within each quarter, we obtain a quarterly measure that aligns with the timing and frequency of the SPF survey. Panel A of Figure 6 presents the unconditional average of the resulting quarterly risk-adjusted inflation distributions over our sample period.

The survey-based distribution also requires additional processing. For each quarter, we convert the bucket probabilities provided by the survey into a cumulative distribution function and interpolate these values using a smooth spline function. For both ends of the tail, we adopt a decaying tail function, fitting it to match the probability of the last unbounded bucket. The resulting inter/extrapolated distribution function is defined over a continuous domain and

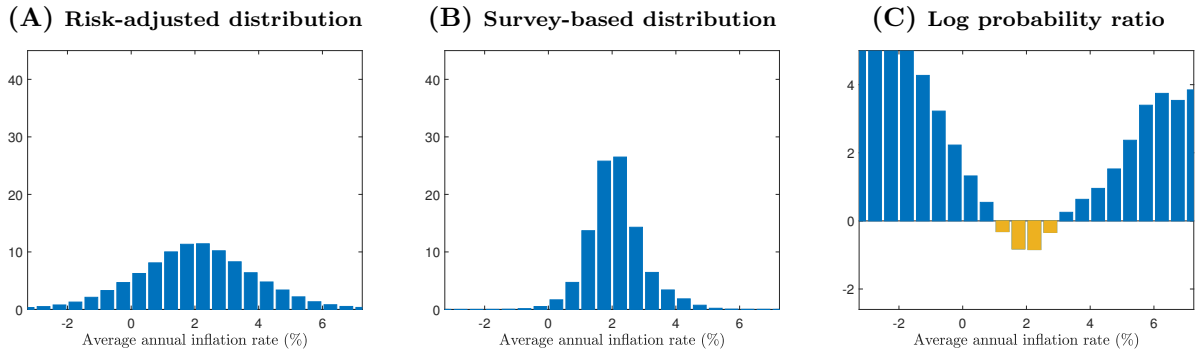


Figure 6: Unconditional inflation densities (next calendar year). This figure presents the unconditional inflation densities. Panel A provides the risk-adjusted distribution $\mathbb{Q}$, which is derived from options prices. Panel B shows the survey-based distribution $\mathbb{P}$ of inflation derived from the Survey of Professional Forecasters (SPF). Panel C displays the log of the probability ratio between the risk-adjusted and survey-based distributions $\log(\mathbb{Q}/\mathbb{P})$, with yellow bars indicating good inflation and the blue bars indicating bad inflation. The distributions illustrate the probabilities of different inflation outcomes over the next calendar year. The sample period is from 2009Q4 to 2024Q2. All values are in percentages.

reproduces the same bucket probabilities as the survey data. This fitting exercise serves two key purposes. First, it allows us to study the survey-based distribution over finer intervals, which is particularly useful in the earlier part of the sample where inflation buckets are wider. Second, it transforms the two extreme unbounded buckets into a series of bounded buckets with 0.5% intervals, making them directly comparable to the risk-adjusted distribution.[10] The unconditional average of the survey-based distributions is displayed in Panel B of Figure 6.

Panel C depicts the log ratio of risk-adjusted to survey-based probabilities, which corresponds to the log of the probability distortion factor in equation (4). Thus, a given inflation range can be interpreted as "bad" if the log probability ratio is positive (blue bars), or "good" if the log probability ratio is negative (yellow bars). As shown in Panel C, the log probability ratio exhibits an unconditional U-shaped pattern, with positive values (indicating bad inflation) in the tails and negative values (indicating good inflation) in the middle ranges. This pattern suggests that investors tend to exhibit aversion to both high and low inflation environments, consistent with the findings of Kitsul and Wright (2013). In fact, our analysis suggests that

[10] Another potential source of discrepancy is that the SPF survey is based on the GDP price index, whereas inflation options reference the CPI. Although the two inflation measures are not identical, they are highly correlated over the sample. As shown in the Internet Appendix, a scatter plot of the two series reveals that most observations lie on or near the 45-degree line, suggesting that the two measures move closely together throughout the period.

investors favor inflation outcomes around 2% and that the good inflation region falls between 1% and 3%, unconditionally over our sample period.

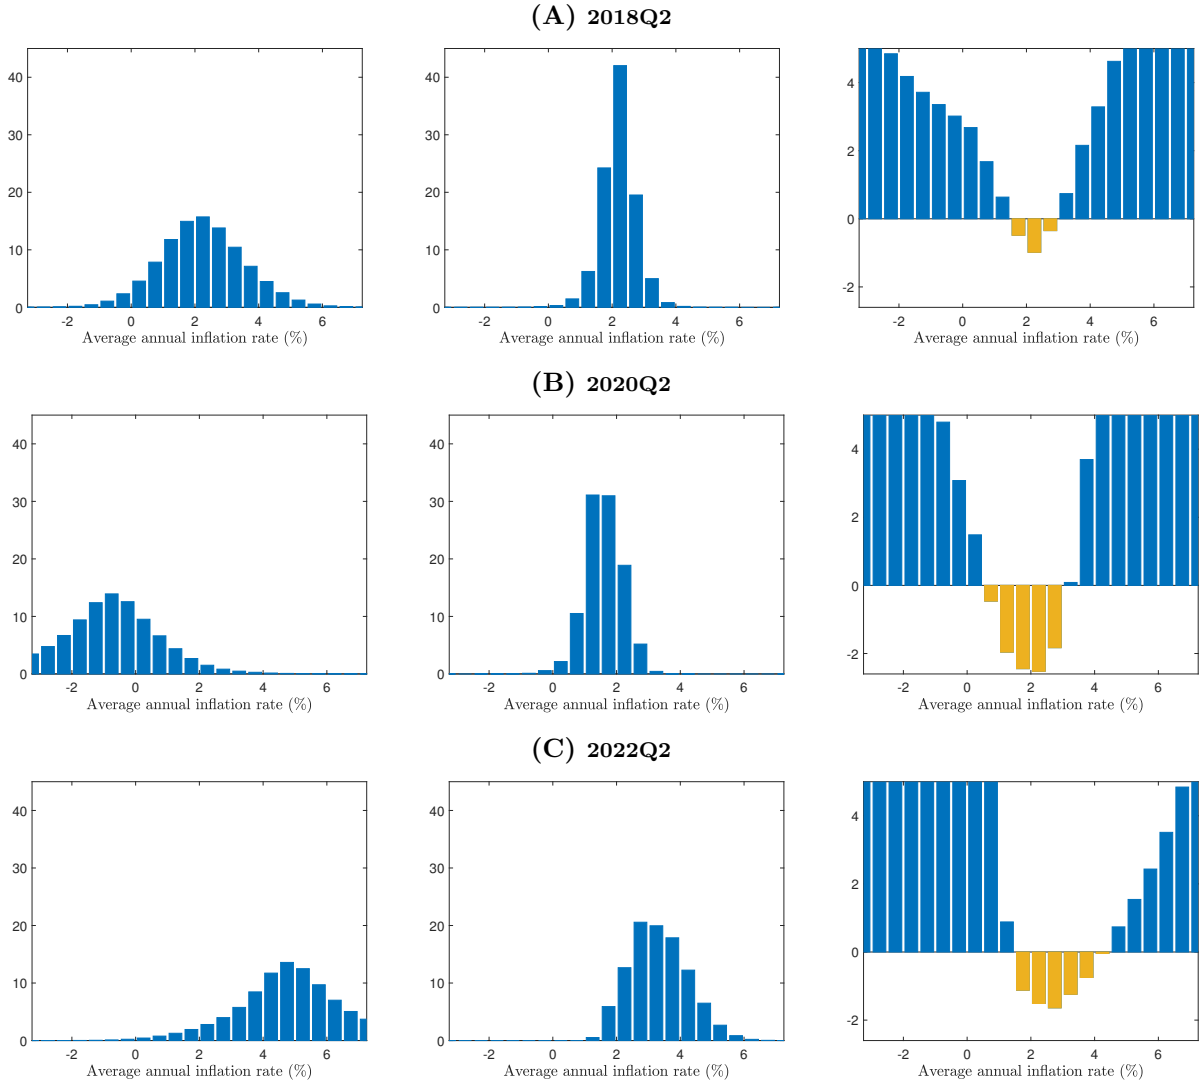


Figure 7: Conditional inflation densities (next calendar year). This figure presents the conditional inflation densities for the second quarter of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). Each panel contains three graphs: the risk-adjusted distribution $\mathbb{Q}$ (left), the survey-based distribution $\mathbb{P}$ (center), and the log probability ratio $\log(\mathbb{Q}/\mathbb{P})$ (right). In Panel C, the yellow bars indicate good inflation, while the blue bars indicate bad inflation. The distributions illustrate the probabilities of different inflation outcomes over the next calendar year. The sample period is from 2009Q4 to 2024Q2. All values are in percentages.

As discussed in Sections 2 and 3, there is significant time variation in both the survey-based and risk-adjusted probabilities. This time variation also affects the probability ratio, reflecting shifts in investors' preferences toward inflation over time. To illustrate, Figure 7 presents the risk-adjusted probabilities (left graphs), survey-based probabilities (center graphs), and the

log ratio between them (right graphs) for three different quarters: the second quarter of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). Focusing on the log probability ratio (right graphs), during relatively normal times like the second quarter of 2018, the good inflation region is relatively small and tightly centered around 2%. In the second quarter of 2020, when deflation risk became a greater concern, the good inflation region substantially expands, especially toward the lower end of the spectrum; at such times, even inflation rates of 0.5-1% are considered good economic outcomes. In contrast, during the inflationary period of 2022, the good inflation region shifts more to the right, extending to the 4-4.5% interval. Despite these variations across the three periods, it is noteworthy that the U-shaped pattern in the probability ratios persists consistently.

The time variation in the log probability ratio provides further insight into how investors' preferences for inflation outcomes evolve. Figure 8 illustrates the time-varying nature of the good inflation range over the sample period. For each quarter, the yellow region represents future inflation values at which the risk-adjusted probabilities fall below their survey-based counterparts, indicating favorable economic outcomes perceived by investors (i.e., good inflation region). Conversely, the blue region highlights inflation values where risk-adjusted probabilities exceed survey-based probabilities, signaling outcomes that investors view negatively (i.e., bad inflation region).

This figure clearly shows that the good inflation range is not fixed throughout the sample period. During periods of low inflation, such as the aftermath of the 2008 financial crisis or the onset of the COVID-19 crisis, the good inflation range tends to shift downward. In contrast, during the inflation surge of 2022, it shifts significantly upward, to the point where the 2% inflation rate – historically viewed as an ideal target by central banks – falls outside this range. As inflation begins to decline toward the end of the sample period in 2023, the good inflation range shifts downward once again, demonstrating how investors' perceptions adapt to the economic environment.

How can we interpret these patterns? One way to rationalize the observed dynamics is by considering what future inflation realizations signal about future economic growth. For instance,

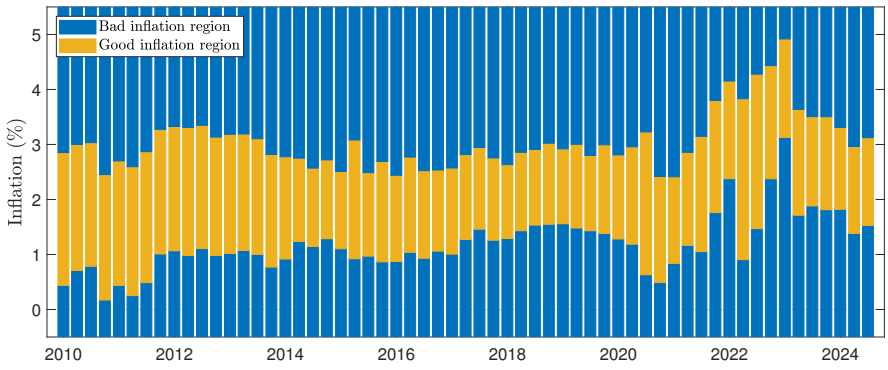


Figure 8: Good and bad inflation regions (next calendar year). This figure illustrates the time series of the good (yellow bars) and bad inflation (blue bars) regions. The distributions illustrate the probabilities of different inflation outcomes over the next calendar year. The sample period is from 2009Q4 to 2024Q2. All values are in percentages.

in the inflationary environment of 2022, high inflation over the next year is clearly perceived as bad news, suggesting a significant risk of stagflation. However, low inflation projections during this period could have mixed interpretations. On the one hand, they might indicate a return to the normalcy seen in 2018; on the other hand, they could reflect a potential scenario with *hard landing*, where the Fed's policies drive the economy into recession, resulting in low inflation rates. Simply put, a sharp decline from the high inflation rates of summer 2022 to lower levels (e.g., 1%) is unlikely to be interpreted as positive news, as it is more likely associated with a hard-landing scenario. This reasoning helps explain why investors' preferred range of inflation varies over time. It is not simply the level of inflation that matters, but the changes and the economic signals they convey, which play a significant role in shaping investors' preferences for good and bad inflation ranges.

4.2 Inflation risk premium and its decomposition

Another way to gauge the variation in investors' preferences toward inflation is by analyzing the inflation risk premium. This premium is defined as the difference between the risk-adjusted and survey-based expected inflation rates:

$$ip_{t,t\to t+T} = \mathbb{E}_t^{\mathbb{Q}_{t+T}}\left[\sqrt[T]{\Pi_{t\to t+T}}\right] - \mathbb{E}_t\left[\sqrt[T]{\Pi_{t\to t+T}}\right]. \tag{5}$$

Using the conditional distributions of inflation under the risk-adjusted and survey-based probability measures, we can calculate the inflation risk premium for the next calendar year.[11] Figure 9 displays the variation in the December-to-December inflation risk premium over the sample period. Note that a positive premium suggests that investors demand compensation for holding nominal securities bearing inflation risk, as they view high inflation as unfavorable, perceiving it to be correlated with low growth. In contrast, a negative premium indicates that investors view high inflation favorably, associating it with high growth on average.

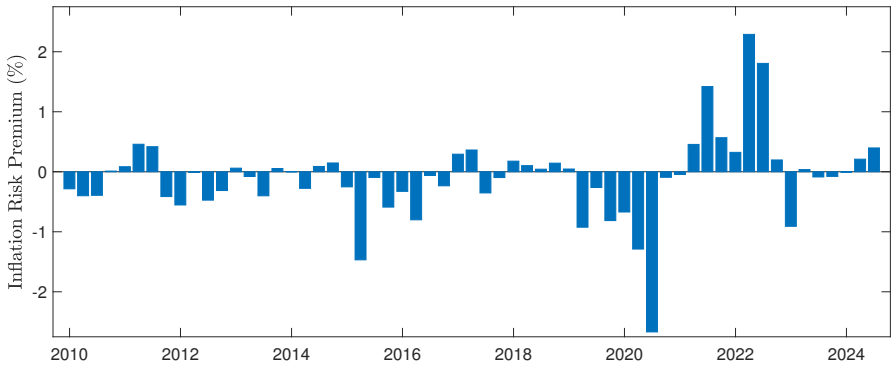


Figure 9: Time series of the inflation risk premium (next calendar year). This figure presents the time series of the inflation risk premium, calculated as the difference between risk-adjusted and survey-based expected inflation. The distributions illustrate the probabilities of different inflation outcomes over the next calendar year. The sample period is from 2009Q4 to 2024Q2. All values are in percentages.

While the inflation risk premium fluctuates significantly over time, Figure 9 illustrates that its level is generally negative until 2021. The most significant negative spike (-2.67%) occurs in the second quarter of 2020 during the COVID-19 pandemic, when heightened uncertainty and a substantial drop in consumer spending led to lower growth expectations alongside lower inflation expectations. Beginning in 2021, the inflation risk premium abruptly turns positive, marking a significant shift in investor sentiment. During this post-pandemic period, the U.S. experienced its highest inflation rates in over 40 years, prompting fears of stagflation reminiscent

[11] As a sanity check, we regress breakeven inflation at various maturities on survey-based expected inflation and our estimated inflation risk premium. The breakeven inflation rate for a given maturity is defined as the yield difference between nominal Treasury securities and TIPS of the same maturity, reflecting the risk-adjusted expectation of future inflation over that horizon. Hence, it can be understood as the sum of expected inflation and the inflation risk premium. We find statistically significant slope coefficients and very high R^2 values from the regressions, especially when the breakeven inflation rate and the two explanatory variables are all measured over the same 5-year horizon. The results remain robust when the variables are first-differenced to account for the potential presence of unit roots.

of the 1970s. The largest inflation risk premium in the sample, 2.29%, is recorded in the first quarter of 2022, as the Fed began to implement aggressive rate hikes and tightening policies aimed at combating inflation.

These patterns reflect the changing cyclicality of inflation risk perceived by investors. The negative inflation risk premium prior to 2021 indicates that inflation risk was largely seen as procyclical (high inflation, high growth) during this period, implying that nominal securities like Treasury bonds served as hedging assets. The positive inflation risk premium observed post-2021, on the other hand, suggests that inflation risk was regarded as countercyclical (high inflation, low growth), making nominal securities risky like stocks. This shift in the sign of the inflation risk premium aligns closely with the time-varying patterns of the stock-bond correlation documented in the literature (e.g., Campbell, Sunderam, and Viceira, 2017). In particular, Campbell, Pflueger, and Viceira (2020) find a persistently negative correlation between Treasury bond returns and stock returns in the first two decades of the 21st century, which is indicative of procyclical inflation: during economic downturns, lower stock returns come with lower inflation, which raises real bond prices, resulting in the negative stock-bond correlation. More recently, Lombardi and Sushko (2023) document that this correlation turned positive in mid-2021, almost exactly when the inflation risk premium becomes positive in Figure 9; both quantities clearly point to changing investor perceptions toward countercyclical inflation risk in the new post-pandemic era.

As the next exercise, we decompose the inflation risk premium into compensation for different inflation outcomes. For this analysis, we adopt the methodology of Beason and Schreindorfer (2022), which uses options and returns data to decompose the equity risk premium over different parts of the return state space. Applying this to our context, the inflation risk premium associated with a range $[\pi_L, \pi_U]$ can be defined as:

$$ip_{t,t\to t+T}^{[\pi_L,\pi_U]} = \int_{\pi_L}^{\pi_U} x \left[\mathbb{Q}_{t+T} \left(\sqrt[T]{\Pi_{t\to t+T}} - 1 = x \right) - \mathbb{P} \left(\sqrt[T]{\Pi_{t\to t+T}} - 1 = x \right) \right] dx. \tag{6}$$

Comparing equations (5) and (6) makes it clear that $ip_{t,t\to t+T}^{[-\infty,\infty]} = ip_{t,t\to t+T}$. Unlike Beason

and Schreindorfer (2022), who focus on the unconditional equity premium due to the difficulty of accurately estimating conditional distributions for stock returns under the physical measure, we can analyze both the conditional and unconditional inflation risk premiums using survey-based inflation densities.

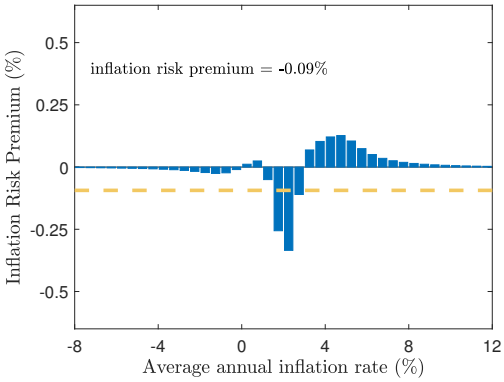


Figure 10: Decomposition of the unconditional inflation risk premium (next calendar year). This figure presents the unconditional decomposition of the inflation risk premium. The yellow dashed line represents the average inflation risk premium, and the blue bars represent the contribution of different inflation ranges to the overall premium. The distributions illustrate the probabilities of different inflation outcomes over the next calendar year. All values are in percentages.

Figure 10 presents the decomposition results for the unconditional inflation risk premium. According to equation (6), the inflation risk premium for a given range depends on the likelihood of the inflation outcomes within that range (i.e., the quantity of risk) as well as the relative difference between risk-adjusted and survey-based probabilities (i.e., the price of risk). This becomes clear when we rearrange the right-hand side of equation (6) as follows:

$$ip_{t,t\to t+T}^{[\pi_L,\pi_U]} = \int_{\pi_L}^{\pi_U} \underbrace{x}_{\text{Inflation outcome}} \times \underbrace{\mathbb{P}\left(\sqrt[T]{\Pi_{t\to t+T}} - 1 = x\right)}_{\text{Quantity of of inflation risk}} \times \underbrace{\left[\frac{\mathbb{Q}_{t+T}}{\mathbb{P}}\left(\sqrt[T]{\Pi_{t\to t+T}} - 1 = x\right) - 1\right]}_{\text{Price of inflation risk}} dx.$$

Therefore, as we move toward the less likely tail ranges (with a smaller quantity of inflation risk), the premium contribution diminishes as shown in Figure 10. This equation also reveals that the sign of each range's contribution is governed by two factors: (i) the sign of the inflation outcomes $x \in [\pi_L, \pi_U]$ and (ii) the sign of the price of inflation risk, which is determined by whether the range falls within the good or bad inflation region. Recall that, by definition,

$[\mathbb{Q}_{t+T}/\mathbb{P} - 1]$ is negative in the good inflation region and positive in the bad inflation region.

As a result, the contributions to the total inflation risk premium can vary substantially not only in magnitude but also in sign across different ranges. In Figure 10, the contributions are positive in the unfavorably high inflation range (both (i) and (ii) positive) but turn negative in the favorable range around 2% where (i) is positive and (ii) is negative. In the lower inflation range beyond the good region, the contributions are again positive (both (i) and (ii) positive), but they become negative in the deflation region because inflation outcomes themselves are now negative; that is, (i) is negative and (ii) is positive. This implies that despite the significance of premiums associated with certain ranges, the total inflation risk premium can remain small due to the offsetting of positive and negative contributions. In fact, the figure shows that different inflation ranges are associated with premiums ranging from -33 bp for $[2\%, 2.5\%)$ to 13 bp for $[4.5\%, 5\%)$, but the total inflation risk premium is only -9 bp. Given the U-shaped pattern observed in the projection of the pricing kernel, the inflation risk premium – integrated over the entire inflation range – can mask important insights into the pricing of inflation risk.

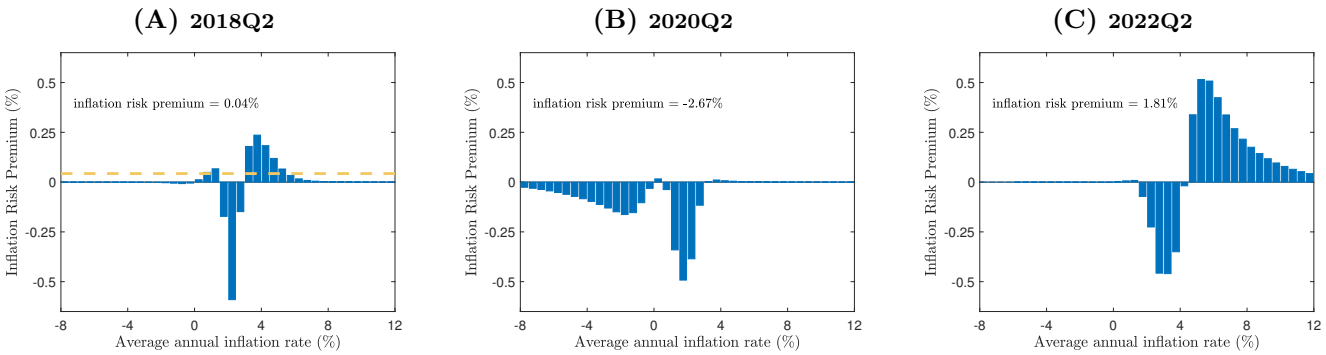


Figure 11: Decomposition of the conditional inflation risk premium (next calendar year). This figure presents the conditional decomposition of the inflation risk premium for the second quarter of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). The yellow dashed line represents the average inflation risk premium, and the blue bars represent the contribution of different inflation ranges to the overall premium. The distributions illustrate the probabilities of different inflation outcomes over the next calendar year. All values are in percentages.

Figure 11 shows the decomposition of the conditional inflation risk premium for three different quarters: the second quarter of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). In 2018, the pattern of contributions across inflation ranges closely resembles the unconditional

results, with the deflation region contributing virtually nothing to the overall premium. By 2020, however, this region with negative inflation outcomes generates a substantially negative premium, reflecting heightened deflation concerns in the aftermath of the COVID-19 outbreak. High inflation outcomes contribute little due to their low probabilities. In 2022, the decomposition shifts again in response to investors' growing concerns about rising inflation: substantial premiums are associated with inflation outcomes above 4%, while contributions from low inflation outcomes below 1% are essentially nonexistent.

4.3 Pricing of long-term inflation risk

Next, we turn to the implied pricing kernel, projected on inflation over a longer term. Figure 12 presents the unconditional inflation densities over the next five years. Panel A shows the risk-adjusted distribution implied by 5-year inflation options, Panel B displays the survey-based distribution of inflation derived from the SPD, and Panel C illustrates the log probability ratio between the two distributions. Compared to Figure 6, which shows the results for the next calendar year, Panel C presents a qualitatively similar but a narrower good inflation region (below zero; yellow bars).

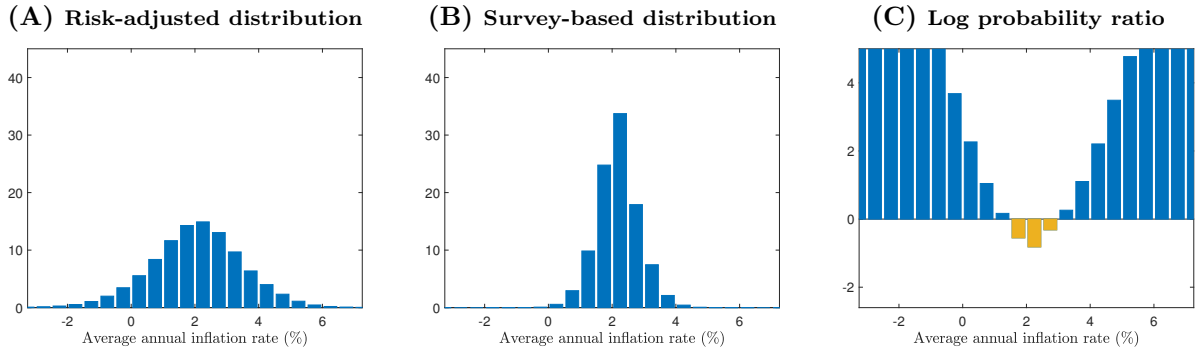


Figure 12: Unconditional inflation densities (5-year horizon). This figure presents the unconditional inflation densities. Panel A provides the risk-adjusted distribution $\mathbb{Q}$, which is derived from options prices. Panel B shows the survey-based distribution $\mathbb{P}$ of expected inflation derived from the Survey of Primary Dealers (SPD). Panel C displays the log of the probability ratio between the risk-adjusted and survey-based distributions $\log(\mathbb{Q}/\mathbb{P})$, with yellow bars indicating good inflation and the blue bars indicating bad inflation. The distributions illustrate the probabilities of different inflation outcomes over the next five years. The sample period is from December 2014 to June 2024. All values are in percentages.

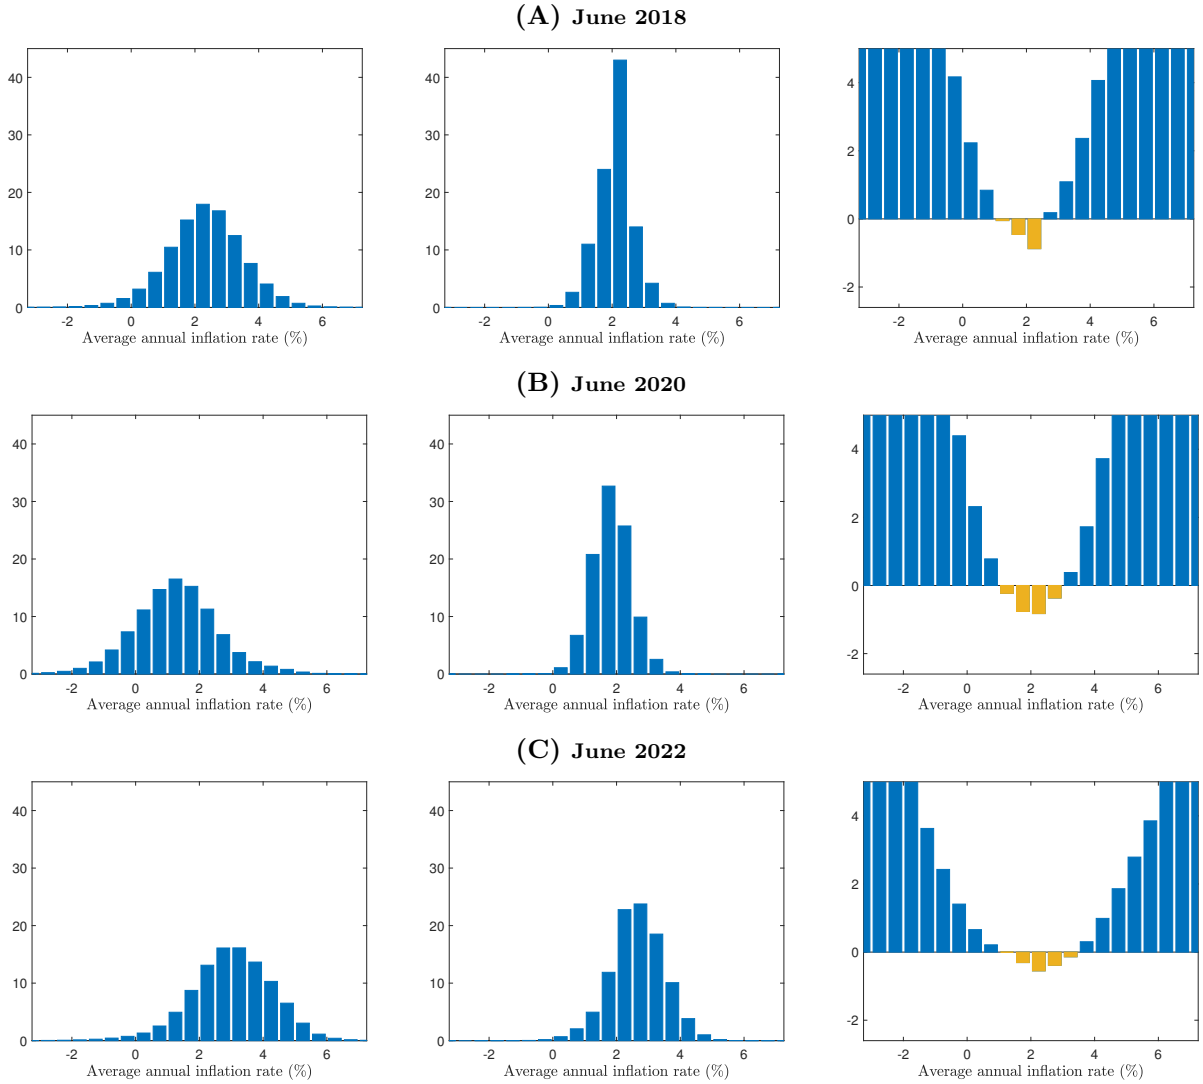


Figure 13: Conditional inflation densities (5-year horizon). This figure presents the conditional inflation densities collected before the fourth FOMC meeting for three different years: June 2018 (Panel A), June 2020 (Panel B), and June 2022 (Panel C). Each panel contains three graphs: the risk-adjusted distribution $\mathbb{Q}$ (left), the survey-based distribution $\mathbb{P}$ (center), and the log probability ratio $\log(\mathbb{Q}/\mathbb{P})$ (right). In Panel C, the yellow bars indicate good inflation, while the blue bars indicate bad inflation. The distributions illustrate the probabilities of different inflation outcomes over the next five years. The sample period is from December 2014 to June 2024. All values are in percentages.

To understand shifts in investors' preferences at different points in time, Figure 13 illustrates the conditional inflation densities collected before the fourth FOMC meeting in June of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). Each panel contains three graphs: the risk-adjusted probabilities (left), survey-based probabilities (center), and the log probability ratio (right). In 2018, the good inflation region (yellow bars) is within a narrow range with little

deviation from 2%. This range begins to deviate more and becomes wider in 2020, followed by a slight rightward shift in 2022. Overall, the good inflation regions for the 5-year horizon are much tighter and stable, compared to those for the next calendar year. This makes sense: when long-term inflation expectations are anchored, investors' preferred range of inflation should vary less in the long run.

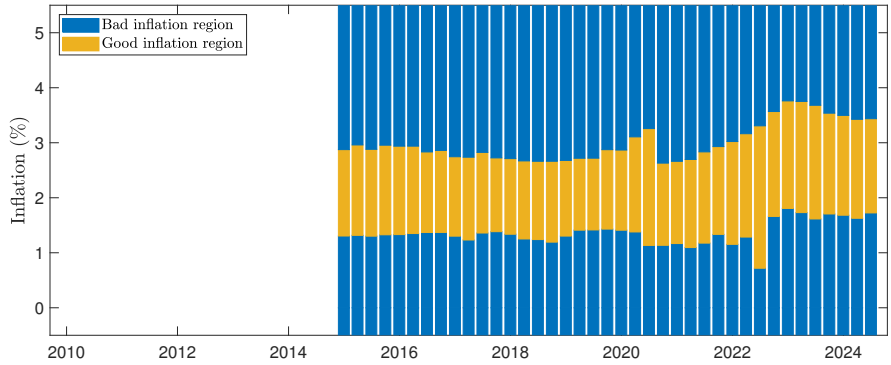


Figure 14: Good and bad inflation regions (5-year horizon). This figure illustrates the time series of the good (yellow bars) and bad inflation (blue bars) regions. The distributions illustrate the probabilities of different inflation outcomes over the next five years. The sample period is from December 2014 to June 2024. All values are in percentages.

This point is further highlighted in Figure 14, which depicts the time-varying ranges of good and bad inflation over the next five years, with each bar representing a quarter. The yellow-shaded portions of the bars indicate the good inflation region, whereas the blue-shaded portions indicate the bad inflation region. As can be seen in the figure, the good inflation region remains relatively stable throughout the sample period, though it shifts upward and slightly widens around 2022. Unlike in Figure 8, the Fed's 2% inflation target is always within the good region when the 5-year horizon is considered.

Figure 15 turns to the time series of the 5-year inflation risk premium. The inflation risk premium fluctuates between positive and negative values throughout our sample period, which begins in December 2014. Negative inflation risk premiums dominate from December 2014 through 2020, especially during crisis periods such as the COVID-19 pandemic, when the premium drops sharply. The inflation risk premium turns positive in 2021 and remains so through 2022. When juxtaposed with the time series in Figure 9, we see that while the overall patterns

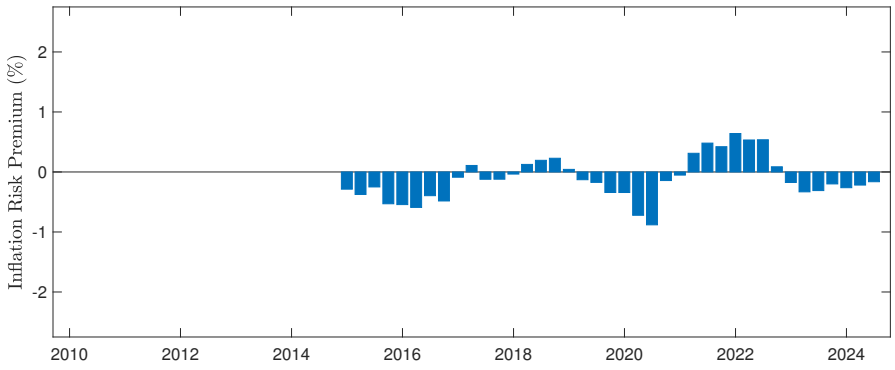


Figure 15: Time series of the inflation risk premium (5-year horizon). This figure presents the time series of the inflation risk premium, calculated as the difference between risk-adjusted and survey-based expected inflation. The distributions illustrate the probabilities of different inflation outcomes over the next five years. The sample period is from December 2014 to June 2024. All values are in percentages.

are similar, the time series for the 5-year horizon is smoother, exhibiting smaller magnitude and less variation. The response to major events is less pronounced, as the longer forecast horizon tempers the immediate reactions observed in the shorter horizon. For example, while the inflation risk premium for the 1-year horizon sharply spikes to its highest level in 2022, the one for the 5-year horizon shows a more sustained and less dramatic increase. Toward the end of the sample period in 2024, the 5-year inflation risk premium returns to negative territory.

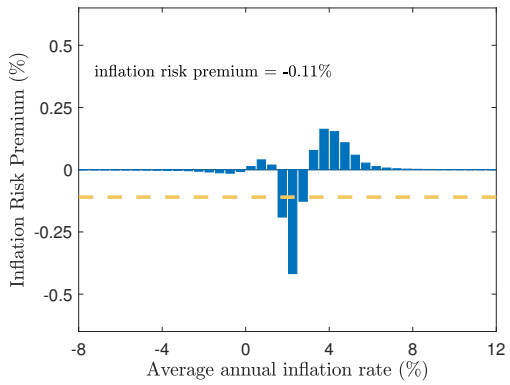


Figure 16: Decomposition of the unconditional inflation risk premium (5-year horizon). This figure presents the unconditional decomposition of the inflation risk premium. The yellow dashed line represents the average inflation risk premium, and the blue bars represent the contribution of different inflation ranges to the overall premium. The distributions illustrate the probabilities of different inflation outcomes over the next five years. All values are in percentages.

Figure 16 depicts the decomposition of the 5-year unconditional inflation risk premium. The unconditional inflation risk premium is small at -11 bp. Moderate inflation ranges near

2% result in the largest negative contributions to the risk premium, which overpower positive contributions from higher inflation ranges. Note that the contributions from low inflation ranges are very small: although investors dislike future deflation (as shown in Panel C of Figure 12), the likelihood of such an outcome is too small to generate a meaningful premium. Finally, Figure 17 provides the decomposition of the conditional inflation risk premium for the next five years, illustrating how different inflation ranges contribute to the overall premium during the second quarter of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). The patterns are similar to those observed for the next calendar year (Figure 11). Nevertheless, the patterns are much smoother for the 5-year horizon and tails contribute little, reflecting more stable inflation expectations in the long run.

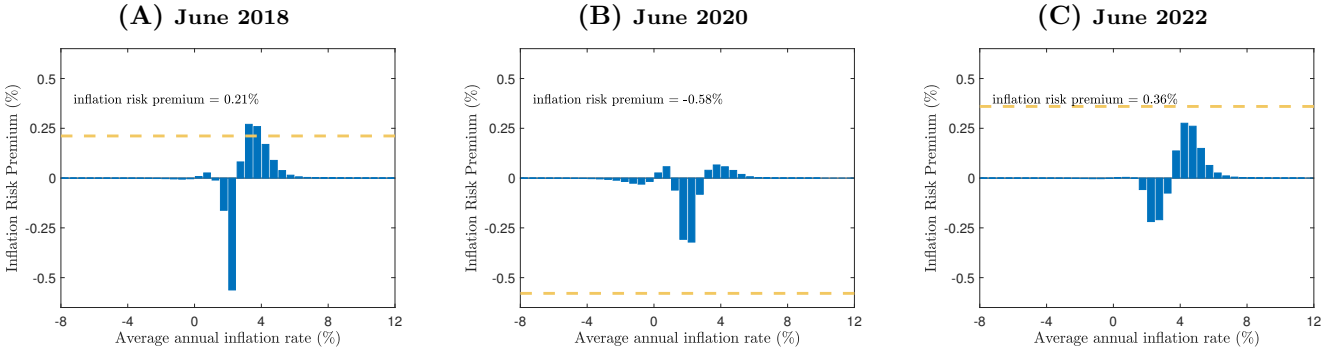


Figure 17: Decomposition of the conditional inflation risk premium (5-year horizon). This figure presents the conditional decomposition of the inflation risk premium for the second quarter of 2018 (Panel A), 2020 (Panel B), and 2022 (Panel C). The yellow dashed line represents the average inflation risk premium, and the blue bars represent the contribution of different inflation ranges to the overall premium. The distributions illustrate the probabilities of different inflation outcomes over the next five years. All values are in percentages.

5 Economic model with learning

We rationalize the empirical patterns we document using a model in which investors cannot precisely observe the true state of the economy. Due to this information friction, investors form their beliefs about two bad economic states: an inflationary recession and a deflationary recession. The fear of these states generates a conditional pricing kernel that is U-shaped when projected on future inflation. Investors' preferences for different inflation ranges vary over time

as they learn and update their beliefs, which aligns with our empirical observations.

5.1 Model setup and solution

We consider an endowment economy with complete markets, where investors have recursive preferences as described by Epstein and Zin (1989) and Weil (1989). Time is discrete, and a unit period between time t and time $t+1$ now represents a month. The stochastic discount factor or the real pricing kernel M_{t+1} is given by

$$M_{t+1} = \exp\left(\theta \log \delta - \frac{\theta}{\psi}\Delta c_{t+1} + (\theta - 1) r_{c,t+1}\right), \tag{7}$$

where $\Delta c_{t+1} \equiv \log\left(\frac{C_{t+1}}{C_t}\right)$ denotes log aggregate consumption growth and $r_{c,t+1}$ represents the log return on a security that delivers aggregate consumption as dividends (commonly referred to as the consumption claim). The parameters δ and ψ correspond to the rate of time preference and the elasticity of intertemporal substitution, respectively. The coefficient $\theta = \frac{1-\gamma}{1-1/\psi}$ reflects investors' preferences for the timing of uncertainty resolution, where γ represents relative risk aversion. Since the Euler equation must hold for real prices of nominal assets (see, e.g., Piazzesi and Schneider, 2006; Wachter, 2006; Bansal and Shaliastovich, 2013), the nominal pricing kernel is specified by

$$M^{\$}_{t+1} = e^{-\pi_{t+1}} M_{t+1}, \tag{8}$$

where $\pi_{t+1} \equiv \log\left(\frac{CPI_{t+1}}{CPI_t}\right)$ denotes the log inflation rate.

5.1.1 Consumption and inflation dynamics and investors' learning

We assume that the true state of the economy follows a Markov-switching process with three potential regimes: $s_{t+1} \in \{n, d, i\}$. State n represents normal or favorable economic conditions, characterized by relatively high consumption growth and moderate inflation. In contrast, the other two regimes capture bad economic conditions with low consumption growth. State d

reflects a deflationary recession where low growth comes with very low and potentially negative inflation, exemplified by the Great Depression in the 1930s. State i corresponds to an inflationary recession where low growth coincides with a sharp increase in inflation, as seen during the stagflation of the 1970s. The transition matrix $\mathcal{P}$ governs the dynamics of the regime-switching process:

$$\mathcal{P} = \begin{bmatrix} p_{nn} & p_{nd} & p_{ni} \\ p_{dn} & p_{dd} & p_{di} \\ p_{in} & p_{id} & p_{ii} \end{bmatrix},$$

where $p_{ss'} = P(s_{t+1} = s' | s_t = s)$ represents the transition probability from state s to state s'.

We model the different consumption and inflation dynamics across the three states by assuming their conditional means are state-dependent:

$$\Delta c_{t+1} = \mu^c_{s_{t+1}} + \sigma^c \epsilon^c_{t+1}, \tag{9}$$

$$\pi_{t+1} = (1 - \varphi^\pi)\mu^\pi_{s_{t+1}} + \varphi^\pi \pi_t + \sigma^\pi \epsilon^\pi_{t+1}, \tag{10}$$

where ϵ^c_{t+1} and ϵ^π_{t+1} are iid standard normal shocks. Specifically, given $s_{t+1} \in \{n, d, i\}$, log consumption growth follows a normal distribution with mean $\mu^c_{s_{t+1}}$ and standard deviation σ^c, where $\mu^c_d < 0 < \mu^c_n$ and $\mu^c_i < 0 < \mu^c_n$. Log inflation follows a mean-reverting Gaussian process with persistence φ^π, mean $\mu^\pi_{s_{t+1}}$, and volatility σ^π. By definition, the deflationary recession regime has the lowest mean log inflation, whereas the inflationary recession regime has the highest: $\mu^\pi_d < \mu^\pi_n < \mu^\pi_i$.

Despite these distinct consumption and inflation dynamics, investors cannot precisely identify which economic state they are in, as s_{t+1} is not observable. The information friction arises from the fact the conditional means ($\mu^c_{s_{t+1}}$ and $\mu^\pi_{s_{t+1}}$) are not separately observable from the transitory shocks (ϵ^c_{t+1} and ϵ^π_{t+1}) although realized consumption and inflation (Δc_{t+1} and π_{t+1}) themselves are observable. Under this imperfect information, investors form their beliefs based on the historical time series of consumption $\Delta c_{-\infty:t+1} = \{\Delta c_k | k \leqslant t+1\}$ and inflation

$\pi_{-\infty:t+1} = \{\pi_k | k \leqslant t+1\}$.

Let $\xi_{s,t}$ denote investors' time-t belief that the economy is in state $s \in \{n, d, i\}$ at time t:

$$\xi_{s,t} \equiv \xi_{s,t|t} = P(s_t = s|\mathcal{F}_t),$$

where investors' time-t information set $\mathcal{F}_t$ consists of the past consumption and inflation realizations: $\mathcal{F}_t = \{\Delta c_{-\infty:t}, \pi_{-\infty:t}\}$. This belief is updated to $\xi_{s,t+1} = P(s_{t+1} = s|\mathcal{F}_{t+1})$ when investors observe new realizations of consumption growth Δc_{t+1} and inflation π_{t+1} at time $t+1$. Recent empirical evidence suggests that this belief updating process may not be perfectly Bayesian. Based on SPF inflation forecasts, Coibion and Gorodnichenko (2015) document that ex-post forecast errors are predicted by ex-ante forecast revisions with a positive coefficient, which can be directly linked to the degree of information rigidity. Consistent with this, our model allows for a potential departure from the Bayesian benchmark:

$$\xi_{s,t+1} = (1-\lambda) \times \xi^{\text{Bayes}}_{s,t,t+1} + \lambda \times \xi_{s,t},$$

where $\xi^{\text{Bayes}}_{s,t,t+1}$ represents investors' new belief if they were perfectly rational and updated their previous beliefs $(\xi_{n,t}, \xi_{d,t}, \xi_{i,t})$ in a Bayesian fashion based on the new information $(\Delta c_{t+1}, \pi_{t+1})$. According to Bayes' rule,

$$\xi^{\text{Bayes}}_{s,t,t+1} = \frac{P(s_{t+1} = s|\mathcal{F}_t) P(\Delta c_{t+1}, \pi_{t+1}|s_{t+1} = s, \mathcal{F}_t)}{\sum_{s' \in \{n,d,i\}} P(s_{t+1} = s'|\mathcal{F}_t) P(\Delta c_{t+1}, \pi_{t+1}|s_{t+1} = s', \mathcal{F}_t)}, \tag{11}$$

where the denominator equals $P(\Delta c_{t+1}, \pi_{t+1}|\mathcal{F}_t)$ but can be expressed as above due to the law of total probability. With a positive degree of information rigidity ($\lambda > 0$), investors' belief updating is sticky and puts λ as a weight on the past belief $\xi_{s,t}$. This leads to a positive slope coefficient when forecast errors are regressed on forecast revisions. As highlighted by Coibion and Gorodnichenko (2015), this regression coefficient enables us to pin down the parameter value of λ. In the absence of information rigidity ($\lambda = 0$), investors' learning dynamics return to the Bayesian case where forecast errors and forecast revisions are uncorrelated.

With the probability of observing the new consumption and inflation realizations (Δc_{t+1} and π_{t+1}) in the denominator, the numerator of equation (11) considers the likelihood that the observed realizations were generated under state s. This is a product of two probabilities. The first one is $P(s_{t+1} = s|\mathcal{F}_t)$, the probability that the true state of the economy is s at time $t+1$, which purely depends on the transition dynamics of the state between time t and time $t+1$, given the time-t information set. Denoting this as $\xi_{s,t+1|t} = P(s_{t+1} = s|\mathcal{F}_t)$, we have

$$\begin{bmatrix}\xi_{n,t+1|t} & \xi_{d,t+1|t} & \xi_{i,t+1|t}\end{bmatrix} = \begin{bmatrix}\xi_{n,t} & \xi_{d,t} & \xi_{i,t}\end{bmatrix} \times \mathcal{P}.$$

The second one is $P(\Delta c_{t+1}, \pi_{t+1}|s_{t+1} = s, \mathcal{F}_t)$, the probability of observing the new consumption and inflation realizations if the true state is indeed $s_{t+1} = s$. This term is crucial for investors' learning, as it determines how the new information is incorporated into their beliefs. It follows from the consumption and inflation dynamics in equations (9) and (10) that this probability is given by

$$P(\Delta c_{t+1}, \pi_{t+1}|s_{t+1} = s, \mathcal{F}_t) \quad \propto \quad \exp\left\{-\frac{1}{2}\left[\left(\tilde{\epsilon}^c_{s,t+1}\right)^2 + \left(\tilde{\epsilon}^\pi_{s,t+1}\right)^2\right]\right\},$$

where $\tilde{\epsilon}^c_{s,t+1} = \frac{\Delta c_{t+1} - \mu^c_s}{\sigma^c}$ and $\tilde{\epsilon}^\pi_{s,t+1} = \frac{\pi_{t+1} - (1-\varphi^\pi)\mu^\pi_s - \varphi^\pi \pi_t}{\sigma^\pi}$.

5.1.2 Equilibrium price-consumption and price-dividend ratios

We solve the model by finding the log price-consumption ratio $pc_t = \log\left(P^c_t/C_t\right)$ where P^c_t denotes the price of the consumption claim. In equilibrium, the log return on the consumption claim $r_{c,t+1} = \log\left(1 + e^{pc_{t+1}}\right) - pc_t + \Delta c_{t+1}$ must satisfy the Euler equation:

$$\mathbb{E}_t\left[\exp\left(\log M_{t+1} + r_{c,t+1}\right)\right] \quad = \quad 1, \tag{12}$$

where $\mathbb{E}_t$ represents investors' time-t expectation under their subjective probability measure. By plugging the expression for the pricing kernel (equation (7)) and that for log consumption growth (equation (9)) into the Euler equation, we recursively define the price-consumption ratio

as follows:

$$pc_t = \log \delta + \frac{1}{\theta} \log \mathbb{E}_t \left[e^{(1-\gamma)\Delta c_{t+1}} \left(1 + e^{pc_{t+1}}\right)^{\theta} \right].$$

The Markov property dictates that pc_t is a function of $\xi_{s,t}$ for $s \in \{n, d, i\}$. Since the three beliefs always sum to 1 at any given time, we can specify the solution in terms of just the two beliefs associated with the deflationary and inflationary recession regimes: $pc_t = pc(\xi_{d,t}, \xi_{i,t})$. We numerically solve for pc over a two-dimensional grid of $\xi_{d,t}$ and $\xi_{i,t}$ by searching for the fixed point of equation (13); the nonlinearity introduced by investors' learning dynamics does not permit a closed-form solution. See Ghaderi, Kilic, and Seo (2022) and Ghaderi, Kilic, and Seo (2024) for details on the numerical procedure.

As the model counterpart of the aggregate stock market, we consider a security that pays aggregate dividends, which are proxied by levered consumption $D_t = e^{\mu^d + \sigma^d \epsilon^d_{t+1}} C_t^{\phi}$. The log return on this dividend claim is expressed as $r_{d,t+1} = \log\left(1 + e^{pd_{t+1}}\right) - pd_t + \phi \Delta c_{t+1} + \mu^d + \sigma^d \epsilon^d_{t+1}$ and should satisfy the following Euler equation in equilibrium:

$$\mathbb{E}_t \left[\exp\left(\log M_{t+1} + r_{d,t+1} \right) \right] \;=\; 1, \tag{13}$$

where $pd_t = \log\left(P_t^d / D_t\right)$ denotes the log price-dividend ratio. Equation (13) leads to the following recursive relation between pd_t and pd_{t+1}:

$$pd_t = \theta \log \delta - (\theta - 1) pc_t + \mu^d + \frac{1}{2}(\sigma^d)^2 + \log \mathbb{E}_t \left[e^{(\phi-\gamma)\Delta c_{t+1}} \left(1 + e^{pc_{t+1}}\right)^{\theta - 1} \left(1 + e^{pd_{t+1}}\right) \right],$$

which is also solved numerically. Consequently, we can characterize the price-dividend ratio as a function of the two beliefs: $pd_t = pd(\xi_{d,t}, \xi_{i,t})$.

5.1.3 Term structures of model quantities

In this section, we describe a recursive method to compute the term structures of various model quantities based on our model. We start with the price of a risk-free nominal zero-coupon bond,

which is a function of three key state variables: the current inflation rate π_t and investors' two beliefs $\xi_{d,t}$ and $\xi_{i,t}$. Thus, we denote the bond price as $D(\pi_t, \xi_{d,t}, \xi_{i,t}; K)$, where K represents the number of periods until maturity. The pricing relation implies:

$$D(\pi_t, \xi_{d,t}, \xi_{i,t}; K) = \mathbb{E}_t \left[M^{\$}_{t \to t+K} \right] = \mathbb{E}_t \left[\prod_{k=1}^{K} M^{\$}_{t+k} \right] = \mathbb{E}_t \left[M^{\$}_{t+1} \mathbb{E}_{t+1} \left[\prod_{k=2}^{K} M^{\$}_{t+k} \right] \right],$$

where the last equality holds due to the law of iterative expectations. This expression provides us with a recursive formula for the term structure of zero-coupon bond prices:

$$D(\pi_t, \xi_{d,t}, \xi_{i,t}; K) = \mathbb{E}_t \left[M^{\$}_{t+1} D(\pi_{t+1}, \xi_{d,t+1}, \xi_{i,t+1}; K-1) \right].$$

Similarly, we obtain the term structure of inflation expectations using a recursive formula. For instance, let $\pi^e(\pi_{t+1}, \xi_{d,t}, \xi_{i,t}; K) = \mathbb{E}_t \left[\sum_{k=1}^{K} \pi_{t+k} \right]$ denote the expected log cumulative inflation rate over the next K periods. This quantity can be written as

$$\pi^e(\pi_t, \xi_{d,t}, \xi_{i,t}; K) = \mathbb{E}_t \left[\pi_{t+1} + \mathbb{E}_{t+1} \left[\sum_{k=2}^{K} \pi_{t+k} \right] \right] = \mathbb{E}_t \left[\pi_{t+1} + \pi^e(\pi_{t+1}, \xi_{d,t+1}, \xi_{i,t+1}; K-1) \right],$$

which enables us to find its value recursively. Similarly, the term structure for gross or annualized inflation can be calculated in a recursive fashion.

Next, we turn to the term structure of the inflation density $\mathbb{P}\left(\sum_{k=1}^{K} \pi_{t+k} = x \,\middle|\, \mathcal{F}_t \right)$. By the law of total probability,

$$\begin{aligned} \mathbb{P}\left(\sum_{k=1}^{K} \pi_{t+k} = x \,\middle|\, \mathcal{F}_t \right) &= \int \mathbb{P}\left(\sum_{k=1}^{K} \pi_{t+k} = x \,\middle|\, \mathcal{F}_{t+1} \right) dG\left(\Delta c_{t+1}, \pi_{t+1} \,\middle|\, \mathcal{F}_t \right), \\ &= \int \mathbb{P}\left(\sum_{k=2}^{K} \pi_{t+k} = \underbrace{x - \pi_{t+1}}_{= x'} \,\middle|\, \mathcal{F}_{t+1} \right) dG\left(\Delta c_{t+1}, \pi_{t+1} \,\middle|\, \mathcal{F}_t \right), \end{aligned}$$

where G represents the joint distribution of Δc_{t+1} and π_{t+1} conditional on $\mathcal{F}_t$. Denoting the

density by $\rho(x, \pi_t, \xi_{d,t}, \xi_{i,t}; K)$, the equation leads to the following recursion:[12]

$$\rho(x, \pi_t, \xi_{d,t}, \xi_{i,t}; K) = \int \rho(x', \pi_{t+1}, \xi_{d,t+1}, \xi_{i,t+1}; K-1)\, dG\left(\Delta c_{t+1}, \pi_{t+1} \,\middle|\, \mathcal{F}_t\right).$$

Once we obtain the conditional density of the average log inflation rate, we can easily transform it into the density of the average annualized inflation rate (which can be compared to the data) using a change of variable technique.

Finally, note that our recursive method also directly applies to inflation expectations and densities under the risk-adjusted probability measure. The only adjustment is to re-weight the probability distribution under $\mathbb{P}$ with the Radon-Nikodym derivative or the probability distortion factor $\frac{M_{t\to t+K}}{\mathbb{E}_t[M_{t\to t+K}]} = \frac{M_{t\to t+K}}{D(\pi_t, \xi_{d,t}, \xi_{i,t}; K)}$.

5.2 Calibration and asset pricing moments

Our model is calibrated at a monthly frequency, with the parameter values listed in Table 3. Panel A shows three preference parameters that describe investors' risk preferences. The coefficient of relative risk aversion γ is chosen to be 5, and the elasticity of intertemporal substitution ψ is set to 1.5. The subjective monthly discount rate is 0.9990, which corresponds to an annual discounting of 1.2%. These choices are standard and within reasonable ranges in the literature.

Panel B of Table 3 outlines the characteristics of the three regimes in our model: normal (n), deflationary recession (d), and inflationary recession (i). In the normal state, consumption grows at an average annual rate of 3%, while inflation averages 2%. For the recessionary states (d and i), we aim to capture severe economic downturns, akin to the Great Depression or the stagflation of the 1970s. Therefore, we set the average annual consumption growth at -5% and -3% for these two regimes, respectively. Despite similarities in consumption growth, these

[12] Under our inflation dynamics in equation (10), the following identify holds:

$$\rho(x, \pi_{t+1}, \xi_{d,t+1}, \xi_{i,t+1}; K) = \rho\left(x - \varphi\frac{1-\varphi^K}{1-\varphi}(\pi_{t+1} - \pi_t), \pi_t, \xi_{d,t+1}, \xi_{i,t+1}; K\right),$$

for any pair of π_t and π_{t+1}. This allows us to reduce the dimensionality of our numerical procedure, as it suffices to find the density ρ only for one value of π_t.

regimes differ markedly in inflation behavior: deflationary recessions exhibit an average annual deflation rate of -3%, whereas inflationary recessions are characterized by a high inflation rate of 12%. We calibrate endowment volatility to an annual rate of 1.5% across all regimes, reflecting the relatively stable consumption growth volatility observed in the post-war period. The log inflation volatility in the absence of a regime switch is set at 0.25% annually. In our model, inflation acts as an informative signal influencing learning dynamics. While headline inflation is often swayed by transient fluctuations in food and energy prices, investors are expected to focus on the more stable, persistent components of the price index to form their beliefs. Lastly, we set the persistence of monthly inflation at 0.5, in line with historical data.

Panel A: Preferences		
Relative risk aversion, γ	5	—
Elasticity of intertemporal substitution, ψ	1.5	—
Time discount factor, δ	0.9990	Annual discounting of 1.2%
Panel B: Consumption and inflation		
Mean consumption growth in state n, μ_n^c	0.0025	Annual growth rate of 3%
Mean consumption growth in state d, μ_d^c	-0.0042	Annual growth rate of -5%
Mean consumption growth in state i, μ_i^c	-0.0025	Annual growth rate of -3%
Volatility of consumption growth, σ^c	0.0043	Annual volatility of 1.5%
Mean log inflation in state n, μ_n^π	0.0017	Annual rate of 2%
Mean log inflation in state d, μ_d^π	-0.0025	Annual rate of -3%
Mean log inflation in state i, μ_i^π	0.0100	Annual rate of 12%
Volatility of log inflation, σ^π	0.0009	Annual volatility of 0.25%
Persistence of log inflation, φ^π	0.5000	—
Panel C: Regime switch and learning		
Transition probability from state n to d, p_{nd}	0.0004	Annual probability of 0.5%
Transition probability from state d to n, p_{dn}	0.0208	Average duration of 4 years
Transition probability from state n to i, p_{ni}	0.0042	Annual probability of 5.0%
Transition probability from state i to n, p_{in}	0.0104	Average duration of 8 years
Degree of information rigidity, λ	0.85	0.55 on a quarterly basis
Panel D: Dividends		
Leverage parameter, ϕ	3	—
Independent dividend growth, μ^d	-0.0017	Annual growth rate of -2%
Independent dividend volatility σ^d	0.0144	Annual growth volatility of 5%

Table 3: Model calibration parameters. This table presents the parameter values from the model calibration. The model is calibrated at a monthly frequency. For ease of interpretation, equivalent annualized values are also provided where applicable.

Given the extreme properties of consumption growth and inflation under the two recessionary regimes, these states naturally occur infrequently. In Panel C, the probability of transition-

ing from the normal regime to a deflationary recession (p_{nd}) is set at 0.5% annually, implying that such events are exceedingly rare, with an expected occurrence roughly once every two centuries. In contrast, the transition probability from the normal regime to an inflationary recession (p_{ni}) is set at 5.0%, reflecting the fact that the U.S. economy has experienced multiple inflationary recessions in the post-war period. This probability is ten times higher than that of a deflationary recession. The probabilities of returning to the normal state from the recessionary regimes reflect the typical duration of such downturns. For the deflationary regime, we set p_{dn} at 25% annually, suggesting an average recession length of four years, consistent with historical episodes such as the Great Depression and the Long Depression of the 1870s. Inflationary recessions, however, are modeled to be twice as persistent, with p_{in} set at 12.5%, consistent with the extended inflationary periods seen during the late 1960s and 1970s. For simplicity, we do not allow for direct transitions between the two recessionary regimes in our calibration.

Lastly, we set the parameter governing the stickiness of beliefs λ to 0.85. This choice is informed by the findings of Coibion and Gorodnichenko (2015), who, using quarterly regressions of forecast errors on forecast revisions, estimate that in updating inflation expectations, forecasters put approximately 54% weight on their prior beliefs, leaving 46% for new information. At the monthly frequency of our calibration, the weight assigned to the new Bayesian update is $1 - \lambda = 0.15$. This aligns well with the evidence of Coibion and Gorodnichenko (2015) as, ignoring the non-linearities of the updating procedure, the weight assigned to new information in a quarter would be around $3 \times 0.15 = 0.45$.

Finally, to model the equity market dynamics, we set the leverage parameter ϕ to 3, which is consistent with the standard range commonly used in the literature (see, e.g., Bansal and Yaron, 2004; Bansal, Kiku, and Yaron, 2012; Wachter, 2013). We calibrate μ^d, the mean dividend growth rate that is independent of consumption growth, to -2%, ensuring that the average growth rate of dividends does not deviate significantly from consumption growth, as observed in the data. Additionally, we assign an independent annual volatility of 5% for σ^d.

To demonstrate the quantitative validity of our calibration, we analyze the population char-

acteristics of the model by simulating a long sample spanning 100,000 years. As shown in Table 4, the model preserves low consumption volatility over this long sample, with a value of 1.70%, closely aligned with post-war data (1.52%). Importantly, the model successfully addresses key asset pricing puzzles, generating a high equity premium (6.73% versus 6.57% in the data) and high market volatility (13.21% versus 16.46% in the data), while maintaining a low risk-free rate (0.59% versus 0.39% in the data). Additionally, the model produces an upward-sloping nominal term structure with a term premium of 0.57%, consistent with empirical evidence (0.88% in the data). While the mean and volatility of the inflation rate are slightly higher in the simulated sample (4.67% and 4.43%, respectively) than in post-war data (3.54% and 2.81%, respectively), this discrepancy is expected given the rare nature of the two recessionary states.

Lastly, to verify the model's ability to capture realistic belief dynamics, we regress inflation forecast errors on quarterly forecast revisions, following the method of Coibion and Gorodnichenko (2015), and find results consistent with theirs (1.32 versus 1.19). Overall, our model effectively captures standard asset pricing moments alongside inflation dynamics and its expectations.

	$\sigma(\Delta c_{t+1})$	$\sigma(\Delta d_{t+1})$	$\mathbb{E}[\pi_{t+1}]$	$\sigma(\pi_{t+1})$	β^{CG}	$\mathbb{E}[r_{f,t}]$	$\mathbb{E}[r_{d,t+1} - r_{f,t}]$	$\sigma(r_{d,t+1})$	$\mathbb{E}[y^{10y} - y^{2y}]$
Data	1.52	7.04	3.54	2.81	1.19	0.36	6.57	16.46	0.88
Model	1.70	7.15	4.67	4.43	1.32	0.59	6.73	13.21	0.57

Table 4: Standard asset pricing moments. This table reports the moments of consumption growth, inflation, and asset returns in the data and in the model. $\sigma(\Delta c_{t+1})$ represents the standard deviation of log consumption growth. $\mathbb{E}[\pi_{t+1}]$ and $\sigma(\pi_{t+1})$ represent the average and standard deviation of the annual inflation rate, respectively. β^{CG} corresponds to the slope coefficient in the regression of inflation forecast errors on forecast revisions as in Coibion and Gorodnichenko (2015). $\mathbb{E}[r_{f,t}]$ is the average log real risk-free rate, while $\mathbb{E}[r_{d,t+1} - r_{f,t}]$ denotes the average excess log return on the market. $\sigma(r_{d,t+1})$ is the standard deviation of the log market return. The term premium, defined as the difference between 10-year and 2-year Treasury yields, is represented by $\mathbb{E}[y^{10y} - y^{2y}]$. With the exception of the term premium, which is based on data from 1976 to 2023, all other data moments reflect the post-war sample period from 1947 to 2023. The model moments are calculated from a long simulation spanning 100,000 years. All values are in percentages.

5.3 Model implications

As discussed in Section 4 with equation (4), the change of measure from $\mathbb{P}$ to $\mathbb{Q}_{t+T}$ skews the probabilities, underweighting favorable events and overweighting unfavorable ones. As a result, a positive (negative) log probability ratio signals bad (good) inflation outcomes from the investor's perspective. Figure 18 evaluates the implications of our model by examining the value of this ratio for different inflation outcomes over the next year. The figure aims to illustrate how preferences for various inflation ranges change over time within the model. Blue bars in the figure represent positive log probability ratios, while yellow bars indicate negative ratios.

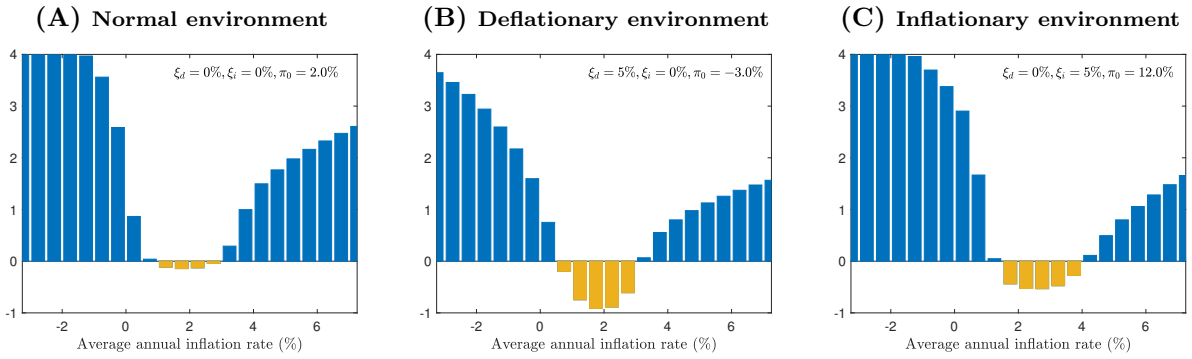


Figure 18: Model-implied inflation densities. This figure presents the model-implied conditional inflation probability ratios across three scenarios: normal times with moderate inflation (Panel A), heightened risk of deflationary depression with low inflation (Panel B), and heightened risk of inflationary depression with high inflation (Panel C). Each panel shows the log probability ratio between the risk-adjusted distribution and the physical distribution $\log(\mathbb{Q}/\mathbb{P})$. The distributions capture the probabilities of different inflation outcomes over the next year. All values are in percentages.

The figure highlights three distinct economic scenarios. In Panel A, investors are assumed to assign a 100% probability to the "normal" regime. In contrast, Panels B and C depict situations where the representative investor assigns a 5% probability to the occurrence of either a deflationary (Panel B) or inflationary (Panel C) recession. Additionally, we assume that the monthly inflation rate (π_0) takes the average value associated with each respective regime. These scenarios are chosen to establish a mapping between the model and historical episodes from the second quarter of 2018, 2020, and 2022, as illustrated in Figure 7, allowing us to analyze time variation in inflation preferences.

In Panel A, we observe that the good inflation region is relatively narrow, centered around 2%. The model can accommodate a substantial increase in the probability ratio, and thus the marginal utility, as inflation moves toward both low and high extremes. The stickiness of investor beliefs results in a high level of persistence in beliefs and a substantially low perceived chance of extreme inflation outcomes. However, if such extreme inflation outcomes were to occur, the model implies that investors would need to experience a series of "bad" shocks, which would persistently increase the probabilities of recessionary states (ξ_d or ξ_i). This would, in turn, lead to a significant rise in marginal utility. Note that while the U-shape persists under the Bayesian benchmark, achieving such high levels of marginal utility as observed in the data would be more difficult with reasonable parameters.

Panel B illustrates a scenario with a heightened risk of a deflationary recession, similar to the economic environment of summer 2020 when concerns about a deflationary depression intensified. Consistent with the data, this leads to an expansion of the good inflation range, along with a leftward shift in the U-shape. The widening of the good range can be attributed to increased overall uncertainty regarding inflation outcomes, which causes a relatively larger increase in physical probabilities.

In contrast, Panel C depicts a scenario with a heightened risk of an inflationary recession, reminiscent of the environment from mid-2021 to late 2022. While there is a similar expansion of the good inflation range, this range is now shifted rightward. In the model, this shift occurs due to changes in the conditional probability of inflation over the next year. Practically, a higher belief in the likelihood of an inflationary state corresponds to higher levels of the inflation state variable and expectations of future inflation. Given the persistent nature of inflation shocks in both the model and the data, this shifts the entire conditional distributions under both the risk-neutral and physical measures. Despite these differences across the three scenarios, the U-shaped pattern in the probability ratios remains consistent, as also observed in the data.

How does the inflation risk premium change under these three scenarios, and which ranges of inflation outcomes contribute most to it in the model? We address these questions in Figure 19. Following the approach of Beason and Schreindorfer (2022), we decompose the conditional

inflation risk premium in the model for each scenario. The goal, once again, is to establish a mapping between these model scenarios and key historical episodes shown in Figure 11.

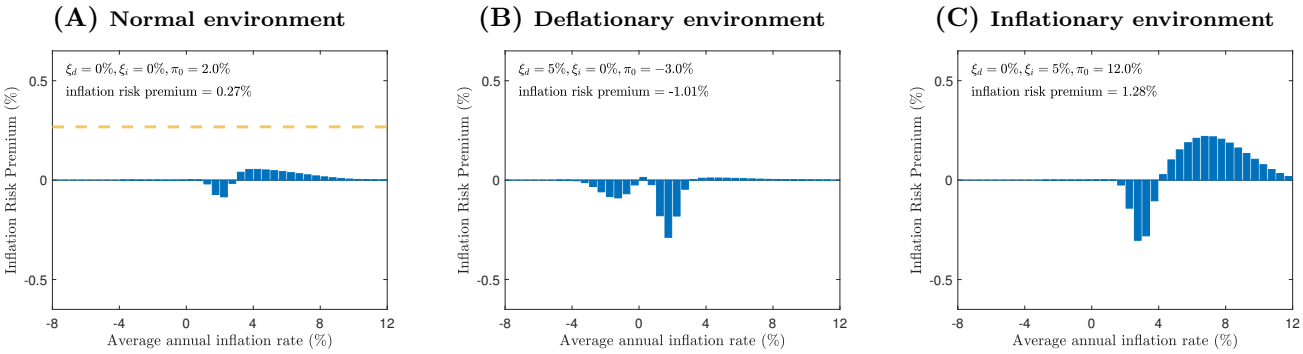


Figure 19: Model-implied decomposition of the inflation risk premium . This figure presents the model-implied conditional decomposition of the inflation risk premium across three scenarios: normal times with moderate inflation (Panel A), heightened risk of deflationary depression with low inflation (Panel B), and heightened risk of inflationary depression with high inflation (Panel C). Each panel shows the inflation risk premium over the next year, together with its decomposition into compensation for different inflation outcomes. All values are in percentages.

In the normal environment depicted in Panel A, the inflation risk premium is relatively small, with a modest overall positive premium of 0.27%. The decomposition shows that the model provides minimal compensation for extreme inflation outcomes. Due to the persistent nature of beliefs, investors in this scenario assign low probabilities to such extreme outcomes. This scenario mirrors the relatively calm environment observed during the summer of 2018, when inflation expectations were well-anchored around 2%.

In the deflationary environment (Panel B), the inflation risk premium turns negative (-1.01%). This situation corresponds to periods such as the summer of 2020, when fears of a deflationary recession rose sharply due to the COVID-19 pandemic. While the good region around 2% negatively contributes to the premium, the figure also highlights a significant negative contribution from inflation rates below 0%. This reflects investors' heightened concerns about deflation. As previously explained, because the inflation outcomes are negative, the overall premium turns negative, despite investors' anxieties about these potential outcomes. Lastly, in Panel C, the model produces a substantial positive inflation risk premium of 1.28% under the heightened inflationary risk scenario. This positive contribution to the risk premium

is concentrated in the higher inflation ranges, particularly those exceeding 4%. The good range between 2% and 4% negatively contributes to the inflation risk premium. These patterns are largely consistent with what we observed in Figure 11.

It is worth highlighting that a traditional model treating inflation risk as a monotonic factor cannot capture the U-shaped pattern of the projected pricing kernel. In the Internet Appendix, we demonstrate this limitation using a special case of our model with only two regimes – the normal regime n and one recession regime (either d or i). For instance, consider the case where higher inflation is unanimously bad, as it correlates with lower growth (i.e., n vs i). In this setting, the projected pricing kernel becomes an increasing function of future inflation. Conversely, if higher inflation is unanimously good, we obtain a decreasing function. Neither case is consistent with the empirical U-shaped pattern documented in Section 4.

This misspecification of the pricing kernel has critical implications for how the inflation risk premium is generated and decomposed across different inflation ranges. Under an increasing (decreasing) pricing kernel, both high and low inflation outcomes contribute positively (negatively) to the total inflation risk premium. Not only are these results inconsistent with the decomposition patterns observed in the data, but they also impose tight restrictions on how the inflation risk premium emerges from the quantity and price of inflation risk. In particular, because the premiums from high and low inflation outcomes have the same sign and do not offset one another, a small overall inflation risk premium can arise only from a small quantity or a small price of inflation risk – unlike the case with a U-shaped pricing kernel. To better capture the origins of the inflation risk premium, it is important to realistically model both ends of the inflation distribution and their relationship with future growth.

6 Conclusion

This paper provides novel insights into the complex nature of inflation risk and its pricing in financial markets. By comparing survey-based and risk-adjusted probability distributions of future inflation, we extract investors' time-varying preferences for different inflation ranges. Our

findings reveal a U-shaped pattern in the probability ratio between these distributions, confirming that investors generally dislike both high and low inflation environments while favoring outcomes around the Fed's 2% target. Importantly, we demonstrate that these preferences are not static but evolve significantly over time, reflecting changing economic conditions and investor perceptions. These empirical findings are further supported by our economic model, which incorporates learning about unobservable economic regimes.

Our exercise on the decomposition of the inflation risk premium offers a more nuanced understanding of how different inflation ranges contribute to the overall premium. This approach allows us to identify which inflation scenarios investors find particularly concerning at different points in time, providing valuable insights that are often masked when focusing solely on the aggregate premium. The analysis of specific periods, such as the deflationary concerns in 2020 and the inflationary pressures in 2022, illustrates how our methodology can capture shifts in investor sentiment and economic outlook.

References

Afrouzi, Hassan, Spencer Y Kwon, Augustin Landier, Yueran Ma, and David Thesmar, 2023, Overreaction in expectations: Evidence and theory, *Quarterly Journal of Economics* 138, 1713–1764.

Aït-Sahalia, Yacine, and Jefferson Duarte, 2003, Nonparametric option pricing under shape restrictions, *Journal of Econometrics* 116, 9–47.

Andrei, Daniel, and Michael Hasler, 2023, Can the Fed control inflation? Stock market implications, Working paper.

Ang, Andrew, Geert Bekaert, and Min Wei, 2007, Do macro variables, asset markets, or surveys forecast inflation better?, *Journal of Monetary Economics* 54, 1163–1212.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, An empirical evaluation of the long-run risks model for asset prices, *Critical Finance Review* 1, 183–221.

Bansal, Ravi, and Ivan Shaliastovich, 2010, Confidence risk and asset prices, *American Economic Review: Papers and Proceedings* 100, 537–41.

Bansal, Ravi, and Ivan Shaliastovich, 2013, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies* 26, 1–33.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long-run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Beason, Tyler, and David Schreindorfer, 2022, Dissecting the equity premium, *Journal of Political Economy* 130, 2203–2222.

Bianchi, Francesco, Martin Lettau, and Sydney C Ludvigson, 2022, Monetary policy and asset valuation, *Journal of Finance* 77, 967–1017.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer, 2020, Overreaction in macroeconomic expectations, *American Economic Review* 110, 2748–2782.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2019, Diagnostic expectations and stock returns, *Journal of Finance* 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2018, Diagnostic expectations and credit cycles, *Journal of Finance* 73, 199–227.

Bouchaud, Jean-Philippe, Philipp Krueger, Augustin Landier, and David Thesmar, 2019, Sticky expectations and the profitability anomaly, *Journal of Finance* 74, 639–674.

Breeden, Douglas T., and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.

Campbell, John Y, Carolin Pflueger, and Luis M Viceira, 2020, Macroeconomic drivers of bond and equity risks, *Journal of Political Economy* 128, 3148–3185.

Campbell, John Y., Adi Sunderam, and Luis M. Viceira, 2017, Inflation bets or deflation hedges? The changing risks of nominal bonds, *Critical Finance Review* 6, 263–301.

Clements, Michael P, 2018, Are macroeconomic density forecasts informative?, *International Journal of Forecasting* 34, 181–198.

Coibion, Olivier, and Yuriy Gorodnichenko, 2012, What can survey forecasts tell us about information rigidities?, *Journal of Political Economy* 120, 116–159.

Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Information rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–2678.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar, 2018, How do firms form their expectations? New survey evidence, *American Economic Review* 108, 2671–2713.

Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele, 2020, Inflation expectations and firm decisions: New causal evidence, *Quarterly Journal of Economics* 135, 165–219.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber, 2022, The expected, perceived, and realized inflation of US households before and during the COVID19 pandemic, Working paper.

David, Alexander, and Pietro Veronesi, 2013, What ties return volatilities to price valuations and fundamentals?, *Journal of Political Economy* 121, 682–746.

Diebold, Francis X, Anthony Tay, and Kenneth Wallis, 1997, Evaluating density forecasts of inflation: The survey of professional forecasters, Working Paper.

Dykstra, Richard L., 1983, An algorithm for restricted least squares regression, *Journal of the American Statistical Association* 78, 837–842.

Epstein, Larry, and Stan Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig, 2017, Deflation risk, *Review of Financial Studies* 30, 2719–2760.

Ghaderi, Mohammad, Mete Kilic, and Sang Byung Seo, 2022, Learning, slowly unfolding disasters, and asset prices, *Journal of Financial Economics* 143, 527–549.

Ghaderi, Mohammad, Mete Kilic, and Sang Byung Seo, 2024, Why do rational investors like variance at the peak of a crisis? A learning-based explanation, *Journal of Monetary Economics* 142, 103513.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken, 2012, Inflation expectations, real rates, and risk premia: Evidence from inflation swaps, *Review of Financial Studies* 25, 1588–1629.

Hilscher, Jens, Alon Raviv, and Ricardo Reis, 2022, Inflating away the public debt? An empirical assessment, *Review of Financial Studies* 35, 1553–1595.

Hilscher, Jens, Alon Raviv, and Ricardo Reis, 2024, How likely is an inflation disaster?, Working paper.

Kenny, Geoff, Thomas Kostka, and Federico Masera, 2014, How informative are the subjective density forecasts of macroeconomists?, *Journal of Forecasting* 33, 163–185.

Kitsul, Yuriy, and Jonathan H Wright, 2013, The economics of options-implied inflation probability density functions, *Journal of Financial Economics* 110, 696–711.

Li, Haitao, and Feng Zhao, 2009, Nonparametric estimation of state-price densities implicit in

interest rate cap prices, *Review of Financial Studies* 22, 4335–4376.

Lombardi, Marco Jacopo, and Vladyslav Sushko, 2023, The correlation of equity and bond returns, *BIS Quarterly Review*.

Mankiw, N Gregory, and Ricardo Reis, 2002, Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve, *Quarterly Journal of Economics* 117, 1295–1328.

Mertens, Thomas M, and John C Williams, 2021, What to expect from the lower bound on interest rates: Evidence from derivatives prices, *American Economic Review* 111, 2473–2505.

Piazzesi, Monika, and Martin Schneider, 2006, Equilibrium yield curves, *NBER Macroeconomics Annual* 21, 389–442.

Rich, Robert, and Joseph Tracy, 2010, The relationships among expected inflation, disagreement, and uncertainty: Evidence from matched point and density forecasts, *Review of Economics and Statistics* 92, 200–207.

Wachter, Jessica A., 2006, A consumption-based model of the term structure of interest rates, *Journal of Financial Economics* 79, 365–399.

Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *Journal of Finance* 68, 987–1035.

Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 402–421.