The Dynamics of Long-Run Inflation Expectations: A Market-Based Perspective

Anna Cole, Julian Kozlowski, and Joseph Martorana

Abstract

This article analyzes market-based probability distributions for long-run inflation expectations derived from inflation derivatives. We construct forward-looking distributions for five-year-ahead inflation to assess the likelihood that inflation will fall above, below, or near the Federal Reserve's 2 percent target. By examining the mean, volatility, and skewness of these distributions, we document how expectations have evolved since the onset of the COVID-19 pandemic. To assess the reliability of market-based measures, we compare our results with alternative data sources. We highlight the elevated probability of inflation exceeding the 2 percent target that persisted shortly after the COVID-19 pandemic. The findings underscore the importance of market-based tools in capturing nuanced inflation dynamics and informing policy and financial decisions.

JEL code: E31

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2025, Vol. 107, No. 13, pp. 1-14.
https://doi.org/10.20955/r.2025.13

1. INTRODUCTION

Long-run inflation expectations are a key determinant of monetary policy decisions and financial market dynamics. While most existing measures focus on point estimates—such as the expected average inflation rate—we complement this literature by constructing full probability distributions of future inflation. This approach captures not only the central tendency but also the likelihood of a wide range of inflation outcomes, including extreme scenarios.

Point forecasts convey limited information about uncertainty and tail risks. In contrast, probability densities offer a richer perspective, describing the full distribution of possible outcomes and the relative likelihood of each. This allows policymakers and market participants to better understand the risks of extreme inflation—either very high or very low—that may materialize over time.

In this article, we analyze the forward distribution of inflation, focusing on its mean, variance, and skewness. These moments provide insight into the level, uncertainty, and asymmetry of inflation expectations. Our analysis relies on market-based information extracted from the prices of inflation derivatives. Specifically, we

Julian Kozlowski is a senior economic policy advisor, Anna Cole is a research associate, and Joseph Martorana is a former research associate at the Federal Reserve Bank of St. Louis. The authors thank Carlos Garriga, Ken Heineche, Alberto Musalem, Jason Schmidt, and Mark Wright for useful comments and suggestions.

use inflation caps and floors, which are options that provide payouts when inflation rises above or falls below a specified threshold. These derivatives embed investor beliefs about future inflation rates.

In addition to estimating these distributions, we assess the data quality and key assumptions underlying the market-based approach. We evaluate how these assumptions affect inference, particularly in the post-COVID period when inflation uncertainty has increased. By comparing pre-pandemic stability with the volatility of recent years, we highlight the value and limitations of these market-derived measures.

Literature Review An extensive literature covers the methods of using tradable securities to measure market expectations. Feldman et al. (2015) argue that policymakers face considerable uncertainty regarding policy outcomes and should accordingly weigh the net benefits of a policy in a given possible future against the benefits of that policy in alternative futures. They also assert that as policymakers implement policy on behalf of constituents or households, they should consider households' views of the net costs or benefits of specific policies. In determining the appropriate stance of monetary policy, central bankers would have to account for households' views of inflation risks. Rather than collect these views through surveys, central bankers can obtain them from financial markets where households and financial institutions trade securities to hedge against these inflation risks.

Some researchers have focused on ways to obtain probability distributions from such securities. Aside from the well-known inflation-indexed Treasuries, or Treasury inflation-protected securities (TIPS), inflation derivatives can also be used to isolate the probability of inflation occurring at various rates. Kitsul and Wright (2012) (KW12) use daily price quotes for inflation caps and floors for various strike prices from 2009 to 2012 to create custom portfolios of those securities, where each portfolio is centered on a particular strike price, representing a given inflation rate. The market price of each portfolio can be used to determine the market-implied risk-neutral probability of inflation occurring at its strike rate. With a collection of portfolios with strike prices of 1 percent, 2 percent, 3 percent, and so on, KW12 create a distribution of probabilities over a range of inflation rates. In a later article, Kitsul and Wright (2013) (KW13) use the second derivative of the pricing equation for inflation caps with respect to their strike price, as the second derivative represents the risk-neutral probability density function (Breeden and Litzenberger, 1978; Ross, 1976).

These methods have also been used in other work where distributions of expected inflation reveal what single point estimates of inflation forecasts cannot. Mertens and Williams (2021) construct inflation distributions according to KW13 to study how the distributions change with the natural rate of interest. Hilscher, Raviv, and Reis (2022) build on KW13 by estimating tail probabilities for very high or low persistent inflation using inflation caps and floors. Importantly, they show that while the average of inflation expectations in the U.S. and the eurozone remained anchored in 2021 and 2022, the tails of the distributions did not. They argue that a focus on average inflation expectations led policymakers to believe that inflation would be transitory at the time, whereas accounting for the tails may have led to alternative conclusions.

The rest of the article proceeds as follows. Section 2 introduces the mechanics and pricing of inflation caps and floors. Section 3 explores the data used in our analysis, including changes in cap and floor prices over time and their implications for the market-based probabilities. Section 4 estimates the inflation distribution using parametric methods and discusses its robustness through comparisons with alternative measures. Section 5 concludes.

2. INFLATION OPTIONS: CAPS AND FLOORS

This section introduces the mechanics of inflation caps and floors, which are key derivatives used by investors to manage inflation risk. We also explain how these instruments are priced and later use their market prices to derive market-based inflation probabilities.

2.1 Inflation Caps and Floors

Inflation caps and floors are European-style options designed to provide payouts based on inflation exceeding or falling below a predetermined threshold, referred to as the *strike price*. Inflation caps ensure a payout if inflation surpasses the strike price, shielding investors from unexpectedly high inflation. Conversely, if inflation remains below the strike price, no payment is made. Inflation floors provide a payout if inflation falls below the strike price, protecting investors from deflationary (or disinflationary) pressures.

The payoff for a cap or floor with a unit notional, a strike price k, and a maturity of h years is given by

$$\text{Payoff}_{\text{cap}(h,k)} = \max\{(1+\hat{\pi}_{t+h})^h - (1+k)^h, 0\}, \tag{1}$$

$$\text{Payoff}_{\text{floor}(h,k)} = \max\{(1+k)^h - (1+\hat{\pi}_{t+h})^h, 0\}. \tag{2}$$

Figure 1
Inflation Cap and Floor, Strike $k = 2$ Percent

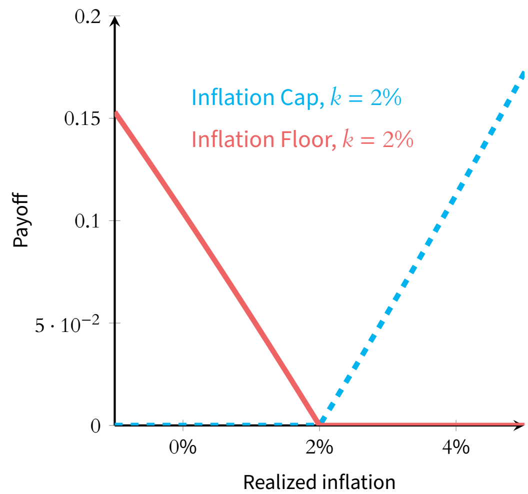

Here, $\hat{\pi}_{t+h}$ represents the observed inflation rate—specifically, the headline consumer price index (CPI), nonseasonally adjusted—over the contract's maturity.

Figure 1 shows the payoff of both an inflation cap (dashed blue line) and floor (solid red line) with a strike price of 2 percent.

An investor who owns the cap receives a payout only once inflation exceeds 2 percent and receives no payout otherwise. Their payout is the difference between the observed inflation rate and the strike price, multiplied by the contract notional. In simple terms, the payoff is the higher of zero or the difference between the actual average annualized inflation rate and strike price k over the period from t to $t + h$. Accordingly, it can be inferred that investors who own inflation caps with a strike price of 2 percent expect that inflation will exceed an annual average rate of 2 percent over the contract's term because they receive a payment only in that case. Conversely, an investor who owns an inflation floor with a strike price of 2 percent receives a payment only if average inflation falls below 2 percent.

These types of inflation options have been traded over the counter since 2009, with the liquidity of the caps and floors market varying over time. Chipeniuk and Walker (2021) note that the volume of trades for caps and floors generally increased from 2011 to 2017 but then declined in 2018. We discuss these liquidity concerns and their implications for data reliability in Section 3.3.2.

2.2 Risk-Neutral Pricing

The price of inflation caps and floors is the discounted expected value of their future payoffs under the risk-neutral measure Q:[1]

$$P_{cap}(k) = e^{-hy(h)}\mathbb{E}^{Q}\left[\max\left\{(1+\bar{\pi}_{t+h})^{5} - (1+k)^{h}, 0\right\}\right], \tag{3}$$

$$P_{floor}(k) = e^{-hy(h)}\mathbb{E}^{Q}\left[\max\left\{(1+k)^{h} - (1+\bar{\pi}_{t+h})^{h}, 0\right\}\right]. \tag{4}$$

Equations (3) and (4) show the price of inflation caps $P_{\text{cap}}(k)$ and floors $P_{\text{floor}}(k)$, with a term of h years and a strike price of k. $y(h)$ is the risk-free discount rate for maturity h, and $\bar{\pi}_{t+h}$ is the average annualized inflation rate over the contract's term. These prices are the present value of its risk-neutral expected payoff at maturity, with the payoff structure described in Section 2.1. The price of a cap or floor is essentially a premium paid to the contract's seller, based on the contract notional or the size of the investment that the buyer is looking to protect against inflation. For example, if the price of a cap is 100 basis points (1 percent), and if the buyer wants to protect a notional of \$10,000 against inflation above the strike price, then the buyer pays \$100, or 1 percent of \$10,000 to the contract seller.

1. Risk-neutral probabilities reflect the probabilities of future outcomes as implied by asset prices under the assumption that investors are indifferent to risk. They are not actual beliefs but rather adjusted probabilities that incorporate investors' risk preferences. In the context of inflation options, they correspond to the probability distribution of future inflation under the forward martingale measure, which discounts payoffs using the yield on zero-coupon bonds of the corresponding maturity (see Kitsul and Wright, 2013).

Figure 2
Butterfly Portfolio

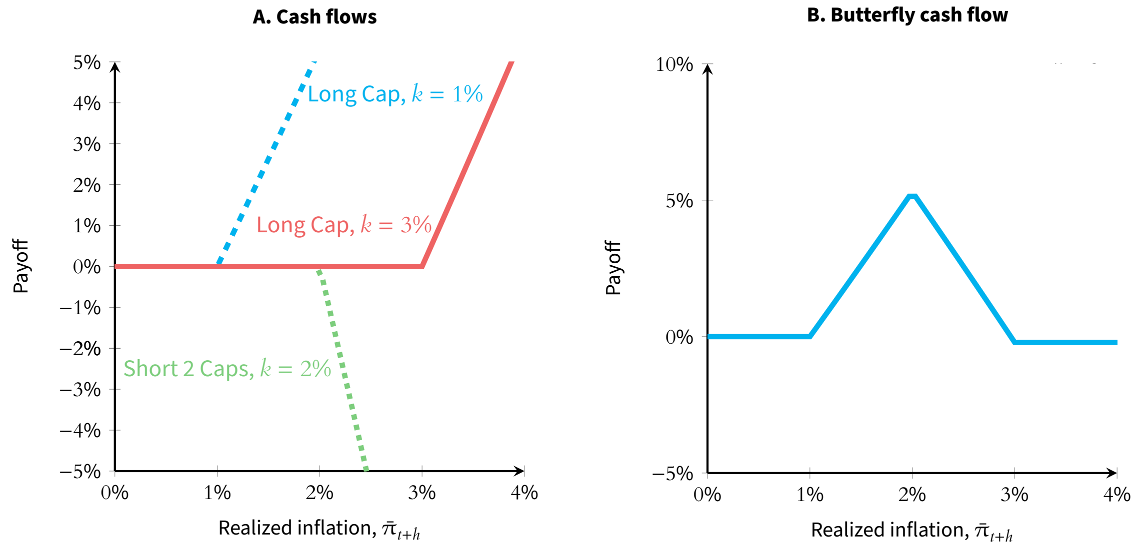

2.3 Butterfly Portfolio

Butterfly portfolios are a strategic combination of securities that allows investors to construct synthetic instruments that provide payouts within a specific range of inflation outcomes. For instance, an investor seeking a payout if inflation falls between 1 and 3 percent over the next year could achieve this by taking the following positions: a long position on a 1 percent inflation cap, a long position on a 3 percent inflation cap, and a short position on two 2 percent inflation caps.

The cash flows generated by each of these positions are illustrated in Panel A of Figure 2. If inflation is below 1 percent, no payout is received. For inflation between 1 and 2 percent, the payout comes solely from the 1 percent cap (represented by the dashed blue line). For inflation of 3 percent or higher, payouts are received from both the 1 percent cap and the 3 percent cap (solid red line), but the investor incurs payment obligations to the holders of the 2 percent caps that were shorted (dotted green line). As shown in Panel B of the figure, the net cash flows from this strategy result in a payout exclusively when inflation falls between 1 and 3 percent.

The price of this bundle of caps is the present value of the risk-neutral probability of inflation falling between 1 and 3 percent over the term of the caps and floors. By constructing similar butterfly portfolios for various ranges of strike prices, we can derive the risk-neutral discounted probability distribution of inflation across different strike intervals. Using the appropriate discount rate also allows us to back out the associated risk-neutral probability of different scenarios and construct a risk-neutral distribution of expected inflation.

3. DATA

In this section we present data on prices of inflation caps and floors and the implied risk-neutral distribution. We review how these contracts have changed since 2020 and assess how the implied moments of the risk-neutral distribution compare with other measures of inflation expectations.

3.1 Prices of Inflation Caps and Floors

Figure 3 shows how the quoted prices of a handful of caps and floors changed from January 2020 through December 2024. Panel A shows daily price quotes from Bloomberg of five-year zero-coupon caps with strike prices of 3, 4, 5, and 6 percent (their Bloomberg tickers are "USIZC35," "USIZC45," "USIZC55," and "USIZC65," respectively). These caps and their floor counterparts are indexed against the nonseasonally adjusted headline U.S. CPI for All Urban Consumers: U.S. City Average, measured by the Bureau of Labor Statistics (BLS).

Prices were quite low in the early part of 2020, though the price of the 3 percent cap began to climb noticeably before the year was over. This implies that markets assigned a very low probability of the inflation rate exceeding 3 percent over the next five years. However, during 2021, the price of inflation caps started to increase. By early 2022, when annual CPI inflation had surpassed 8 percent, the 3 percent cap had a quoted price of approximately 400 basis points, while the 4 percent cap had a quoted price of more than 100 basis

Figure 3
Quoted Prices of Inflation Caps and Floors

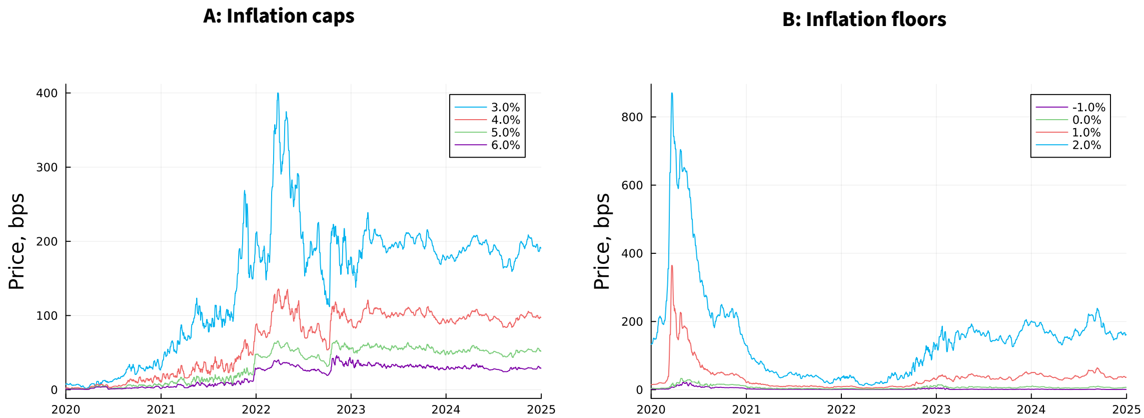

SOURCE: Daily data obtained from Bloomberg from January 1, 2020, through December 31, 2024. Notes: Tickers for 3, 4, 5, and 6 percent five-year zero-coupon caps are "USIZC35," "USIZC45," "USIZC55," and "USIZC65," respectively. Tickers for –1, 0, 1, and 2 percent five-year zero-coupon floors are "USIZFZ5," "USIZFA5," "USIZF15," and "USIZF25," respectively.

points. By 2023, cap prices had come down from their peaks but were still elevated compared with where they were in 2020. This may reflect changes in average inflation expectations over the next five years, the volatility of inflation expectations, or changes in the discount rate. We investigate these issues further in Section 3.2.

Panel B of Figure 3 shows the quoted daily prices of zero-coupon five-year inflation floors with strike prices of –1, 0, 1, and 2 percent (these are also obtained from Bloomberg using the tickers "USIZFZ5," "USIZFA5," "USIZF15," and "USIZF25," respectively). Unsurprisingly, the pattern here is the opposite of the caps. The prices of the inflation floors were much higher in early 2020. Investors generally expected low inflation, as the U.S. was in the depth of the COVID-19 recession. However, as CPI inflation began to pick up, the price of a 2 percent inflation floor declined precipitously from over 800 basis points in 2020 to under 100 basis points by 2022. The prices of all other floors similarly fell by early 2022. Since CPI inflation reached its recent peak, the prices of inflation floors have increased but to nowhere near their previous levels. This may also reflect changes in average inflation expectations, their volatility, or the discount rate.

Figure 4 shows the relationship between the strike prices and contract prices of zero-coupon inflation caps and floors. Two curves are provided for each type of contract: The red curve shows the prices of a five-year contract on March 10, 2022, for different strike prices, while the blue curve shows the prices of identical contracts on December 31, 2024.

Panel A of Figure 4 shows that the price-strike curve of inflation caps has clearly shifted down since early 2022, consistent with inflation cooling over the following two years. The contract price of a 3 percent cap has fallen the most, from over 300 basis points in 2022 to approximately 200 basis points in 2025. As inflation expectations came down from their 2022 highs, the prices of caps have followed. The prices of caps with higher strike prices have shown little movement, and this may suggest that even as annual CPI inflation increased over 2021 and 2022, investors mostly did not expect inflation over the five-year horizon to average more than 5 percent.

Panel B shows that the prices of inflation floors have increased since 2022 as investors have viewed it more likely that future inflation will decelerate. As inflation is more likely to be below 2 percent now than it was two years ago, the 2 percent inflation floor commands a premium. The price of this contract grew dramatically, from less than 25 basis points in 2022 to over 150 basis points in 2024. However, even as inflation has fallen since the Federal Reserve began its interest rate tightening cycle, investors still consider inflation below 0 percent to be highly unlikely. Thus, the prices of the 0 percent and –1 percent floors have hardly increased. Overall, the shifts of the price-strike curves of inflation caps and floors are consistent with inflation moving back toward the Federal Reserve's 2 percent inflation target.

3.2 Market-Based Probabilities

While the price of an inflation cap or floor contract can generally imply what market participants expect to happen over its term, the price of a butterfly portfolio can reveal more details about the distribution of investors' inflation expectations. This kind of inference from asset prices also applies to investor expectations regarding

Figure 4
Price-Strike Curve of Inflation Caps and Floors

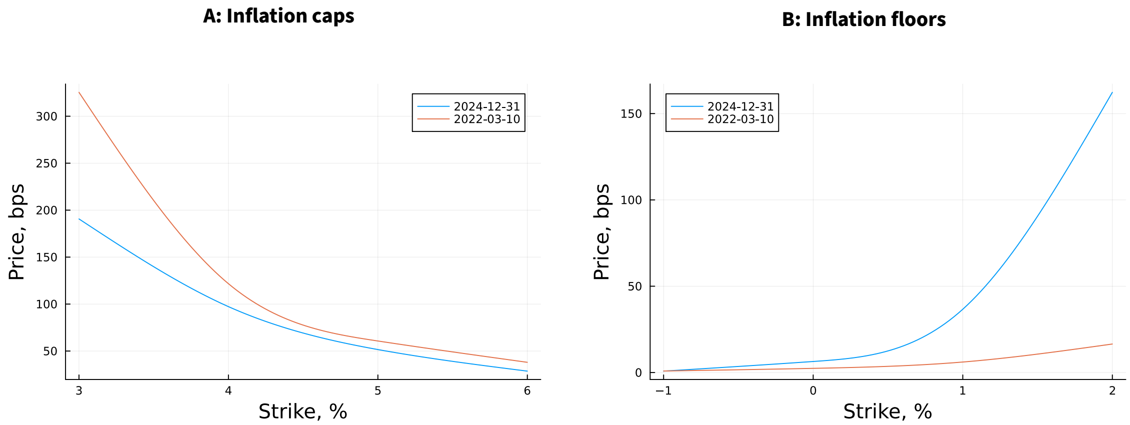

SOURCE: Bloomberg.

other economic developments. The Federal Reserve Bank of Minneapolis tracks expectations for a range of indicators on its Current and Historical Market-Based Probabilities webpage. It maintains weekly statistics for expectations regarding 5- and 10-year U.S. Treasury futures; the performance of the S&P 500 over the next 6 and 12 months; future dollar-to-euro, dollar-to-yen, and dollar-to-pound-sterling exchange rates; and future prices for commodities such as gold, silver, corn, soybeans, real estate, and wheat. Here, we focus on the expectations of five-year-ahead CPI inflation.

Researchers at the Minneapolis Fed calculate the mean, standard deviation, skewness and kurtosis of the expectations of five-year-ahead CPI inflation. They estimate the market-based probability distributions using price quotes for inflation caps and floors. By following the procedure described in Section 2.3, they create several butterfly portfolios centered on a wide range of strike prices and derive the distribution. Below, we show how the distribution's mean, volatility, and skewness have changed over time.

3.2.1 Mean

The average five-year inflation expectations obtained from the market-based probability distribution is very similar to inflation expectations measured by alternative sources or methodologies. Figure 5 shows the average expected inflation rate calculated by the Minneapolis Fed using market-based probabilities (blue line) alongside alternative measures. The green line shows the five-year break-even inflation rate, which is the difference between the yields of a conventional five-year Treasury note and a five-year TIPS. It is called the break-even rate because it is the inflation rate that would have to be observed for the real yield of a conventional security to equal the yield of the inflation-indexed security. The daily data for the break-even rate from January 4, 2010, through December 31, 2024, have been obtained from FRED (Federal Reserve Economic Data), maintained by the Federal Reserve Bank of St. Louis.

The purple line shows the results of the Federal Reserve Bank of Cleveland's model of inflation expectations. This series is obtained from the Cleveland Fed's Inflation Expectations webpage. The red line shows the implied expectation according to zero-coupon five-year inflation swaps, which is the implied fixed leg of the swap given current investor expectations of inflation. The daily data from January 4, 2010, through December 31, 2024, come from Bloomberg using the last price of the security trading under the ticker "USSWIT5 Curncy."

Prior to 2020, inflation expectations according to the market-based probabilities hovered around 2 percent, suggesting a stable outlook for inflation over the coming five years. However, in early 2020 at the height of the COVID-19 recession, the expected average annual rate of inflation turned sharply down, before fully retracing its decline by the end of the year. Expected inflation continued to climb through 2021 and early 2022 but peaked around the same time that observed inflation began to cool. In the time since early 2023, expectations have returned to near 2 percent but are still slightly elevated compared with where they were pre-COVID.

The expectations measured by the break-even rate and inflation swaps have closely followed that same pattern. This is not surprising since the cash flows of a five-year inflation swap contract can be replicated using the same TIPS and conventional securities as used in calculating the break-even rate. As noted by Fleming and Sporn (2013), the break-even rate and the swap rate should be equal in the absence of market frictions, though in reality they are typically not. Likewise, an inflation cap is a derivative of the same underlying Treasury

Figure 5
Average Five-Year Inflation Expectations

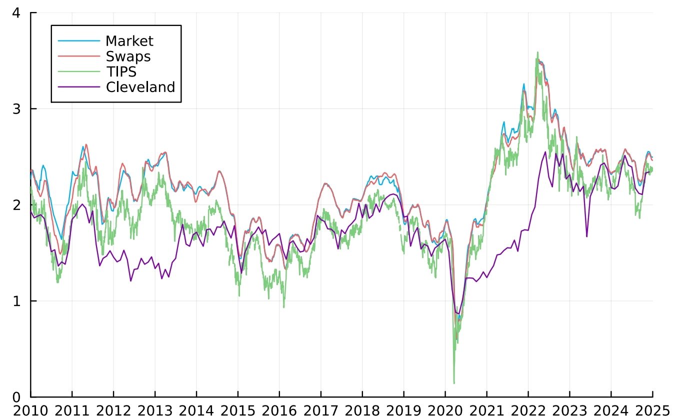

SOURCE: Federal Reserve Bank of Minneapolis, Federal Reserve Bank of Cleveland, U.S. Treasury, Federal Reserve Economic Data, Bloomberg. Notes: Daily and weekly data have been aggregated to a monthly frequency by averaging and are from January 2010 through December 2024.

securities. Accordingly, all three measures of inflation expectations should be reasonably close to each other.

The expectations provided by the Cleveland Fed are also similar but differ substantially at a few points. Most notably, expectations were slower to increase after 2020 than alternative measures but have since fallen back in line. The Cleveland Fed incorporates an array of data into its model of inflation expectations, including Blue Chip forecasts of CPI, current CPI, previous vintages of CPI, inflation expectations measured by the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters, and Treasury yields—all in addition to inflation swaps. This modeling approach explains why these inflation expectations can differ substantially from caps, swaps, and the break-even rate even though all four measures broadly move in the same direction.

3.2.2 Uncertainty

The standard deviation of inflation expectations obtained from the market-based probabilities, or the degree of uncertainty surrounding future inflation, is also consistent with other measures of volatility. Figure 6 presents the monthly averages of the weekly standard deviations obtained from the Minneapolis Fed's market-based probability distributions. We compare them with alternative measures of volatility: (1) the Bloomberg measure of volatility for a five-year 2 percent zero-coupon inflation cap, as reported by Bloomberg Inflation Bond/Swap Settings on the last business day of each month; (2) the monthly one- and three-year-ahead inflation uncertainty measured by the New York Fed's Survey of Consumer Expectations (SCE); and (3) the interquartile range of inflation expectations provided by the New York Fed's Survey of Primary Dealers (SPD), conducted in advance of each meeting of the Federal Open Market Committee, typically scheduled every six weeks.

For months in which a survey was not conducted and for which an observation is not available, we use the average of values of the preceding and following months. Each measure has been indexed to its value during January 2020 so that the figure presents volatility from January 2015 through December 2024, relative to the volatility of January 2020. Months for which data are not available have been populated with the averages of values of the preceding and following months.

The figure shows that the volatility measure of the market-based probabilities provided by the Minneapolis Fed (green line in Figure 6) had been broadly stable from 2015 through 2019. Relative volatility in 2015 had hovered between 1.50 and 1.70, though it slowly decreased beginning in 2016 and generally stayed below 1.30 after that. However, with the onset of the pandemic by the middle of 2020, relative volatility had jumped from 1 to 1.30 and more or less continued to increase (albeit unevenly) beyond that point. By July of 2022, at approximately the same time as when inflation had peaked, it had reached an average of 1.80. By the end of 2022, it had jumped even higher to above 2.30, where it has stayed ever since.

Most other measures of volatility had followed a near-identical pattern up until 2022. Relative one- and three-year-ahead inflation uncertainty as measured in the SCE started 2015 slightly higher (between 1.20 and 1.50) but came down gradually all the way through 2019. Volatility reported by Bloomberg had also decreased from 2015 through 2019. All three measures then began to climb in 2020 and had roughly doubled

Figure 6
Volatility of Inflation Expectations

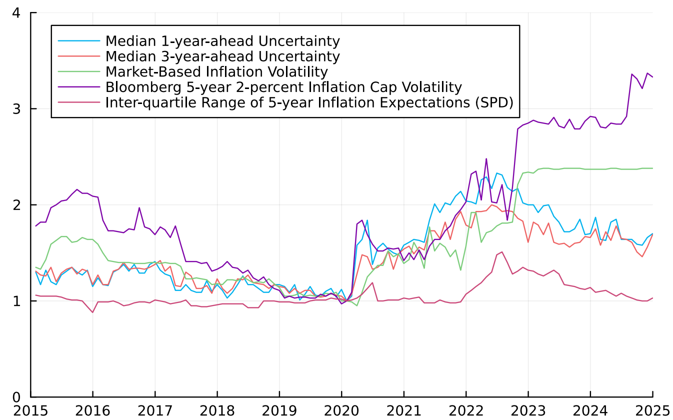

SOURCE: Data are from January 2015 through December 2024. Notes: Weekly “Market-Based Inflation Volatility” is provided by the Federal Reserve Bank of Minneapolis, and data have been aggregated to a monthly frequency. “Inflation Cap Volatility” is as reported by Bloomberg Inflation Bond/Swap Settings on the last business day of each month for a zero-coupon 5-year 2 percent inflation cap. Median “1-year-ahead Uncertainty” and “3-year-ahead Uncertainty” is obtained from the Federal Reserve Bank of New York’s SCE, where uncertainty is the interquartile range of median expectations across all respondents. “Inter-quartile Range of 5-year Inflation Expectations” is obtained from the Federal Reserve Bank of New York’s SPD. Months for which data are not available have been populated with the averages of the preceding and following series values.

by mid-2022. However, the interquartile range of expectations reported in the SPD behaved differently: It was essentially constant from 2015 through 2019, and even through 2021. Only in early 2022 did it begin to climb, and to a far lesser degree. By mid-2022, relative uncertainty measured by the SPD was only approximately 1.50.

The behavior of each series has been considerably more diverse since late 2022. The relative volatility reported by Bloomberg continued to increase and has been above 3.25 as of December 2024. Conversely, the uncertainty measured in the SCE and SPD have decreased substantially. In the SCE, one- and three-year-ahead inflation uncertainty have each recently fallen to 1.70 in December 2024, from 2.00 and 2.30 in June 2022, respectively. While this is, of course, significantly below the peak volatility levels, expectations are still more volatile than at any point from 2015 to 2019. Last, uncertainty expressed in the SPD has retreated from its June 2022 peak back to pre-pandemic levels.

3.2.3 Tail Risks

The final moment we discuss is the skewness of the inflation market-based probability distribution. As shown in Figure 7, the distribution of five-year expectations was generally symmetric in the several years leading up to 2020. In the mid-2010s, the skewness was always less than 1. Then, beginning in 2016, it was almost always negative, suggesting greater risks of low inflation or inflation in the left tail of the distribution.

Though the skewness was firmly negative in the few years leading up to 2020, it quickly increased with the onset of the pandemic. The risks of inflation in the low end of the distribution were suddenly replaced with risks of high inflation. However, that increase in skewness was far from steady. In line with increased inflation volatility discussed previously, the skewness oscillated wildly over 2020 through early 2022. At inflation’s peak, the skewness of the market-based probability distribution was 3, indicating significant perceived risk of high (or right-tail) inflation. Skewness has come down after inflation came off its peak but has been above 1 since late 2022, suggesting that markets still perceive a greater risk of inflation falling in the right tail of the expectations distribution.

3.3 How Reliable Are Market-Based Probabilities?

Next, we discuss the advantages and disadvantages of taking the market-based approach to measuring risk-neutral probabilities of inflation. We first compare this method with density forecasts of inflation and then outline the competing views regarding the reliability of the market data underlying the risk-neutral probabilities.

Figure 7
Skewness of the Market-Based Probability Distribution of Inflation

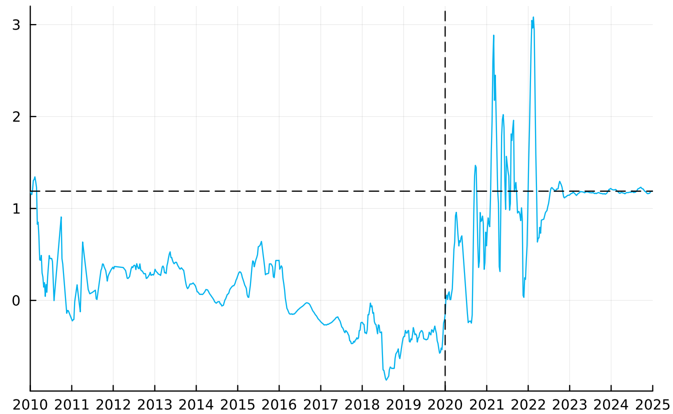

SOURCE: Minneapolis Fed. Data are from January 7, 2010, to December 18, 2024.

3.3.1 Comparison with Density Forecasts

Density forecasting is one alternative approach to creating a distribution of possible future outcomes. It has been used extensively by McCracken, Owyang, and Sekhposyan (2021) and Amburgey and McCracken (2024), among others, for real-time scenario analysis and for understanding the role of financial conditions in the growth of consumption, investment, government spending, and exports. Density forecasts differ from conventional forecasts in that the latter provide a single point estimate of some future value, while the former provide a distribution of possible values. However, density forecasts are similar to conventional forecasts in several ways and therefore share some of the strengths and weaknesses of conventional forecasts compared with market-based expectations.

A notable advantage of density forecasts over extracting distributions from asset prices is that density forecasts can be created for a wide variety of indicators, whereas analysts are limited to creating market-based probability distributions of only those indicators that are clearly tied to financial assets. For instance, sufficient data are available only for five-year inflation caps and floors contracts and not for one-year contracts, so it is much more difficult to create probability distributions over extended periods of time for one-year inflation with the methodology that we use. However, an analyst can create a density forecast of any indicator with available data and supplement it with data of various related predictors. In other words, density forecasts are more readily available for a host of indicators because suitable input data are more readily available.

However, data availability does not imply reliability. In fact, because conventional and density forecasts rely heavily on available historical data, the forecasts may not capture new market dynamics, policy changes, or unprecedented economic events. For example, the relationship between inflation and unemployment seemed to have changed significantly in the 2010s from what had been observed over prior decades (Engemann, 2020; Kuttner and Robinson, 2010; Occhino, 2019). Forecasting models that were trained using data from the 1980s and 1990s may have predicted inflation rates that were much higher than observed, especially as unemployment reached record lows in the years immediately preceding the onset of the COVID-19 pandemic.

Last, forecasts are often delayed, as historical data become available with a lag and analysts must wait for the latest relevant data. For instance, the BLS releases its first estimate of CPI inflation for a given month approximately two weeks after that month has ended. The Bureau of Economic Analysis takes roughly four weeks before its Personal Income and Outlays report shows price inflation as measured by personal consumption expenditures for the previous month.

These issues can be circumvented by using market-based expectations and asset prices. An analyst does not need to rely on historical data to see what markets expect will happen in the future, as that information is fully contained in current asset prices. Additionally, asset prices are typically updated at a high frequency. When investors trade contracts in public markets, the data regarding the trades are available almost immediately. Thus, analysts can quickly understand what market participants expect for the future. And because data are available at a much higher frequency, big shifts in the economic or policy environment will quickly be apparent.

But the market-based approach to measuring expectations also has its shortcomings. Most importantly, while asset prices in theory provide clear signals of what investors expect, in reality these signals can be muddled by market frictions and regulatory changes. Market liquidity is also a source of concern regarding the validity of the expectations measured by the market-based approach. If only a handful of trades are completed in a given period, then the asset prices may not fully represent what broader markets expect to unfold over the coming years. The next section evaluates concerns about the liquidity of inflation caps and floors.

3.3.2 Market Liquidity

There are two primary sources for inflation caps and floors market data. The first source is Bloomberg, which uses (a modified version of) the Black-Scholes model for determining the prices of caps and floors. The key inputs to the model are inflation's mean and volatility, and Bloomberg also uses data obtained from certain contributors. The data from contributors can either be actual, observed market data or indicative data that are model based and not derived from contract trading. The second data source is Fenics Market Data, which also reports either price quotes indicative of what investors might see in markets or actual market data provided by BGC Group.

Though data sources may mix observed and quoted contract prices, it is not clear that price quotes are as useful as observed data. After all, if the data are simply being modeled and there are no market participants, can an analyst actually depend on the fact that the data represent the expectations of market participants? In a speech delivered to the U.S. Treasury Market Conference, New York Fed President John Williams highlighted this difference between quoted prices and realized prices of inflation caps and floors: "Based on our market contacts and public reporting of derivatives transactions, these [price data] aren't data at all. There have been no trades reported in the U.S. inflation options market since early 2021. The so-called data that people are citing are generated by a model, not from investors putting real money on the line as is frequently claimed" (Williams, 2023).

Other economists argue that the data are still reasonably connected to markets. Even though the available data may best be interpreted as a set of model-derived price quotes, the contract prices are linked to inflation-protected and conventional Treasury securities. Therefore, the pricing dynamics that play out in the Treasuries market can be reasonably assumed to also be reflected in the modeled prices of the caps and floors.

As shown in Figures 5 and 6, the inflation expectations and volatility that we observe from the market-based probability distribution closely align with alternative measures of expectations and volatility. This does lend some additional credibility to our market-based probability distributions.

4. SIMPLE ESTIMATION OF DISTRIBUTION WITH TAIL RISK

In this section, we use a simple parametric distribution to use the reported data on mean, standard deviation, and skewness and estimate market-based probabilities of future inflation. We then present our estimated distribution of inflation rates along with other measures of robustness.

4.1 Estimation

To obtain our probability distribution, we employ a modified standard normal distribution. Specifically, we follow Backus, Foresi, and Wu (2004) and use a Gram-Charlier expansion of the normal distribution function. A Gram-Charlier expansion is composed of the standard normal density function along with additional terms to correct for the effects of skewness. This expansion generates a density function while allowing for nonzero skewness.

We follow this procedure as outlined in the Chicago Board Options Exchange's (CBOE) white paper on its volatility index (Chicago Board Options Exchange, 2010). The CBOE estimates the distribution of the S&P 500 using S&P 500 out-of-the-money option prices. The need for this measure emerged from the 1987 financial crash. After the crash, investors in the S&P 500 out-of-the-money option prices had a greater expectation of large and abrupt declines in the S&P 500. Analogous to our economic measure of interest (inflation), the expected distribution of the S&P 500 gained significant tail risk and skewness.

We standardize the five-year CPI inflation by subtracting its mean and dividing the difference by its standard deviation, as given by $\omega = \frac{\pi-\mu}{\sigma}$. We define the density for the standardized variable, $f(\omega)$, as shown below:

$$f(\omega) = \underbrace{\varphi(\omega)}_{\text{standard normal density}} \cdot \underbrace{\left(1 - \gamma\frac{3\omega - \omega^3}{6}\right)}_{\text{Correction for skewness}}, \tag{5}$$

Figure 8
Distribution Results

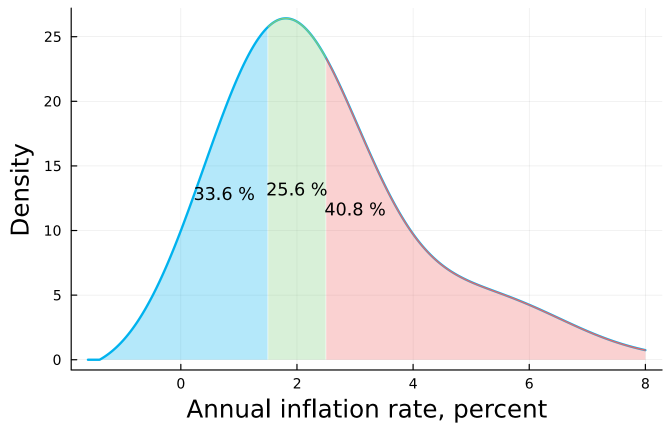

SOURCE: Minneapolis Fed and authors' calculation. The distribution is calculated using data for December 18, 2024.

where $\varphi(\omega)$ is the standard normal density.

We need three moments to compute the equation above: mean μ, standard deviation σ, and skewness γ. We use the estimates from the Federal Reserve Bank of Minneapolis's Current and Historical Market-Based Probabilities discussed in Section 3.2.

4.2 Results

The approximated distribution according to the density function $f(\omega)$ is shown in Figure 8 for December 18, 2024. The blue-shaded portion of the distribution represents the area where inflation is below 1.5 percent, the green-shaded portion represents inflation between 1.5 and 2.5 percent, and the red-shaded portion represents the area in which inflation is above 2.5 percent.

The distribution is heavily skewed to the right, as shown by its long right tail. This is because the skewness is positive, reflecting persistent market fears of inflation being above the target. The estimation implies that there is a 40.8 percent probability that inflation will be above 2.5 percent, while there is a 33.6 percent probability that it will fall below 1.5 percent. There is only about a 25 percent chance that inflation will fall within the range of 1.5–2.5 percent over the next five years.

We estimate the inflation distributions from January 7, 2010, to December 18, 2024, using historical data from the Federal Reserve Bank of Minneapolis. Figure 9 presents the area of our distributions given a date and threshold for the inflation rate. Panel A shows the probability of low inflation: $\pi \leq 0.5$ and $\pi \in [0.5, 1.5]$. Panel B shows the probability that inflation will fall within the Federal Reserve's target range, $\pi \in [1.5, 2.5]$. Panel C shows the probabilities of high inflation: $\pi \in [2.5, 3.5]$ and $\pi \geq 3.5$.

Figure 9
Probability Bins

SOURCE: Minneapolis Fed and authors' calculation. Notes: Data begin on January 7, 2010, and end December 18, 2024. Data are updated weekly.

Figure 10
Expected versus Realized

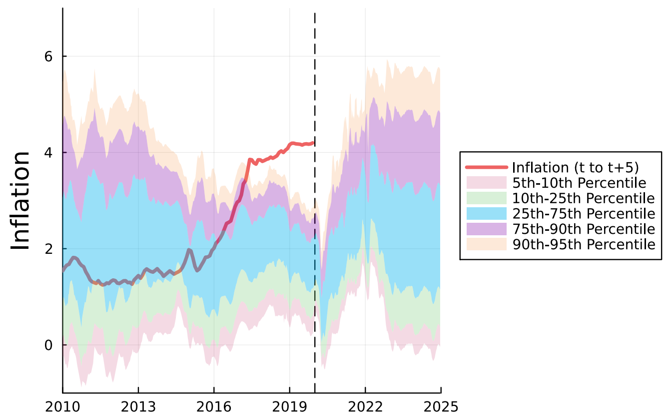

SOURCE: BLS, Minneapolis Fed, and authors' calculation. Notes: December 18, 2024, is the last observation.

As seen in Panel B, there is a general upward trend in the probability that inflation would be on target until the pandemic hit, implying that investors increasingly expected inflation to fall within the target range from 2010 to 2020. The probability of inflation falling between 1.5 and 2.5 percent dives by approximately 40 percentage points around the onset of the COVID-19 pandemic in March 2020. This steep drop in the expectation that inflation would fall in the target range indicates investors' predictions that inflation would plummet as unemployment skyrocketed. This is also reflected in Panel A, as the probability of low inflation climbed dramatically. Panel B also reveals the resulting volatility of expectations following the initial shock of the pandemic to the global economy. Moreover, our probability distributions were much wider during this time of uncertainty. A

4.3 Expectations versus Realizations

We compare the historical risk-neutral distributions to realized CPI inflation provided by the BLS. We obtain a five-year-ahead series at time t by taking the fifth root of CPI at time t divided by CPI at time $t + 5$. The red line of Figure 10 represents this five-year-ahead annualized CPI inflation rate. For example, the last data point is for December 2019, representing the annualized price inflation between December 2019 and December 2024. The shaded regions represent the various areas of our inflation distributions for each date under the risk-neutral measure. Specifically, the different regions are certain percentiles from each distribution. Until 2015, the median expected inflation rate is relatively close to the realized inflation rate. After 2015, the gap between expected and observed inflation widens, with inflation outpacing the rates provided by our distributions. Further, after 2017, there is no overlap between the 25th to 95th percentiles and the realized inflation rate.

4.4 Robustness: Comparison with the SCE

To check the robustness of our estimations, we compare our results with the median expected inflation rate provided by respondents in the SCE (discussed in Section 3.2.2). Table 1 shows the mean and uncertainty (interquartile range) across our distributions and the SCE data for months December 2023 through December 2024. The values in the "Market" column are derived from the distributions based on the market for inflation caps and floors. We only show data from this 12-month period to avoid noise in either set of distributions, market or SCE, created by the pandemic.

The mean and interquartile range of the predicted distributions generally increase and decrease along with the values of the SCE data. The data themselves, however, differ significantly, with around a half-percentage-point discrepancy for a given month. Across all months, the SCE estimates greater expected inflation rates than our measure, indicating that our distributions might be underestimating inflation. Our values for uncertainty also fall below the SCE predictions across all months, indicating that our probability distributions might be narrower than the true distributions, though the values are generally consistent with those provided by the SCE.

Table 1
SCE versus Market Inflation Expectations

		Mean		Uncertainty	
		Market	SCE	Market	SCE
2023	December	2.35	2.54	2.15	3.12
2024	January	2.33	2.54	2.14	3.25
2024	February	2.38	2.89	2.15	3.20
2024	March	2.45	2.62	2.15	3.14
2024	April	2.52	2.82	2.16	3.23
2024	May	2.55	3	2.16	3.09
2024	June	2.49	2.83	2.15	3.29
2024	July	2.45	2.81	2.15	3.26
2024	August	2.31	2.79	2.14	2.73
2024	September	2.21	2.86	2.14	3.03
2024	October	2.37	2.77	2.15	2.86
2024	November	2.54	2.86	2.16	3.01
2024	December	2.52	2.72	2.16	2.78

SOURCE: Minneapolis Fed and authors' calculation, SCE. Notes: December 2024 is the last observation.

5. CONCLUSION

This article highlights the value of market-based probabilities in understanding long-run inflation expectations, offering insights into the dynamics of inflation risks and uncertainties. By analyzing inflation caps and floors, we demonstrate how these instruments capture shifts in market sentiment, particularly during periods of heightened economic uncertainty, such as the COVID-19 pandemic. The results show persistent right-tail risks, reflecting ongoing concerns about elevated inflation levels. While market-based measures provide timely and nuanced perspectives, their reliability depends on market liquidity and data quality. Future research should explore combining market-based probabilities with alternative forecasting methods to enhance robustness and policy-relevant insights.

REFERENCES

Amburgey, Aaron J., and Michael W. McCracken. 2024. Growth-at-risk is investment-at-risk. *Federal Reserve Bank of St. Louis Working Paper Series,* 2023–020B. https://fedinprint.org/item/fedlwp/96594/98794.

Backus, David K, Silverio Foresi, and Liuren Wu. 2004. Accounting for biases in black-scholes. *Available at SSRN,* 585623. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=585623.

Breeden, Douglas T, and Robert H Litzenberger. 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51 (4): 621–651.

Chicago Board Options Exchange, Inc. 2010. Skew white paper. *White Paper Series,* https://cdn.cboe.com/resources/indices/documents/SKEWwhitepaperjan2011.pdf.

Chipeniuk, Karsten O, and Todd B Walker. 2021. Forward inflation expectations: evidence from inflation caps and floors. *Journal of Macroeconomics* 70:103348.

Engemann, Kristie M. 2020. What's the phillips curve (and why has it flattened). *Federal Reserve Bank of St. Louis: Open Vault Blog,* https://www.stlouisfed.org/open-vault/2020/january/what-is-phillips-curve-why-flattened.

Feldman, Ron, Ken Heinecke, Narayana Kocherlakota, Sam Schulhofer-Wohl, and Tom Tallarini. 2015. Market-based probabilities: a tool for policymakers. *Unpublished Manuscript,* https://www.minneapolisfed.org/-/media/files/banking/mpd/optimal_outlooks_dec22.pdf.

Fleming, Michael J., and John R. Sporn. 2013. Trading activity and price transparency in the inflation swap market. *Federal Reserve Bank of New York Economic Policy Review* 19 (1): 45–57.

Hilscher, Jens, Alon Raviv, and Ricardo Reis. 2022. How likely is an inflation disaster? *Available at SSRN,* 4083404. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4083404.

Kitsul, Yuriy, and Jonathan H. Wright. 2012. The economics of options-implied inflation probability density functions. *NBER Working Paper Series,* 18195. http://www.nber.org/papers/w18195.

———. 2013. The economics of options-implied inflation probability density functions. *Journal of Financial Economics* 110 (3): 696–711.

Kuttner, Ken, and Tim Robinson. 2010. Understanding the flattening phillips curve. *The North American Journal of Economics and Finance* 21 (2): 110–125.

McCracken, Michael W., Michael T. Owyang, and Tatevik Sekhposyan. 2021. Real-time forecasting and scenario analysis using a large mixed-frequency bayesian var. *International Journal of Central Banking* 18 (5): 327–367.

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 (8): 2473–2505.

Occhino, Filippo. 2019. The flattening of the phillips curve: policy implications depend on the cause. *Federal Reserve Bank of Cleveland Economic Commentary* 2019 (11).

Ross, Stephen A. 1976. Options and efficiency. *The Quarterly Journal of Economics* 90 (1): 75–89.

Williams, John C. 2023. Elementary, dear data. *Remarks at 2023 U.S. Treasury Market Conference, Federal Reserve Bank of New York,* https://www.newyorkfed.org/newsevents/speeches/2023/wil231116.