
# Conditional Risk-Neutral Densities Around FOMC Meetings

Gianmarco Cinquini[†] Olof Nordin[‡]

Master Thesis

Stockholm School of Economics

Department of Finance

May 18, 2026

## Abstract

We combine Polymarket contract prices with US 2-year Treasury futures options to identify conditional risk-neutral distributions around FOMC meetings. While short-term rate decisions are typically well anticipated, September 2024 presented genuine uncertainty because market pricing remained split between a 25 basis point and 50+ basis point cut until late in the pre-meeting window. The decomposition recovers economically meaningful conditional components not available from either source alone. The larger-cut state implies higher futures prices for most of the pre-meeting window, with a mean separation exceeding what a single overnight-rate difference would suggest. This indicates that the two scenarios priced different expectations for the subsequent policy path. The larger-cut state also has relatively higher conditional volatility across most of the pre-meeting window, suggesting higher uncertainty about the implications of a larger cut.

**Keywords:** Treasury futures, options, FOMC, prediction markets, risk-neutral, RND.

**Acknowledgements:** We are grateful to our supervisor, Tobias Sichert, for his guidance throughout the process of writing the thesis.

[†]42900@student.hhs.se
[‡]42892@student.hhs.se

# Contents

# List of Figures



# List of Tables

# 1 Introduction

The eight annual meetings of the Federal Open Market Committee (FOMC) may be among the most closely followed events in financial markets. At these meetings, the Federal Reserve announces its target-rate decision and communicates information about the likely future path of monetary policy. Treasury yields can therefore move not only because the current policy rate changes, but also because investors revise their expectations about subsequent decisions.

Before each meeting, two different pieces of information affect these prices. The first is the probability of each policy outcome. The second is the distribution of prices conditional on that outcome. Prediction markets price the first. Option surfaces reflect the second, but aggregated across all possible outcomes. Neither source alone separates event probabilities from conditional pricing distributions. In September 2024, for example, the market debated between a 25 and 50 basis point cut. Polymarket contracts reflected two-sided uncertainty through most of the pre-meeting window, and the Treasury option surface embedded the pricing consequences of both scenarios, but no single instrument revealed how 2-year futures would be distributed under each outcome separately.

We ask whether combining these two data sources can recover conditional risk-neutral distributions around FOMC meetings. We follow the event-mixture framework of Hanke et al. (2018), in which prediction-market probabilities are fixed as external mixture weights and option prices identify the conditional distributions under each policy outcome. Hanke et al. (2018) apply this framework to exchange rates around political events such as Brexit and the 2016 US presidential election. Branger et al. (2024) study a related event-mixture setting for wheat futures during the Ukrainian war, but without the same external prediction-market identification strategy. In both cases, however, the event has a more direct effect on the asset price. In the FOMC setting, the transmission is less direct. A rate decision affects Treasury futures not only through the overnight rate change but also through what it signals about the future policy path.

We estimate the model for 14 FOMC meetings from June 2024 to January 2026. In our setting, the framework is most informative when prediction-market probabilities remain genuinely two-sided. September 2024 is the clearest case in the sample. For most of the pre-meeting window, the larger-cut state is associated with higher futures prices, and vice versa. The gap is larger than what a mechanical 25 basis point rate difference would suggest, which points to differences in the expected policy path after the meeting. The larger-cut state also has higher conditional volatility, indicating higher uncertainty about what a 50+ basis point cut would signal.

The remainder of the thesis is organised as follows. Chapter 2 reviews the relevant literature. Chapter 3 presents the theoretical framework. Chapter 4 describes the data and empirical implementation. Chapter 5 presents the results. Chapter 6 concludes.

# 2 Literature Review

The effect of monetary policy decisions on financial markets has been studied extensively. Cook and Hahn (1989) document a positive relationship between changes in the federal funds rate target and Treasury yields in the 1970s, with the response stronger at shorter maturities and declining along the longer end of the yield curve. Kuttner (2001) replicates this analysis for the period 1989–2000 and finds that the relationship had weakened significantly. Using federal funds futures to decompose target changes into anticipated and unanticipated components, he shows that the market response to unanticipated outcomes is large across all maturities, while the effect of the anticipated component is small. Gürkaynak et al. (2005) identify a second component beyond the surprise in the current target rate, namely surprises in the expected future rate path. They show that FOMC statements explain the majority of the variation in longer-maturity yields around announcements, which the current target surprise alone cannot account for.

While these papers describe how policy actions affect realised yields, Beber and Brandt (2006) take an ex-ante perspective and extract option-implied state-price densities from Treasury bond futures options around scheduled macroeconomic announcements. The authors show how the state-price densities change in response to scheduled macroeconomic announcements. The theoretical link between option prices and implied probability distributions, however, is provided by Breeden and Litzenberger (1978), who show that the risk-neutral density of the underlying at maturity can be recovered from the second derivative of the call price function with respect to strike. In practice, this result cannot be applied directly since option prices are observed only at a finite number of strikes and are affected by market noise, so some interpolation is needed before the density can be recovered. Furthermore, the implied density can exhibit bimodality (see e.g., Alexiou et al., 2025), which a single lognormal cannot accommodate and which motivates richer parametric structures for the risk neutral density.

One way to impose this structure is to model the option-implied distribution as a mixture of lognormal densities. Melick and Thomas (1997) apply a mixture of three lognormal densities to crude oil options during the Gulf Crisis, showing how such models can capture non-lognormal features in option-implied distributions. Before an event, the unconditional distribution combines several conditional outcomes, and the mixture provides a parametric structure that can separate them.

The main difficulty is identification. When both mixture weights and conditional distribution parameters are estimated from option prices alone, the decomposition is not identified; different combinations of weights and conditional moments can be observationally equivalent on the option surface while implying different economic interpretations. Leahy and Thomas (1996) and Branger et al. (2024) encounter this problem respectively when applying a mixture of lognormal densities to options around the Quebec sovereignty referendum and wheat derivatives during the Ukrainian war. Hanke et al. (2018) resolve it by using betting-market probabilities as externally fixed mixture weights and estimating the conditional distributions from option prices alone. The probability of the event is

identified outside the option surface, while the option surface identifies the market-implied consequences of each outcome.

Hanke et al. (2018) apply the framework to FX options around Brexit and the 2016 U.S. presidential election. They note that FOMC meetings are a natural extension, but that sufficiently liquid betting markets for such outcomes were not available at the time. The recent growth of prediction markets for FOMC decisions removes this constraint. Prediction-market prices can be interpreted, with the usual caveats, as a proxy for the market-implied probabilities of discrete outcomes, as discussed more generally by Wolfers and Zitzewitz (2004). The application to FOMC meetings has begun to attract attention. Diercks et al. (2026) compare prediction-market distributions with option-implied distributions around FOMC meetings and document that the two information sources contain related but distinct content. The conditional decomposition of the option-implied density under each policy outcome remains an open question in that work.

We follow the framework of Hanke et al. (2018) and apply it to U.S. Treasury futures options around FOMC meetings, with Polymarket contract prices used as externally observed mixture weights and the Treasury option surface used to estimate the conditional futures-price distributions associated with each policy outcome. We contribute to the literature in two ways. First, we apply the framework of Hanke et al. (2018) to US Treasury note futures, an asset class to which, to our knowledge, it has not previously been applied. Second, we use probabilities from Polymarket, a data source that only recently has become available at sufficient liquidity.

# 3 Theoretical Framework

## 3.1 Risk-Neutral Valuation and Option-Implied Densities

In standard derivatives pricing, European options are valued as discounted expected payoffs under a risk-neutral probability measure $\mathbb{Q}$ (Hull, 2021). The price of a derivative is

$$V_t = e^{-r\tau}\mathbb{E}_t^{\mathbb{Q}}\left[h(F_T)\right] \tag{3.1}$$

where $h(F_T)$ is the terminal payoff, $\tau = T - t$ is the time to expiry, and the expectation is taken with respect to the risk-neutral density $q(F_T)$. The measure $\mathbb{Q}$ is a pricing measure rather than a forecast; risk is priced through adjusted probabilities rather than through higher discount rates (Hull, 2021). The risk-neutral density therefore represents the market's risk-adjusted pricing of future outcomes, not a direct estimate of the physical distribution. A finding that one conditional expected price exceeds another should accordingly be read as a statement about risk-neutral pricing, not as a forecast of the realised futures price.

For a European call option written on a futures contract with current price $F$, strike $K$, and time to expiry $\tau$, the risk-neutral price is

$$C(K) = D\int_K^{\infty}(F_T - K)\,q(F_T)\,dF_T \tag{3.2}$$

where $D$ is the discount factor over the option horizon. The call price is an integral over the right tail of the risk-neutral density, weighted by the distance from the strike. Options at different strikes effectively sample different portions of that tail, which is why a cross-section of option prices is informative about the full shape of $q$.

Breeden and Litzenberger (1978) show that differentiating Equation (3.2) twice with respect to $K$ recovers the risk-neutral density,

$$q(K) = \frac{1}{D}\frac{\partial^2 C(K)}{\partial K^2} \tag{3.3}$$

The result can be interpreted as the limiting price of a local state claim, since a butterfly spread around strike $K$ approximates a payoff concentrated near that strike (Bahra, 1997).

In practice, Equation (3.3) cannot be applied mechanically. Options are quoted at a finite number of strikes, and observed prices contain bid-ask spreads and other noise. Direct numerical second differences of observed call prices typically produce negative values or large oscillations, particularly in the tails where liquidity is thin. Some smoothing and interpolation is usually needed, where common choices in the literature are polynomials and spline based methods (see e.g., Birru and Figlewski, 2012; Shimko, 1993; Malz, 1997).

## 3.2 Mixture Densities Around Discrete Events

The risk-neutral density recovered from Equation (3.3) describes the full distribution of prices at expiry, reflecting all possible outcomes, but it is not the most informative one for the question we ask. The object of interest is not the overall distribution of prices after an event, but rather the conditional distribution under each outcome separately. For FOMC, a 25 basis point cut and a 50 basis point cut imply different Treasury futures prices, and the overall risk-neutral density mixes these together with weights equal to the respective event probabilities. Recovering the conditional distributions requires giving the density more structure.

When the event has a finite set of possible outcomes, the unconditional risk-neutral density can be written as a mixture:

$$q(F_T) = \sum_{j=1}^{n} \omega_j \, q_j(F_T) \tag{3.4}$$

where $\omega_j \geq 0$ and $\sum_j \omega_j = 1$, with $\omega_j$ the risk-neutral probability of event state $j$ and $q_j$ the conditional risk-neutral density of $F_T$ given that state $j$ occurs. Each $\omega_j$ describes the event probability, while each $q_j$ describes the distribution of Treasury futures prices conditional on that event state.

Because option prices are linear in densities, this structure passes directly through to option prices. Substituting Equation (3.4) into Equation (3.2) gives

$$C(K) = \sum_{j=1}^{n} \omega_j \, C_j(K) \tag{3.5}$$

where $C_j(K)$ is the price that would prevail if event $j$ were certain ($\omega_j = 1$). The observed option price is therefore a weighted average of conditional option prices, with the same weights as the density mixture. Once the weights are fixed externally, the option surface can be used to estimate the conditional pricing distributions under each outcome.

For the conditional densities $q_j$, we follow Hanke et al. (2018) in using lognormal distributions. A lognormal component is fully characterised by a conditional mean and a conditional volatility, which keeps the number of free parameters manageable while allowing the unconditional mixture to produce skewness, excess kurtosis, and even bimodality. Equivalently, each component is normal in log futures-price returns and lognormal in futures-price levels.

## 3.3 Event-Mixture Model with External Probabilities

The mixture structure in equation (3.5) raises an identification problem. When the mixture weights $\omega_j$ and the conditional parameters are all estimated from option prices alone, the decomposition is not identified from the option surface. Different combinations of weights, conditional means, and volatilities can generate similar option prices, but imply different economic interpretations. Hanke et al. (2018) resolve this problem by taking the mixture weights from an external source rather than from option prices. In their framework, prediction market quotes provide the event probabilities, and option prices are used only to identify the conditional distributions. This separates the two questions the model answers: how likely is each outcome, and what would markets price under each outcome.

Prediction-market contracts pay one dollar if a specified outcome occurs and zero otherwise. Its price is expressed as a fraction of one currency unit and can be interpreted as the market-implied probability of that outcome. More precisely, the price of such a contract can be interpreted as a state price for the event, subject to market frictions and any risk premia embedded in prediction-market prices (Hanke et al., 2018). We use Polymarket contract prices analogously, treating the implied probability of each outcome as an externally observed mixture weight that is fixed before any estimation is performed on the option surface.

For each event, we consider two states. In what follows, we write $\omega$ for the probability of state 1 and $1-\omega$ for the probability of state 2. Furthermore, we write $V$ rather than $C$ because the option-price inputs, constructed in Section 4.2, include both put and call prices across the observed delta quotes. The observed option price for a futures option with strike $K$ is modelled as

$$V^{\text{mix}}(K) = \omega\, V_1(K;\, E_1, \sigma_1) + (1-\omega)\, V_2(K;\, E_2, \sigma_2) \tag{3.6}$$

where $V_j$ denotes the Black (1976) price of the relevant option contract, $\omega$ is the prediction-market probability of state 1 normalised over the two modelled outcomes, $E_j$ is the conditional expected futures price under state $j$, and $\sigma_j$ is the conditional volatility. Following Hanke et al. (2018), the conditional expected futures prices are written as

$$E_j = F \exp\big(\big(\mu_j + \tfrac{1}{2}\sigma_j^2\big)\,\tau\big) \tag{3.7}$$

where $\mu_j$ is the conditional log-drift under state $j$, $F$ is the current futures price, and $\tau$ is time to expiry in years. Then, the following constraint is introduced

$$F = \omega\, E_1 + (1-\omega)\, E_2 \tag{3.8}$$

Given $\omega$, $\mu_1$, $\sigma_1$, and $\sigma_2$, restriction (3.8) determines $\mu_2$, reducing the free parameters from four to three. Since $\mathbb{E}^{\mathbb{Q}}[F_T] = F_t$ (Hull, 2021), this is solved as:

$$\mu_2 = \frac{1}{\tau}\log\left(\frac{1-\omega\exp\big((\mu_1 + \frac{1}{2}\sigma_1^2)\tau\big)}{1-\omega}\right) - \frac{1}{2}\sigma_2^2 \tag{3.9}$$

Following Hanke et al. (2018), the three free parameters $(\mu_1, \sigma_1, \sigma_2)$ are estimated by minimising the sum of squared relative pricing errors across all 11 observed option prices on a given date:

$$\min \sum_{i=1}^{11} \left( \frac{V_i^{\text{model}} - V_i^{\text{mkt}}}{V_i^{\text{mkt}}} \right)^2 \tag{3.10}$$

We give all option quotes equal weight. The use of relative rather than absolute pricing errors follows Hanke et al. (2018), who choose relative errors to avoid implicitly overweighting highly priced options. Without this scaling, the fit would be driven mainly by the highest-priced options near the money, and cheaper out-of-the-money options would receive little weight despite carrying important information about the tails of the distribution.

While the general mixture in equation (3.4) allows for $n$ states, we restrict attention to two. In principle, one could model more outcomes simultaneously, for example a 50 basis point decrease, a 25 basis point decrease, and no change. In practice, this would require estimating additional conditional means and volatilities from only eleven option prices. The prediction-market data also support a two-state specification since for most meetings in the sample, only two outcomes carry non-trivial probability mass over the estimation window. The two state setup therefore reflects both the structure of the available prediction-market data and the need to keep the conditional density estimation stable. Beyond practical constraints, a binary specification is economically meaningful. FOMC decisions are generally made in 25 basis point increments and communication around decisions generally narrows the plausible outcomes to two. We can also see this in the prediction market data (see Figure A.1), where only two outcomes typically have meaningful probability.

Two assumptions are important for how the results should be interpreted. First, Polymarket probabilities are used as mixture weights. This requires prediction-market and option-market participants to price event risk in a broadly consistent way. If the two markets carry different risk premia for the same event, the identifying assumption may not hold exactly (Wolfers & Zitzewitz, 2004). Second, the conditional densities are restricted to be lognormal. This keeps the model tractable, but it may not capture all features of the underlying distribution.

## 3.4 Non-parametric RNDs

To provide a check on the mixture-implied distributions, we also recover an unconditional risk-neutral density directly from the option surface using Breeden and Litzenberger (1978). The purpose is to recover the overall shape of the risk-neutral distribution, against which the mixture-implied unconditional density $\omega q_1 + (1-\omega) q_2$ can be compared.

Because Equation (3.3) requires a smooth, differentiable call price function, direct finite differences on the eleven observed option prices are not sufficient. Shimko (1993) proposes to convert the observed option prices to implied volatilities, fit a smooth curve to the resulting smile, and then convert back to prices before applying Breeden and Litzenberger (1978). Other implementations use more flexible interpolation methods, such as splines, to smooth either the implied volatility curve or the option-price function (Malz, 1997). The common point is that the density is recovered from a fitted pricing curve rather than from raw option quotes. These methods are often called non-parametric or semi-parametric because they avoid imposing a full distributional form, such as a lognormal density.

This type of recovery is most informative in the central region of the distribution, where option quotes are observed and the fitted curve is anchored by market data. The tails, however, are more difficult to analyse. Observed strikes cover only a finite range, so the density recovered inside that range does not by itself describe the full risk-neutral distribution. Some assumptions are needed about the probability mass below the lowest strike and above the highest strike. A simple solution is to extrapolate the implied-volatility smile outside the observed range, but this can impose strong tail behaviour. A more explicit alternative is to recover the central density over the observed range first, and then attach a parametric tail separately.

Several parametric families have been used for the attached tail. Lognormal tails are natural in Black (1976) settings, but they impose a specific shape on the missing region. Extreme-value theory offers more flexible alternatives. The Generalised Extreme Value (GEV) distribution is associated with block maxima, while the Generalised Pareto Distribution (GPD) is the standard model for exceedances beyond a fixed threshold (Coles, 2001). Since the missing tails in an option-implied density are regions beyond strike thresholds, the GPD is a natural candidate when the tail is completed separately from the central fit.

Figlewski (2009) formalises this separation. Rather than treating the extrapolated smile as a single continuous object, the central density is recovered over observed strikes and the tails are attached as a distinct step. Building on this framework, Birru and Figlewski (2012) examine GPD tails as an explicit alternative to constant-volatility extrapolation.

The benchmark is used to check whether the central shape of $\omega q_1 + (1-\omega) q_2$ is consistent with a density recovered directly from the option surface. Tail regions are not the main focus, since both estimates rely on assumptions there rather than on observed quotes. The implementation is described in Section 4.5.

# 4 Data and Estimation

## 4.1 US Treasury Note Futures

The underlying instrument is the US 2-Year Treasury Note future, traded on the Chicago Board of Trade.[1] Each contract represents $200,000 in face value at maturity and is deliverable against an eligible Treasury note from the deliverable basket, with prices quoted relative to a notional 6% coupon via the CME conversion factor system. The contract is highly liquid, with daily trading volumes regularly exceeding one million contracts and open interest of around 4.7 million contracts as of April 2026 (CME Group, 2026). Also, as Hull (2021) notes, the market for Treasury bond futures is considerably more active than the market for any individual Treasury bond.

The 2-year maturity is sensitive to changes in monetary policy expectations. Short- and medium-term Treasury yields respond strongly to unexpected monetary policy actions, while FOMC announcements also affect yields through changes in the expected future policy path (Gürkaynak et al., 2005; Kuttner, 2001). Longer maturities are more sensitive to term premia and longer-run inflation expectations, which may dilute the meeting-specific FOMC signal. Federal funds futures and SOFR futures provide a more direct measure of short-rate expectations, but they are less suitable for the present density decomposition. Because these contracts are closely tied to realised overnight rates, the conditional price response to each FOMC outcome is more mechanical, leaving less residual distributional information for the option surface to identify. Options on SOFR futures also reflect rates over a reference period, making the mapping to a single FOMC meeting less clean. The 2-year futures price, by contrast, reflects expectations over the near-term policy path. Even conditional on a specific rate decision, there is residual uncertainty about subsequent moves, so both the level and the dispersion of the conditional distribution carry economic content.

We download daily implied constant-maturity futures prices from Bloomberg for the 1-month (1M) and 2-month (2M) tenors. Rather than tracking individual contracts, these series are constructed by Bloomberg through interpolation between adjacent delivery contracts, producing a price with a fixed time to expiry on every date. Our sample covers 14 FOMC meetings between June 2024 and January 2026.

For each meeting, we perform estimation over the 30 calendar days preceding the announcement date. We apply this same 30-day event window to both option surfaces. The two surfaces differ in their option horizon (30 calendar days for the 1-month surface and 60 calendar days for the 2-month surface), but the length of the estimation window is identical for both. The 2-month surface is included as a robustness check. We note, however, that as a meeting approaches, this instrument will span the subsequent meeting.

With futures as the underlying instrument, our approach deviates from Hanke et al. (2018), who study options on FX, and stays closer to Branger et al. (2024), who apply

[1] The Chicago Board of Trade (CBOT) is part of CME Group.

the mixture model to commodity futures. This setting, however, still satisfies the three criteria of the Hanke et al. (2018) framework. First, the upcoming event should be anticipated and expected to have an effect on the market prices of the affected asset. Second, there must be a liquid betting market. Third, there are liquid options at multiple strikes.

**Figure 4.1.** 2-year Treasury yield versus the upper bound of the federal funds target range (top), and the spread between the two (bottom). The spread is defined as the 2-year Treasury yield minus the upper bound of the federal funds target range.

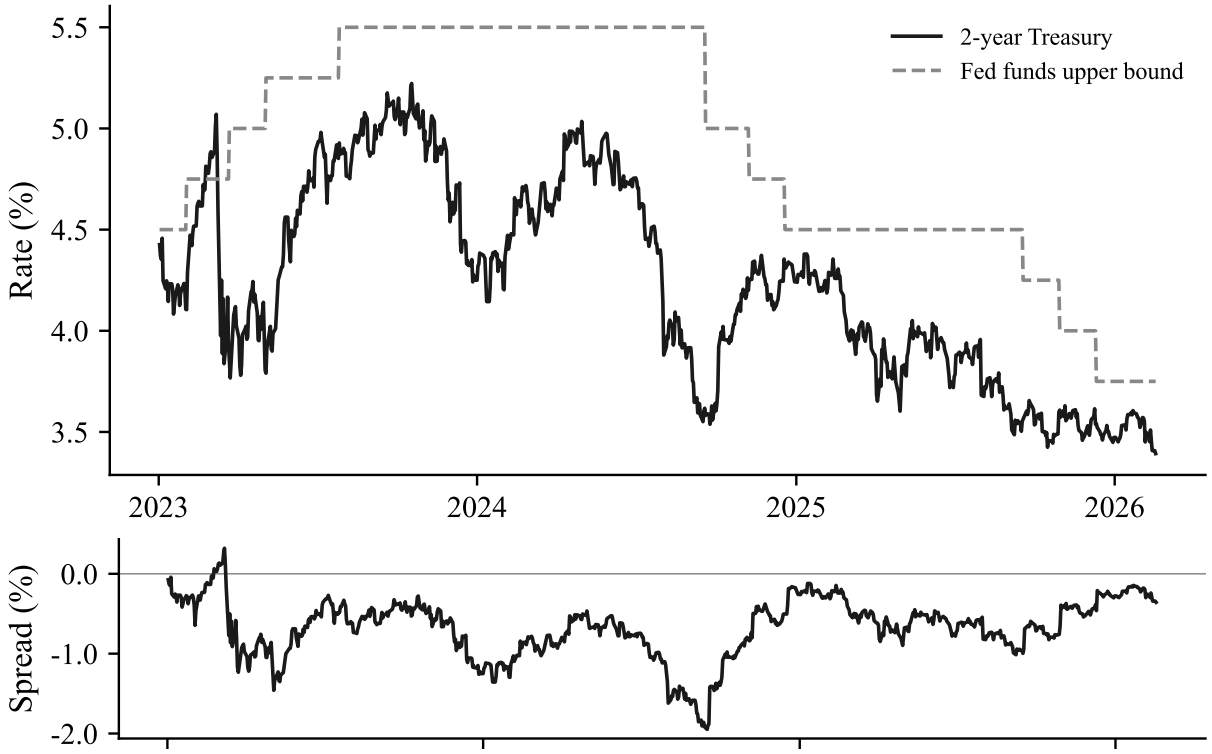


## 4.2 Treasury Futures Options

We download the corresponding 1-month and 2-month constant-maturity implied volatility surfaces from Bloomberg. The 1M surface has a time to expiry of 30 calendar days and the 2M surface has a time to expiry of 60 calendar days. Each surface quotes annualised Black (1976) implied volatilities at eleven delta points: 5DP (delta-put), 10DP, 15DP, 25DP, 35DP, 50D (ATM), 35DC, 25DC, 15DC, 10DC, and 5DC (delta-call).

The listed options on 2-year Treasury futures have an American exercise feature. Bloomberg smooths and converts these into European implied volatilities using a PDE pricer. We therefore use the Bloomberg-reported Black (1976) implied volatilities as the inputs for the European-style pricing formulas below.

We convert each delta quote to a strike under the Black (1976) model. We convert puts to call deltas using $\Delta_C = 1 - \Delta_P$ (Hull, 2021), giving us call deltas of 0.95, 0.90, 0.85, 0.75, 0.65, 0.50, 0.35, 0.25, 0.15, 0.10, and 0.05. Inverting the Black (1976) delta relation gives the corresponding strike:

$$K = F \exp\!\left(-\Phi^{-1}(\Delta_C)\,\sigma\sqrt{\tau} + \tfrac{1}{2}\sigma^2\tau\right) \tag{4.1}$$

where $F$ is the constant-maturity futures price, $\sigma$ the implied volatility at that delta, $\tau$ the time to expiry in years, and $\Phi^{-1}$ the standard normal quantile function.[2]

Once the strike is known, the option price is computed under Black (1976). The five put-wing quotes (5DP through 35DP) are priced as puts and the at-the-money quote and the five call-wing quotes (35DC through 5DC) are priced as calls:

$$\begin{aligned} C &= D\,[F\,N(d_1) - K\,N(d_2)] & P &= D\,[K\,N(-d_2) - F\,N(-d_1)] \\ d_1 &= \frac{\ln(F/K) + \tfrac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}} & d_2 &= d_1 - \sigma\sqrt{\tau} \end{aligned} \tag{4.2}$$

Discounting is applied through the discount factor $D$. We use the USD SOFR OIS swap rate at the matching tenor. The procedure produces up to eleven option prices per date per surface, which serve as inputs to the estimation described in Section 4.4.

**Figure 4.2.** Implied volatility time series for the 1M constant-maturity surface. The solid line shows the at-the-money (50-delta) annualised Black (1976) implied volatility. Dashed and dotted lines show the 10-delta put and 10-delta call implied volatilities, respectively. Vertical lines indicate FOMC meeting dates.

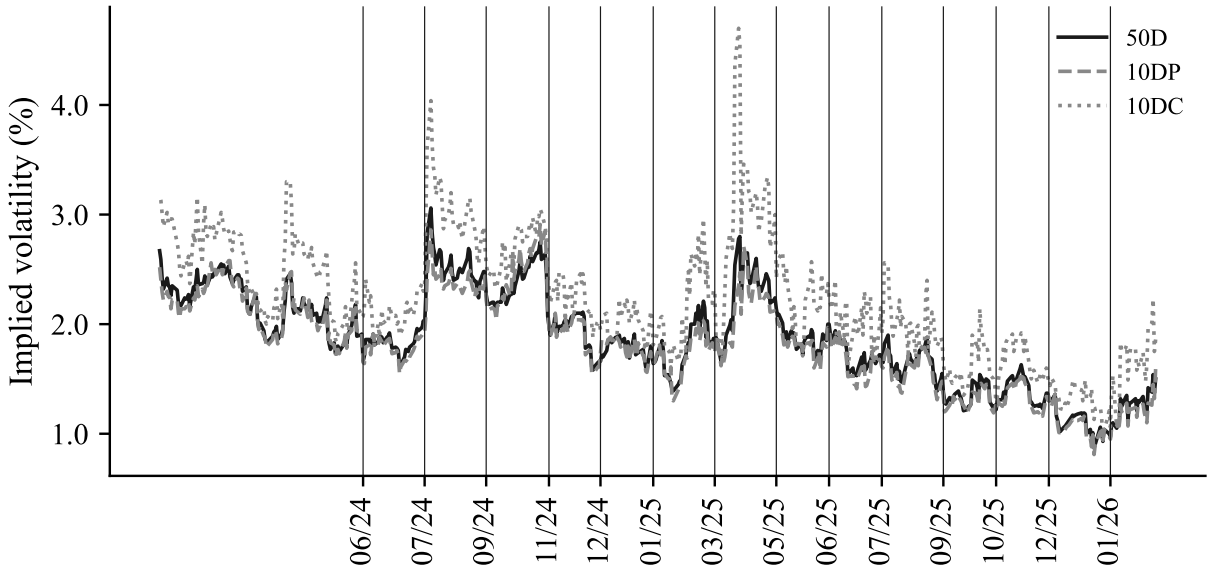


Figure 4.2 shows three features of the implied volatility surface over our sample. The 10-delta call implied volatility is consistently above the 10-delta put, reflecting a skew toward higher futures prices and correspondingly lower yields, consistent with the easing-cycle environment visible in Figure 4.1. The dispersion between the two wings varies over time, with the widest gap occurring before the September 2024 and May 2025 meetings.

[2]We verified this inversion against Bloomberg-reported strikes.

The overall level of implied volatility is lower toward the end of the sample than around the early easing-cycle meetings.

**Figure 4.3.** Implied volatility as a function of delta 30 days prior to September 2024 (left) and May 2025 (right) FOMC meetings. Each line represents the annualised Black (1976) implied volatility across the eleven delta quotes from the Bloomberg 1-month constant-maturity surface.

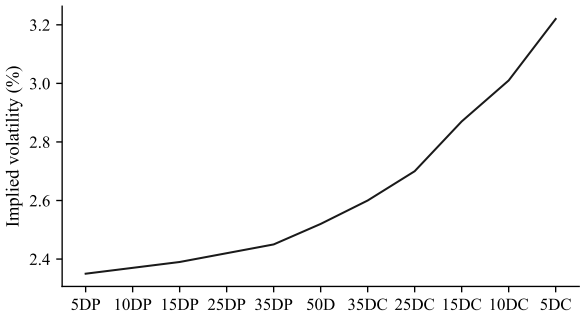


**(a)** September 2024

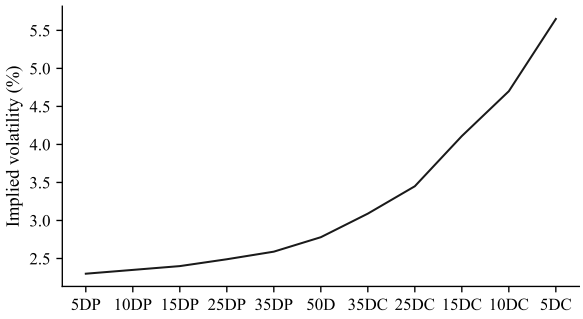


**(b)** May 2025

## 4.3 Prediction-Market Probabilities

We obtain daily prediction-market prices for FOMC rate decisions from Polymarket. For each FOMC meeting in our sample, Polymarket lists contracts corresponding to specific rate outcomes, such as "25 basis point decrease", "50 basis point decrease", and "No Change". Each contract pays one dollar if its stated outcome is realised and zero otherwise.

Our framework requires a binary classification of each meeting into two states: a more dovish state (state 1) and a less dovish state (state 2). For each meeting, we select the two Polymarket outcomes that define this classification. Table 4.1 lists the state assignments for all 14 meetings in the sample. For most meetings, the relevant debate is between a 25 basis point cut and no change. For the September 2024 meeting, state 1 is the Polymarket contract labelled "50+ basis point decrease" and state 2 is the contract labelled "25 basis point decrease", while for the November 2024 meeting, state 1 is the Polymarket contract labelled "50 basis point decrease" and state 2 is the contract labelled "25 basis point decrease".

Let $p_{1,t}$ and $p_{2,t}$ denote the raw Polymarket prices for state 1 and state 2 on date $t$.

$$\omega_t = \frac{p_{1,t}}{p_{1,t} + p_{2,t}} \tag{4.3}$$

The renormalisation in equation (4.3) absorbs any residual probability assigned to outcomes outside the two selected states proportionally into state 1 and state 2. However, usually two outcomes dominate the probabilities, as shown in Figure A.1.

**Table 4.1.** Two-state classification of FOMC prediction-market outcomes. State 1 represents the more dovish outcome and receives mixture weight $\omega_t$. State 2 denotes the more hawkish outcome and receives weight $1 - \omega_t$. Labels are derived from Polymarket contracts.

| FOMC | State 1 | State 2 |
|---|---|---|
| 2024-06-12 | 25 bps decrease | No Change |
| 2024-07-31 | 25 bps decrease | No Change |
| 2024-09-18 | 50+ bps decrease | 25 bps decrease |
| 2024-11-07 | 50 bps decrease | 25 bps decrease |
| 2024-12-18 | 25 bps decrease | No Change |
| 2025-01-29 | 25 bps decrease | No Change |
| 2025-03-19 | 25 bps decrease | No Change |
| 2025-05-07 | 25 bps decrease | No Change |
| 2025-06-18 | 25 bps decrease | No Change |
| 2025-07-30 | 25 bps decrease | No Change |
| 2025-09-17 | 25 bps decrease | No Change |
| 2025-10-29 | 25 bps decrease | No Change |
| 2025-12-10 | 25 bps decrease | No Change |
| 2026-01-28 | 25 bps decrease | No Change |

## 4.4 Estimation Procedure

We estimate the event-mixture model separately for each trading date, FOMC meeting, and constant-maturity surface. The inputs on a given date are the eleven option prices from Section 4.2 and the prediction-market probability $\omega_t$ from Section 4.3. With $\omega_t$ fixed, the free parameters are $(\mu_1, \sigma_1, \sigma_2)$. The estimation criterion is the sum of squared relative pricing errors defined in Equation (3.10).

To reduce sensitivity to initial values, we use a deterministic multi start procedure. The starting values vary the initial drift in both directions around zero and initialise the two conditional volatilities as multiples of the at-the-money implied volatility observed on the same date. We use the Nelder and Mead (1965) method for each starting point and retain the solution with the lowest objective value.

## 4.5 Non-parametric RND Benchmark

As a density benchmark, we recover an unconditional risk-neutral density from the smoothed implied-volatility smile and compare it with the mixture-implied unconditional density $\omega q_1 + (1 - \omega) q_2$. This benchmark does not impose the two-state structure and does not use prediction-market probabilities. The procedure, introduced in Section 3.4, has three parts: polynomial smoothing of the implied-volatility smile, central density recovery using Breeden and Litzenberger (1978), and tail completion using calibrated Generalised Pareto Distribution (GPD) tails.

**Polynomial smoothing.** Using the eleven strike-volatility pairs constructed in Section 4.2, we fit a fourth-degree polynomial to the observed Black (1976) implied volatilities

as a function of strike. The fitted smile is then converted back into Black (1976) prices before the density is recovered. This follows the smoothing logic of Shimko (1993), but uses a fourth-degree polynomial as done by Figlewski (2009).

The polynomial is evaluated on an equally spaced grid of 1,200 strikes spanning the observed strike range. Fitted volatilities are bounded below by a small positive floor and above by 200% as a numerical safeguard against invalid option prices or polynomial blow-up. These bounds are not modelling assumptions about plausible Treasury volatility; they are only used to keep the Black (1976) conversion well defined.

**Observed-strike density.** The density recovered over the observed strike range is then determined by applying Breeden and Litzenberger (1978) in equation (3.3) through the centred finite-difference approximation

$$\hat{q}(K_i) \approx \frac{1}{D} \frac{C(K_{i+1}) - 2C(K_i) + C(K_{i-1})}{(\Delta K)^2} \tag{4.4}$$

Density values that are negative or non-finite because of numerical artefacts are set to zero. The first and last grid points are filled from their nearest interior neighbours for plotting and integration. These copied endpoint values are not used when calibrating the tail slopes.

**Tail masses.** To complete the distribution, we first infer the probability mass outside this range from the slope of the fitted call-price curve at the boundary strikes. Following the logic of Figlewski (2009) and Birru and Figlewski (2012), the call-price slope satisfies

$$\frac{\partial C}{\partial K} = -D \Pr(F_T > K) \tag{4.5}$$

We therefore estimate the left- and right-tail masses as

$$\hat{m}_L = 1 + \frac{C'(K_{\min})}{D}, \qquad \hat{m}_R = -\frac{C'(K_{\max})}{D}, \tag{4.6}$$

where $C'$ denotes the numerical derivative of the fitted call-price curve with respect to strike. These values are clipped only to keep them inside the admissible probability range. If the estimated left and right tail masses do not leave positive probability mass for the central density, the benchmark fit is treated as invalid for that date. This condition does not bind in the reported September estimates.

Once the tail masses have been assigned, the central density is rescaled to integrate to the remaining probability mass,

$$\hat{m}_C = 1 - \hat{m}_L - \hat{m}_R. \tag{4.7}$$

This step is necessary because the central Breeden and Litzenberger (1978) density is recovered only on the observed strike range. Without rescaling, adding parametric tails would double-count probability mass.

**GPD tail completion.** The missing tails are completed using generalised Pareto distributions. Birru and Figlewski (2012) discuss GPD tails as an alternative to extrapolating

constant implied volatility outside the observed strike range. The GPD is also the standard threshold-exceedance model in extreme-value theory (Coles, 2001), which fits the present problem since the missing regions lie beyond the lowest and highest observed strikes.

The GPD tails are calibrated rather than statistically estimated from tail observations. Their total masses are fixed by Equation (4.6). Conditional on those masses, the scale and shape parameters are chosen to match the level and local slope of the recovered central density at each boundary.

Let $\beta_L$ and $\beta_R$ denote the GPD scale parameters for the left and right tails. Since the weighted GPD density at the threshold equals the tail mass divided by the scale parameter, matching the central density level at the boundary gives

$$\beta_L = \frac{\hat{m}_L}{q(K_{\min})}, \qquad \beta_R = \frac{\hat{m}_R}{q(K_{\max})}. \tag{4.8}$$

The shape parameters $\xi_L$ and $\xi_R$ are then calibrated so that the local slope of the GPD tail matches the local slope of the central density at the boundary. In other words, the calibration imposes $f'_{\mathrm{GPD}}(K_b) = q'(K_b)$ at each boundary point $K_b$. The central-density slope is estimated using interior grid points adjacent to the boundary, rather than the copied endpoint values used for plotting.

For stability, the calibrated shape parameters are bounded to $\xi \in [-0.45, 0.49]$. These bounds are not estimated restrictions but are imposed to prevent unstable boundary calibrations and are consistent with standard GPD restrictions: values near or below $-0.5$ are associated with irregular estimation behaviour, while $\xi \geq 0.5$ implies infinite variance (Coles, 2001; Hosking & Wallis, 1987). The lower bound is set slightly above $-0.5$, and the upper bound slightly below 0.5, for numerical safety.

The plotted benchmark density is not renormalised after tail attachment. The left tail, central region, and right tail are instead constructed so that their assigned masses sum to one, and the numerical integral is reported as a mass check in Table A.2.

# 5 Results

## 5.1 Meeting Selection

FOMC meetings have discrete outcomes, but this alone does not tell us which meetings are informative for the mixture model. We begin by estimating an unconditional lognormal model under Black (1976), using a single implied volatility fitted across all eleven option prices. This flat-volatility fit is compared with the two-state mixture model using maximum relative pricing errors. If the option smile reflects two distinct event-contingent outcomes, the mixture model should produce a meaningfully better fit. We also consider the Polymarket-implied probability: when $\omega$ is close to 0 or 1, one component gets negligible weight and its parameters are weakly identified, so the mixture is uninformative regardless of fit quality. We therefore assess each meeting against four criteria: i) a substantial reduction in maximum relative pricing error from the single-lognormal to the mixture model, ii) a genuinely two-sided $\omega$ at the snapshot date, iii) persistence of two-sided uncertainty across the 30-day estimation window (see Figure A.2), and iv) sufficient Polymarket liquidity for the relevant meeting contracts.

As shown in Table 5.1, the single-lognormal fit produces large pricing errors across most meetings, with maximum relative pricing errors typically between 50% and 95%. The mixture model reduces these errors in every case, as expected, since it has more free parameters than the single-lognormal fit. However, for most meetings this reduction coincides with either low Polymarket liquidity (June and July 2024, with volumes below \$3 million), one-sided $\omega$ at the snapshot date (January 2025, March 2025, June 2025), limited persistence of two-sided uncertainty (November 2024 and October 2025), or high maximum relative pricing errors for the mixture model on at least one surface (December 2024, May 2025, July 2025, September 2025, December 2025, January 2026). For November 2024 and October 2025, the Polymarket probabilities in Figure A.2 show that the uncertainty was quickly resolved within the 30-day estimation window, with probabilities turning one-sided shortly after the estimation date.

September 2024 is the only meeting that passes all four criteria. The mixture reduces the maximum relative pricing error from 77.5% to 0.6% on the 1M surface, and from 79.5% to 4.2% on the 2M surface. Table 5.2 also shows a clear separation in the conditional means: the 50+ basis point state is about 0.89 futures price points above the 25 basis point state on the 1M surface, and about 0.90 points above it on the 2M surface.

September 2024 represented the first rate cut in roughly two years, and the uncertainty was not mainly about whether the Fed would cut, but about the size of the cut. Polymarket probabilities for a 50+ basis point cut rose from around 15% in mid-August to around 45% by mid-September. The same uncertainty was visible outside the model; press coverage after the announcement reported that market participants entered the meeting unusually uncertain about the size of the cut, and traders started pricing a 50 basis point cut, rather than the traditional 25 basis point, just days before the meeting (Goldfarb, 2024).

**Table 5.1.** Maximum relative pricing error across FOMC meetings

Table 5.1 reports the maximum relative pricing errors for the lognormal and mixture models across all FOMC meetings in the sample. The volatility surfaces 1M and 2M denote one-month and two-month constant-maturity option surfaces, respectively. Betting volume is Polymarket's total reported trading volume, reported in USD millions. $\omega$ denotes the one-sided risk-neutral probability observed on the corresponding measurement date. Lognormal and mixture-model denote the maximum relative pricing error across the delta quotes.

| FOMC | Volatility surface | Estimation date | Betting volume | $\omega$ | Lognormal | Mixture-model |
|---|---|---|---|---|---|---|
| 2024-06-12 | 1M | 2024-05-13 | 2.3 | 4.8% | 82.1% | 1.7% |
| 2024-06-12 | 2M | 2024-05-13 | 2.3 | 4.8% | 74.7% | 1.3% |
| 2024-07-31 | 1M | 2024-07-01 | 2.8 | 7.0% | 49.1% | 1.9% |
| 2024-07-31 | 2M | 2024-07-01 | 2.8 | 7.0% | 79.5% | 4.1% |
| 2024-09-18 | 1M | 2024-08-19 | 58.7 | 15.4% | 77.5% | 0.6% |
| 2024-09-18 | 2M | 2024-08-19 | 58.7 | 15.4% | 79.5% | 4.2% |
| 2024-11-07 | 1M | 2024-10-08 | 189.5 | 11.4% | 78.8% | 5.3% |
| 2024-11-07 | 2M | 2024-10-08 | 189.5 | 11.4% | 70.3% | 2.0% |
| 2024-12-18 | 1M | 2024-11-18 | 58.8 | 65.5% | 78.3% | 8.9% |
| 2024-12-18 | 2M | 2024-11-18 | 58.8 | 65.5% | 71.9% | 3.3% |
| 2025-01-29 | 1M | 2024-12-30 | 190.9 | 8.6% | 4.7% | 0.6% |
| 2025-01-29 | 2M | 2024-12-30 | 190.9 | 8.6% | 15.4% | 0.8% |
| 2025-03-19 | 1M | 2025-02-18 | 78.0 | 2.2% | 69.7% | 1.3% |
| 2025-03-19 | 2M | 2025-02-18 | 78.0 | 2.2% | 79.9% | 0.8% |
| 2025-05-07 | 1M | 2025-04-07 | 88.4 | 37.4% | 100.0% | 68.5% |
| 2025-05-07 | 2M | 2025-04-07 | 88.4 | 37.4% | 95.9% | 15.8% |
| 2025-06-18 | 1M | 2025-05-19 | 106.9 | 9.6% | 60.1% | 3.7% |
| 2025-06-18 | 2M | 2025-05-19 | 106.9 | 9.6% | 77.0% | 5.4% |
| 2025-07-30 | 1M | 2025-06-30 | 136.6 | 16.8% | 8.8% | 0.8% |
| 2025-07-30 | 2M | 2025-06-30 | 136.6 | 16.8% | 94.2% | 17.6% |
| 2025-09-17 | 1M | 2025-08-18 | 220.6 | 74.5% | 93.3% | 27.4% |
| 2025-09-17 | 2M | 2025-08-18 | 220.6 | 74.5% | 92.3% | 30.7% |
| 2025-10-29 | 1M | 2025-09-29 | 252.5 | 83.9% | 55.6% | 0.5% |
| 2025-10-29 | 2M | 2025-09-29 | 252.5 | 83.9% | 70.9% | 2.6% |
| 2025-12-10 | 1M | 2025-11-10 | 393.9 | 74.7% | 81.4% | 2.7% |
| 2025-12-10 | 2M | 2025-11-10 | 393.9 | 74.7% | 86.9% | 13.3% |
| 2026-01-28 | 1M | 2025-12-29 | 659.5 | 12.6% | 95.4% | 16.2% |
| 2026-01-28 | 2M | 2025-12-29 | 659.5 | 12.6% | 82.6% | 5.7% |

## 5.2 September 2024 FOMC

Table 5.2 reports the estimated parameters for four snapshot dates: August 19, August 27, September 5, and September 16. First, conditional means consistently show a positive separation, which is consistent with the inverse price-yield relationship. We note that the model does not impose $E_1 > E_2$. Second, conditional $\sigma$ is persistently higher in the case of a 50+ basis point cut, suggesting that there was more uncertainty related to this outcome. An explanation could be that the market saw a 50+ basis point decrease as unusual, and was unsure of its impact on the rate path. Table A.3 reports the corresponding pricing errors. On the 1M surface, the daily maximum relative pricing errors are 0.62%, 1.23%, 3.20%, and 5.47% for the four selected dates. On the 2M surface, they are 4.16%, 5.13%, 4.24%, and 2.09%. During the full 30-day window, Table A.1 shows that the average mixture maximum relative pricing error is 3.05% on the 1M surface and 5.05% on the 2M surface, compared with 68.3% and 80.8% for the single-lognormal benchmark.

**Table 5.2.** Parameter estimates - September 2024 FOMC meeting.

Estimated parameters of the mixture model for the September 2024 FOMC meeting across four snapshot dates. The columns 1M and 2M denote one-month and two-month constant-maturity option surfaces, respectively. $\omega$ is the one-sided Polymarket-implied risk neutral probability. $F$ is the observed futures settlement price. Subscript 1 corresponds to a 50+ basis point cut and subscript 2 corresponds to a 25 basis point cut. $E_1$ and $E_2$ are the conditional expected futures prices under each outcome, and $\sigma_1$ and $\sigma_2$ are the corresponding annualised lognormal volatilities.

| Date | Surface | $F$ | $\omega$ | $E_1$ | $E_2$ | $E_1 - E_2$ | $\sigma_1$ | $\sigma_2$ |
|---|---|---|---|---|---|---|---|---|
| 2024-08-19 | 1M | 103.62 | 15.4% | 104.37 | 103.48 | 0.89 | 3.49% | 2.13% |
| 2024-08-27 | 1M | 103.92 | 23.2% | 104.45 | 103.76 | 0.69 | 3.16% | 2.07% |
| 2024-09-05 | 1M | 104.11 | 37.2% | 104.49 | 103.89 | 0.60 | 3.41% | 1.93% |
| 2024-09-16 | 1M | 104.45 | 43.6% | 104.75 | 104.22 | 0.53 | 3.08% | 1.72% |
| 2024-08-19 | 2M | 103.64 | 15.4% | 104.40 | 103.50 | 0.90 | 3.81% | 2.00% |
| 2024-08-27 | 2M | 103.92 | 23.2% | 104.49 | 103.75 | 0.74 | 3.40% | 1.87% |
| 2024-09-05 | 2M | 104.11 | 37.2% | 104.70 | 103.76 | 0.94 | 3.19% | 1.74% |
| 2024-09-16 | 2M | 104.46 | 43.6% | 104.83 | 104.18 | 0.65 | 3.44% | 1.72% |

The spread between the two conditional means narrows over the window. On the 1M surface it goes from 0.886 price points on August 19 to 0.527 on September 16. The 2M surface is less clean, moving from 0.897 to 0.943 before falling to 0.653, though the sign does not change. In yield terms, these correspond to roughly 45 and 27 basis points respectively, using a modified duration of about 1.9 years.[1] Those numbers are larger than a simple 25 basis point overnight-rate difference would suggest. The two scenarios differ not only in the announced rate but in what each implies about the subsequent path, and it is this path divergence that the 2-year yield is likely reflecting (Gürkaynak et al., 2005; Kuttner, 2001). A 50+ basis point cut could signal the start of a faster easing cycle, while a 25 basis point cut could point more to a gradual adjustment.

[1] Specifically, $|\Delta y| \approx |\Delta P|/(P \times D_{\text{mod}})$. This is an approximation. The futures contract is not identical to the note and the correct conversion factor is not applied here.

The narrowing of the spread is also informative. Between August 19 and September 16, the observed futures price rises by about 0.8 price points. At the same time, $E_2$ moves up more than $E_1$. One possible reading is that the 25 basis point scenario itself became more dovish over the window. By mid-September, even the smaller cut was no longer priced as a very cautious first move.

The size of the effect is smaller than in earlier applications of the same type of model. Hanke et al. (2018) study exchange rates around major political events, while Branger et al. (2024) study wheat futures around the Ukraine war. The September FOMC meeting is a less extreme event for the underlying asset. The maximum state separation in Table 5.2 is below 1% of the futures price.[2]

**Figure 5.1.** Conditional volatilities for the September 18, 2024 FOMC meeting.

The figure reports the daily conditional volatilities estimated over the 30 days preceding the September 2024 FOMC meeting. The solid line corresponds to the 50+ basis point decrease state, and the dashed line corresponds to the 25 basis point decrease state. $\sigma_1$ and $\sigma_2$ denote the conditional volatilities of the two mixture components. The 1M and 2M panels use the one-month and two-month constant-maturity option surfaces, respectively.

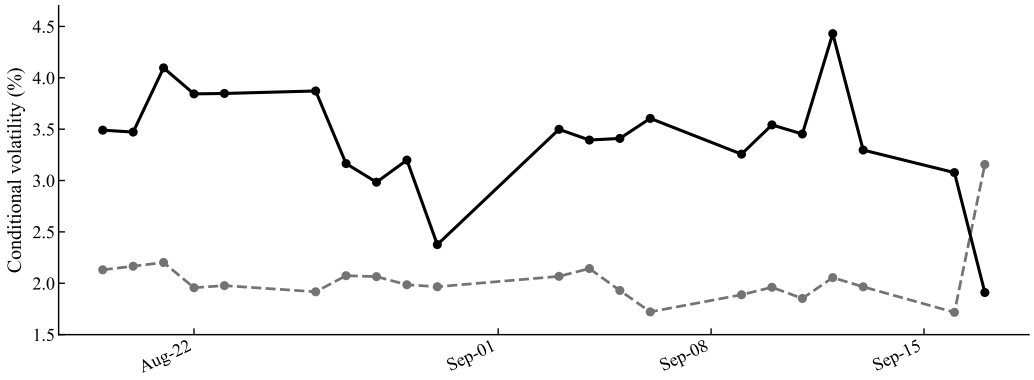


**(a)** 1M surface.

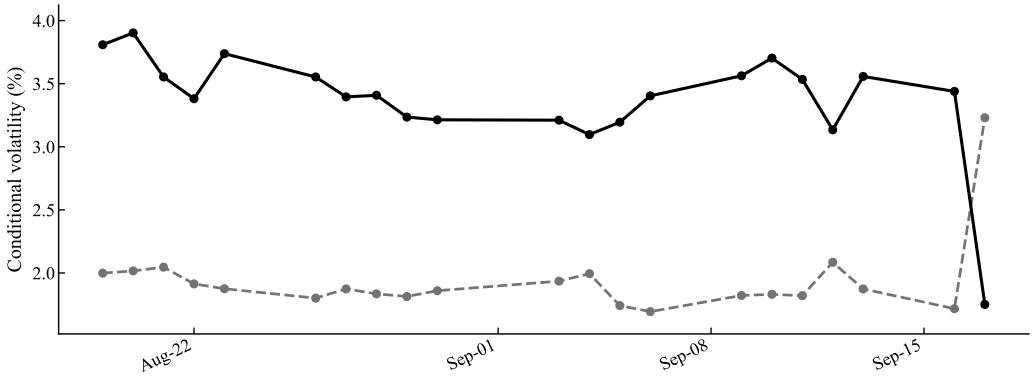


**(b)** 2M surface.

[2]The largest separation is 0.943 price points, on the 2M surface on September 5. At a futures price of 104.11, this represents $0.943/104.11 \approx 0.9\%$ of the price level.

The second pattern is in the conditional volatilities. $\sigma_1$ is higher than $\sigma_2$ on every snapshot date and on both surfaces. On September 16, for example, the 1M estimates are 3.08% and 1.72%. This suggests that the 50+ basis point state was not just the lower-yield state but also the more uncertain state. The relatively higher pre-meeting conditional volatility for a 50+ basis point cut can be interpreted to signal that the market was more uncertain about the interpretation of a 50 basis point cut before it happened. A similar pattern appears in Branger et al. (2024), where the adverse state carries higher conditional volatility. Figure 5.1 shows that this volatility gap is not limited to one date. The $\sigma_1 > \sigma_2$ gap is persistent on both surfaces and is typically between 1 and 2 percentage points, narrowing slightly in early September before compressing in the final observations.

Table 5.3 compares the moments of the mixture-implied density with the non-parametric benchmark. The two methods are close on the first two moments. For example, on September 16 for the 1M surface, the mixture mean is 104.450 and the non-parametric benchmark mean is 104.448, while the standard deviations are 0.764 and 0.760. It suggests that the mixture is not shifting the centre or the width of the unconditional density to fit the two components.

**Table 5.3.** Density moments - September 2024 FOMC

Moments of the risk-neutral density implied by the mixture model and the fitted non-parametric benchmark. Mean and standard deviation are in futures price points. Kurtosis is reported in excess of 3 (normal = 0). Each row corresponds to a snapshot date and constant-maturity option surface.

| Date | Surface | Mean mix | Std mix | Skew mix | Kurt mix | Mean poly | Std bench | Skew bench | Kurt bench | $K_{\min}$ | $K_{\max}$ |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 2024-08-19 | 1M | 103.620 | 0.779 | 0.650 | 1.307 | 103.619 | 0.786 | 0.774 | 1.927 | 100.98 | 107.86 |
| 2024-08-27 | 1M | 103.920 | 0.763 | 0.509 | 0.775 | 103.921 | 0.773 | 0.609 | 1.170 | 101.37 | 107.61 |
| 2024-09-05 | 1M | 104.110 | 0.822 | 0.578 | 0.878 | 104.109 | 0.824 | 0.647 | 1.132 | 101.17 | 107.90 |
| 2024-09-16 | 1M | 104.450 | 0.764 | 0.537 | 0.734 | 104.448 | 0.760 | 0.564 | 0.837 | 101.75 | 107.83 |
| 2024-08-19 | 2M | 103.640 | 1.047 | 0.662 | 1.813 | 103.629 | 1.018 | 0.497 | 1.189 | 99.21 | 109.83 |
| 2024-08-27 | 2M | 103.920 | 1.024 | 0.594 | 1.358 | 103.911 | 1.010 | 0.547 | 1.311 | 99.85 | 109.32 |
| 2024-09-05 | 2M | 104.110 | 1.105 | 0.693 | 0.967 | 104.112 | 1.120 | 0.825 | 1.519 | 100.32 | 109.25 |
| 2024-09-16 | 2M | 104.460 | 1.150 | 0.546 | 1.063 | 104.464 | 1.144 | 0.652 | 1.158 | 100.12 | 109.74 |

The agreement is weaker for skewness and excess kurtosis. Both methods agree on the sign of skewness across all selected dates, but the non-parametric benchmark tends to report larger right skewness, especially on the 1M surface. On September 16 for the 1M surface, the mixture skewness is 0.537, compared with 0.564 for the non-parametric benchmark. Similar differences appear in the kurtosis estimates. Since higher moments are especially sensitive to the tails, these differences should not be interpreted as evidence that either method recovers the true density (Figlewski, 2009; Shimko, 1993).

Table A.2 reports the moments and parameters of the non-parametric benchmark after attaching the tails. The first two moments remain close to the common-grid values, while the higher moments move more noticeably. The clearest example is the August 19 1M surface, where the full-tail skewness is 1.655 and excess kurtosis is 12.281. We therefore use the moment comparison mainly as a robustness check on the mean and standard deviation.

Positive skewness by itself is not enough to justify the mixture, since a single lognormal is also right-skewed. The useful point is that, in the mixture, part of the right tail can be linked to the 50+ basis point state. A single lognormal can fit a skewed shape, but it cannot identify which event state is associated with that part of the distribution. The analysis therefore puts emphasis on the conditional means and volatilities, in line with Hanke et al. (2018) and Branger et al. (2024).

Figure 5.2 shows the daily estimates of $E_1$ and $E_2$ over the 30-day window. The two conditional means move upward with the observed futures price, but they remain separated for most of the window. The 50+ basis point state stays above the 25 basis point state for most of the window, while the gap narrows as the meeting gets closer and the two conditional means converge in the final pre-meeting observation. This is consistent with what Table 5.2 already shows at the snapshot dates, but the daily series makes it harder to attribute the separation to a one-off estimation result.

**Figure 5.2.** Conditional means for the September 18, 2024 FOMC meeting.

The figure reports the daily conditional mean futures prices estimated over the 30 days preceding the September 2024 FOMC meeting. The solid line corresponds to the 50+ basis point decrease state, and the dashed line corresponds to the 25 basis point decrease state. $E_1$ and $E_2$ denote the conditional mean prices under the two policy outcomes. The 1M and 2M panels use the one-month and two-month constant-maturity option surfaces, respectively.

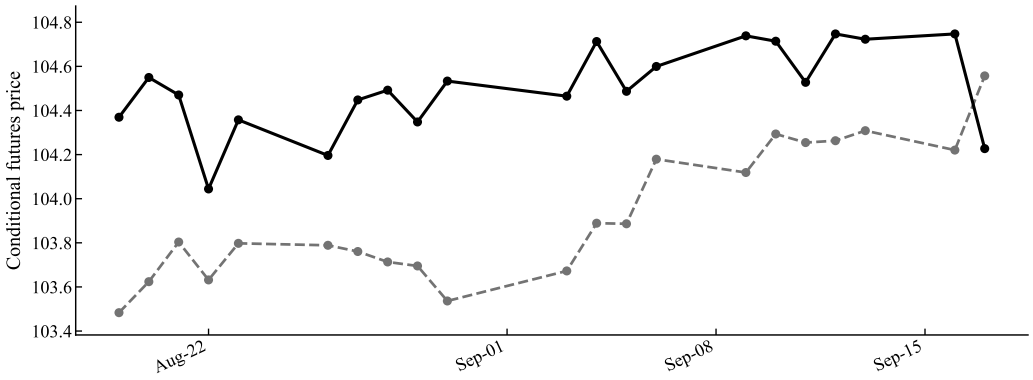


**(a)** 1M surface.

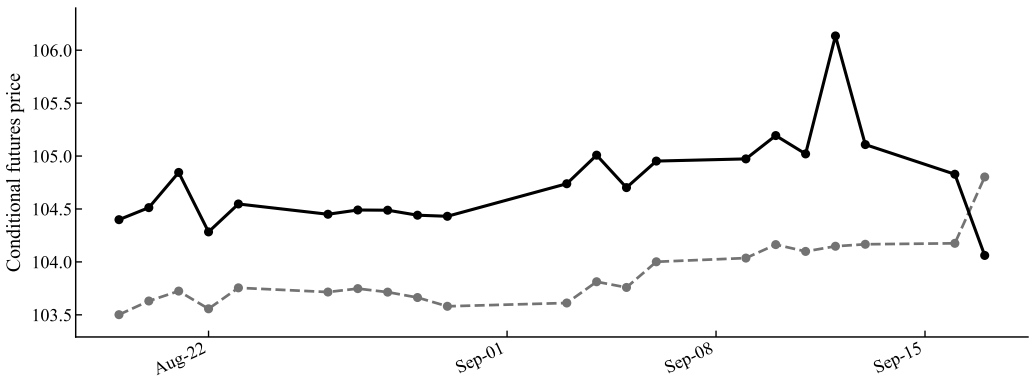


**(b)** 2M surface.

Figure 5.3 shows the risk-neutral densities and the decomposition of the mixture model into two conditional components. The distribution is unimodal, consistent with Branger et al. (2024), who report a similar shape, albeit over a significantly wider price range. This reflects the lower volatility of Treasury futures relative to commodity derivatives during war episodes.

**Figure 5.3.** Risk-neutral densities for the September 18, 2024 FOMC meeting, 1M surface.

Each panel shows the fitted mixture density (solid black), its two weighted components (dotted and dash-dotted), and the polynomial benchmark (dashed). $\omega$ denotes the Polymarket-implied probability of a 50+ basis point rate decrease.

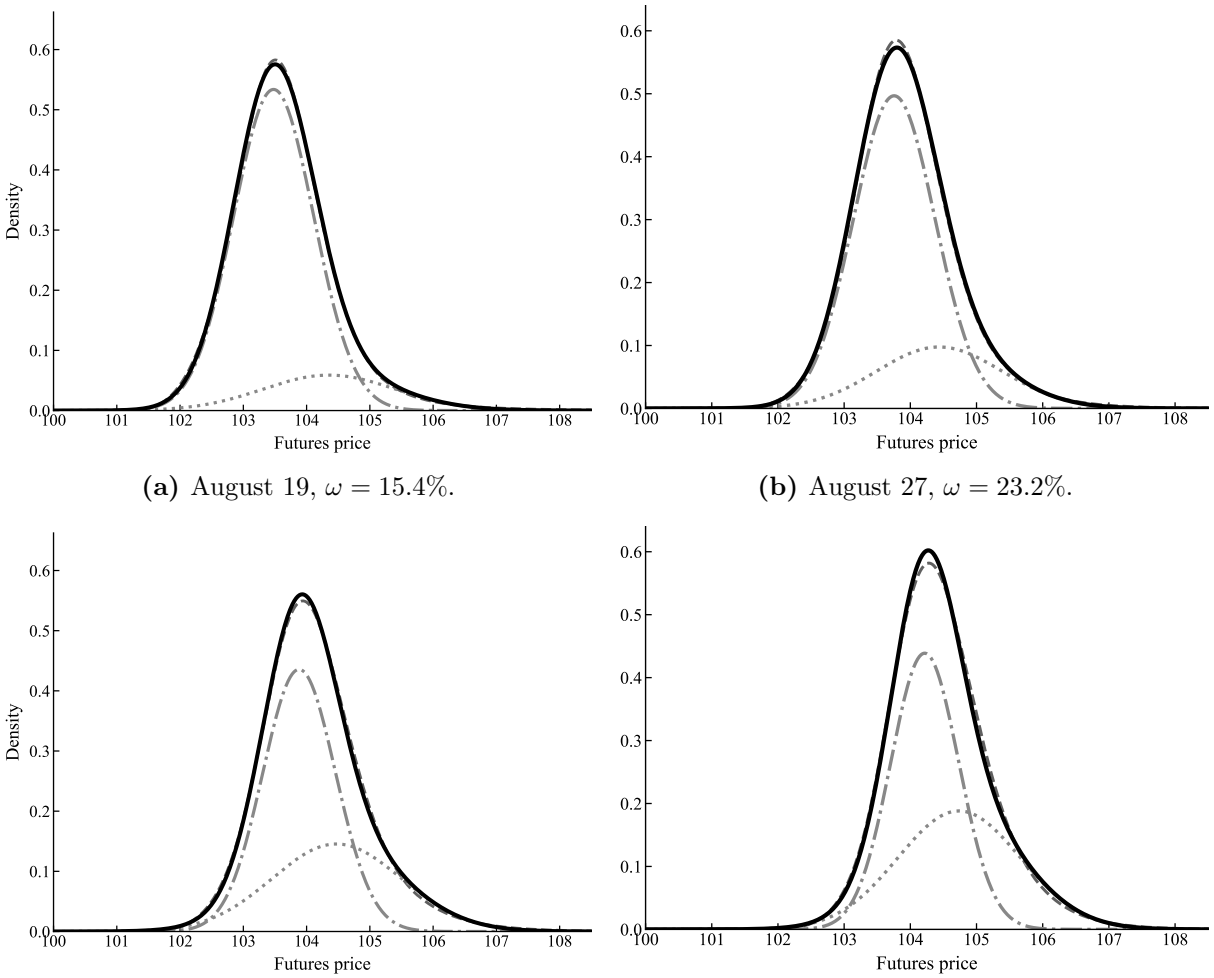


**(a)** August 19, $\omega = 15.4\%$.

**(b)** August 27, $\omega = 23.2\%$.

**(c)** September 5, $\omega = 37.2\%$.

**(d)** September 16, $\omega = 43.6\%$.

Hanke et al. (2018) argue that a useful decomposition should reveal information not available from either options or betting quotes alone. Following this approach, we compare the conditional means with the prediction-market probabilities. Table 5.4 reports the level and change correlations for both surfaces. In levels, the 1M correlations are $\text{corr}(\omega, E_1) = -0.12$ and $\text{corr}(\omega, E_2) = 0.52$. For the 2M surface, the corresponding correlations are $-0.39$ and 0.64. This suggests that the prediction market and the option surface are not giving the same signal, i.e., we are finding information not available from either options or betting quotes alone. In line with the findings of Hanke et al. (2018), we interpret this as the market was able to separate between the probability of the event, and the impact on prices.

**Table 5.4.** Correlations between probabilities and conditional means.

Correlations between the Polymarket-implied event probability and the estimated conditional futures-price components. Level correlations use $\omega$, $E_1$, $E_2$. Change correlations use relative one-period changes.

| Surface | $\text{corr}(\omega,\ E_1)$ | $\text{corr}(\omega,\ E_2)$ | $\text{corr}(\Delta\omega/\omega,\ \Delta E_1/E_1)$ | $\text{corr}(\Delta\omega/\omega,\ \Delta E_2/E_2)$ |
|---|---|---|---|---|
| 1M | -0.121 | 0.516 | -0.293 | -0.071 |
| 2M | -0.393 | 0.635 | -0.788 | -0.074 |

## 5.3 Limitations

We extract risk-neutral densities, which may differ from real world probabilities. This is standard in the RND literature (e.g., Hanke et al. (2018), Branger et al. (2024)).

We limit the study to one- and two-month constant maturity surfaces. FOMC meetings are spaced roughly every six weeks, and prediction markets for upcoming decisions tend to be nearly resolved within that one month window. September 2024 was unusual in that genuine uncertainty persisted into the final days of the meeting. Longer dated option surfaces could potentially capture a richer period of disagreement, but they also introduce exposure to multiple events and make the binary conditioning assumption harder to justify. An idea for future research would be to condition on the expected rate path implied by betting quotes for multiple outcomes rather than a single meeting.

We condition on a binary rate-decision structure, focusing on two dominant outcomes rather than the full set of Polymarket contracts. If the relevant pre-meeting uncertainty is not binary, or not primarily about the rate decision, then the Polymarket quotes we use may not correspond to the event actually shaping the volatility surface. The mixture model may still produce a fit, but the economic interpretation of the conditional components becomes unclear.

The two-year Treasury is a low-volatility instrument where FOMC risk does not dominate the option surface the way a referendum or a war dominates FX and commodity options (see e.g., Hanke et al. (2018), Branger et al. (2024)). Because we use constant maturity options and futures to capture the FOMC meetings, we are naturally inviting post-event noise since the option surface also prices other events such as employment data and CPI releases. While this issue is also relevant for Hanke et al. (2018), our study may be more sensitive to this issue.

More broadly, the RND literature tends to focus on describing how density shapes and moments change around events without testing whether the observed changes are statistically significant (Jackwerth, 1999). While we check pricing errors, we mainly rely on descriptive statistics.

# 6 Conclusion

We combine prediction-market probabilities with Treasury futures options to recover conditional risk-neutral distributions around FOMC meetings. For September 2024, the decomposition produces economically meaningful conditional components. The larger-cut state is more volatile across most of the pre-meeting window, suggesting uncertainty about the implications of a larger cut. The conditional mean separation exceeds what a single overnight rate change would imply, indicating that the two scenarios priced different expectations for the subsequent policy path.

The framework requires sustained two-sided uncertainty in prediction-market probabilities, which limits its scope to meetings where the policy decision is genuinely contested. When this condition holds, the conditional moments are useful ex ante: they describe how markets price different outcomes before the announcement, rather than forecasting the realised outcome. For central banks, risk managers, and traders, the shapes of the conditional densities are informative regardless of what is ultimately announced.

# References


Alexiou, L., Goyal, A., Kostakis, A., & Rompolis, L. (2025). Pricing event risk: Evidence from concave implied volatility curves. *European finance review*, *29*(4), 963–1007.

Bahra, B. (1997). *Implied risk-neutral probability density functions from option prices: Theory and application* (Working Paper). Bank of England.

Beber, A., & Brandt, M. W. (2006). The effect of macroeconomic news on beliefs and preferences: Evidence from the options market. *Journal of monetary economics*, *53*(8), 1997–2039.

Birru, J., & Figlewski, S. (2012). Anatomy of a meltdown: The risk neutral density for the S&P 500 in the fall of 2008. *Journal of Financial Markets*, *15*(2), 151–180. https://doi.org/10.1016/j.finmar.2011.09.001

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, *3*(1-2), 167–179.

Branger, N., Hanke, M., & Weissensteiner, A. (2024). The information content of wheat derivatives regarding the Ukrainian war. *Journal of Futures Markets*, *44*, 420–431. https://doi.org/10.1002/fut.22475

Breeden, D. T., & Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *The Journal of Business*, *51*(4), 621–651. https://doi.org/10.1086/296025

CME Group. (2026). 2-year T-note futures and options: Contract specifications [Accessed: 2026-04-23]. https://www.cmegroup.com/markets/interest-rates/us-treasury/2-year-us-treasury-note.contractSpecs.html

Coles, S. (2001). *An introduction to statistical modeling of extreme values*. Springer.

Cook, T., & Hahn, T. (1989). The effect of changes in the federal funds rate target on market interest rates in the 1970s. *Journal of monetary economics*, *24*(3), 331–351.

Diercks, A. M., Katz, J. D., & Wright, J. H. (2026). *Kalshi and the rise of macro markets* (Working Paper No. 34702). National Bureau of Economic Research. https://doi.org/10.3386/w34702

Figlewski, S. (2009). Estimating the implied risk-neutral density for the US market portfolio. In *Volatility and time series econometrics: Essays in honor of robert engle* (pp. 323–353). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199549498.003.0015

Goldfarb, S. (2024). Stocks slip after fed cuts rates. *The Wall Street Journal*. https://www.wsj.com/finance/stocks/stocks-slip-after-fed-cuts-rates-fd174e26

Gürkaynak, R., Sack, B., & Swanson, E. (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, *1*(1), 55–93.

Hanke, M., Poulsen, R., & Weissensteiner, A. (2018). Event-related exchange-rate forecasts combining information from betting quotes and option prices. *Journal of Financial and Quantitative Analysis*, *53*(6), 2663–2683.

Hosking, J. R. M., & Wallis, J. R. (1987). Parameter and quantile estimation for the generalized Pareto distribution. *Technometrics*, *29*(3), 339–349.

Hull, J. C. (2021). *Options, futures, and other derivatives* (11th Global Edition). Pearson.

Jackwerth, J. C. (1999). Option-implied risk-neutral distributions and implied binomial trees: A literature review. *Journal of Derivatives*, *7*(2), 66–82. https://doi.org/10.2139/ssrn.183705

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, *47*(3), 523–544. https://doi.org/10.1016/S0304-3932(01)00055-1

Leahy, M. P., & Thomas, C. P. (1996). *The sovereignty option: The Quebec referendum and market views on the canadian dollar* (International Finance Discussion Papers No. 555). Board of Governors of the Federal Reserve System.

Malz, A. M. (1997). Estimating the probability distribution of the future exchange rate from option prices. *Journal of Derivatives*, *5*(2), 18–36.

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

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The Computer Journal*, *7*(4), 308–313. https://doi.org/10.1093/comjnl/7.4.308

Shimko, D. (1993). Bounds of probability. *RISK*, *6*(4), 33–37.

Wolfers, J., & Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, *18*, 107–126.

# A Appendix

**Figure A.1.** Each panel shows the full available time series of outcome probabilities for a single FOMC meeting.

**(a)** June 12, 2024. $2.3 million total volume.

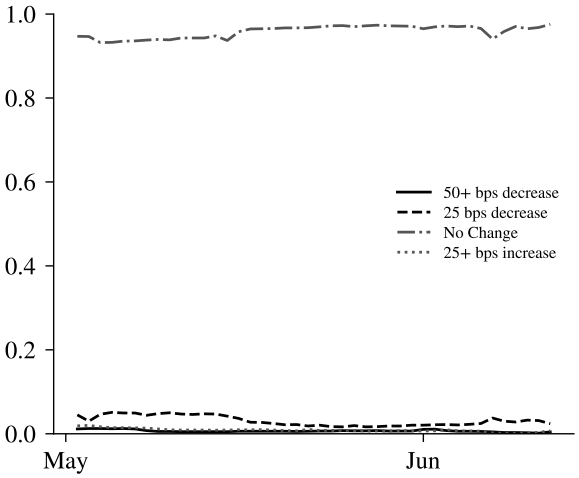


**(b)** July 31, 2024. $2.8 million total volume.

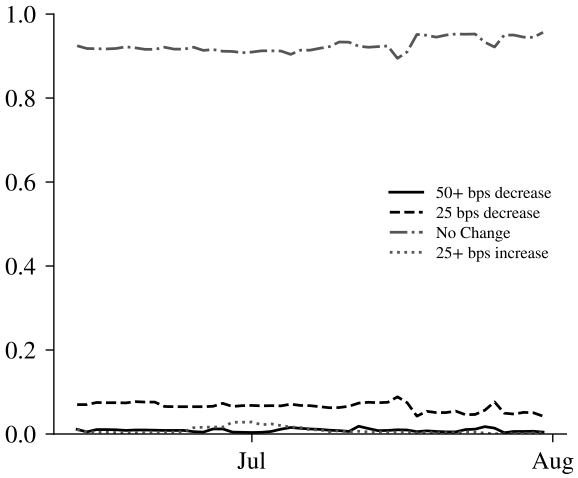


**(c)** September 18, 2024. $58.7 million total volume.

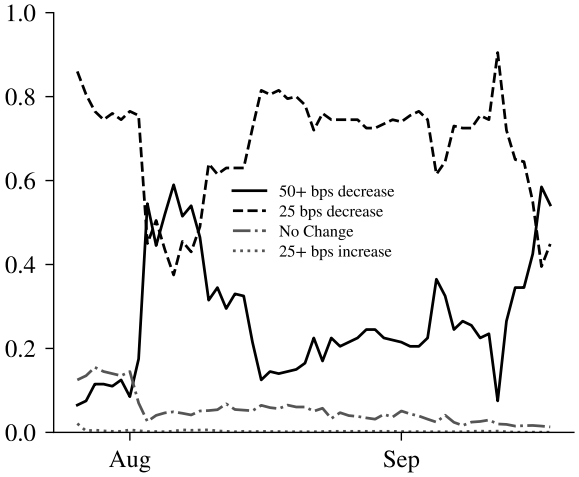


**(d)** November 7, 2024. $189.5 million total volume.

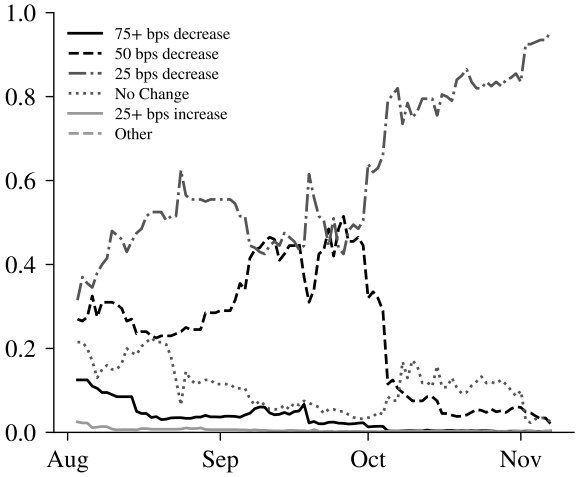


**(e)** December 18, 2024. $58.8 million total volume.

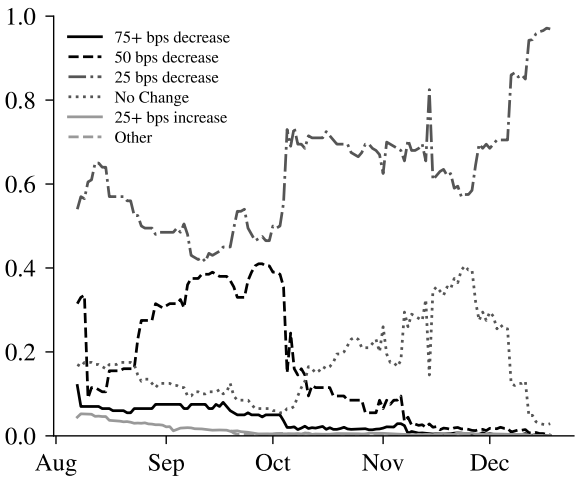


**(f)** January 29, 2025. $190.9 million total volume.

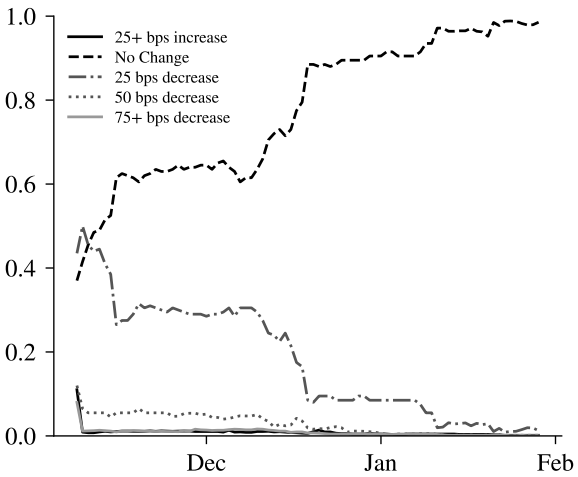

**Figure A.1.** Polymarket FOMC betting quotes, raw probabilities (continued).

**(a)** March 19, 2025. $78.0 million total volume.

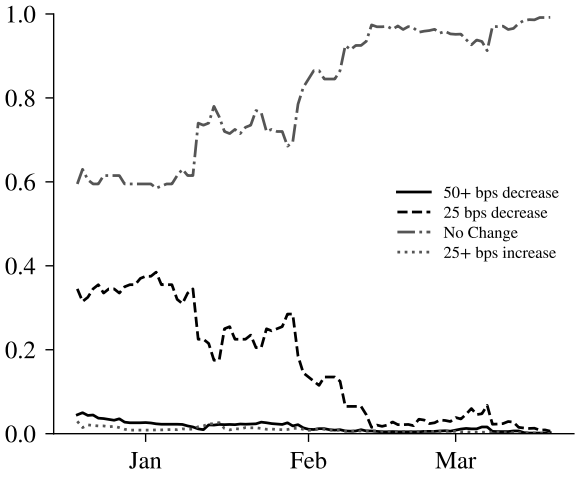


**(b)** May 7, 2025. $88.4 million total volume.

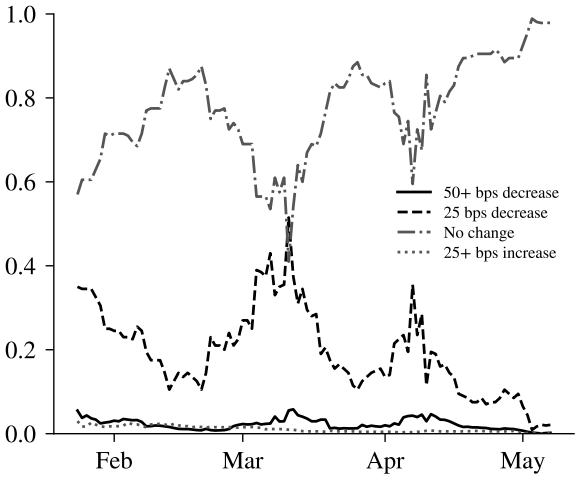


**(c)** June 18, 2025. $106.9 million total volume.

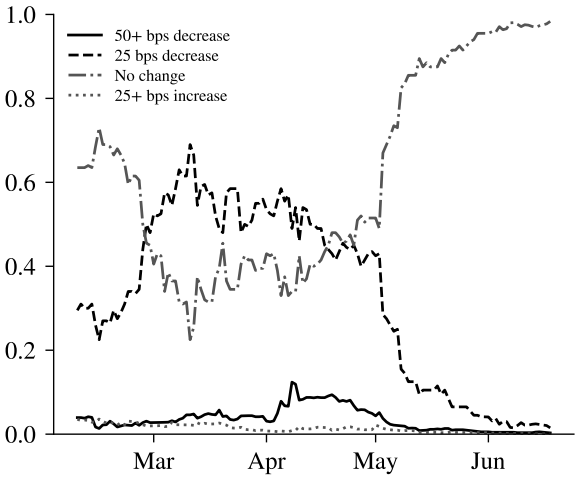


**(d)** July 30, 2025. $136.6 million total volume.

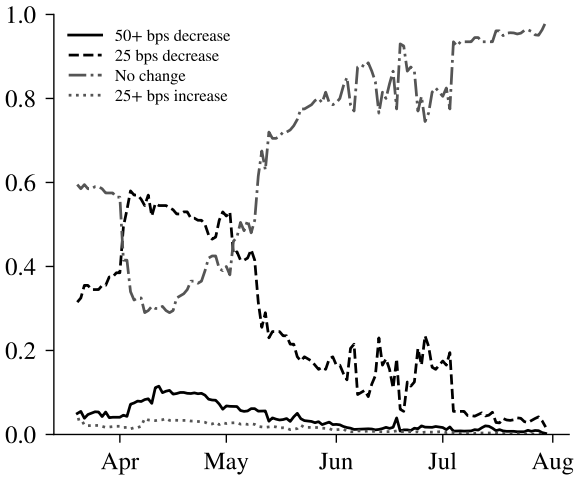


**(e)** September 17, 2025. $220.6 million total volume.

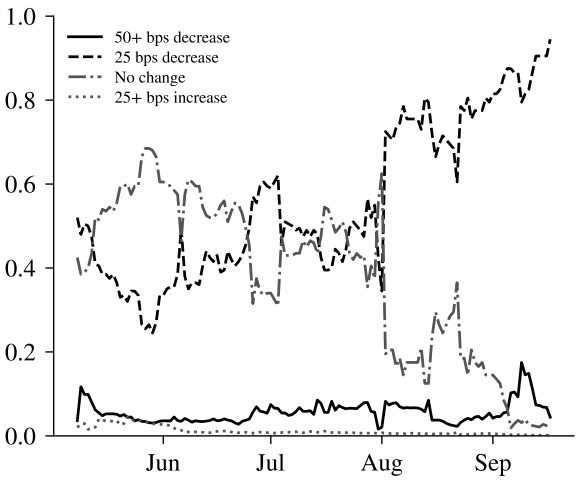


**(f)** October 29, 2025. $252.5 million total volume.

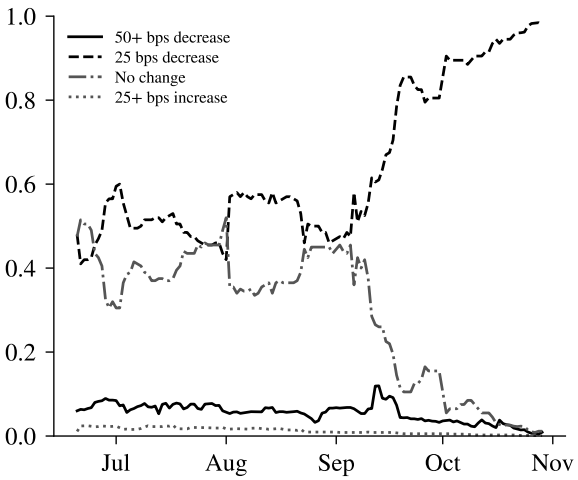

**Figure A.1.** Polymarket FOMC betting quotes, raw probabilities (continued).

**(a)** December 10, 2025. $393.9 million total volume.

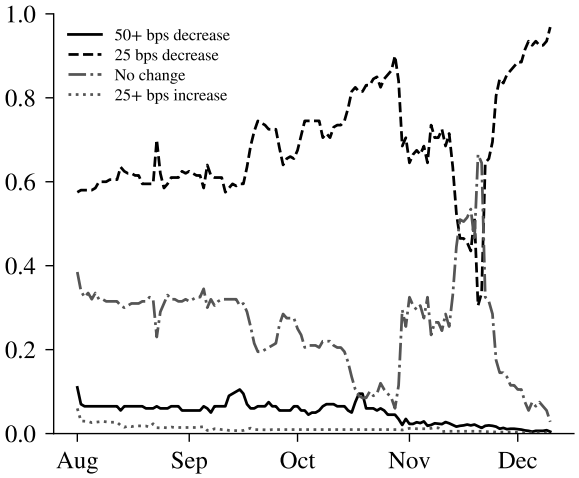


**(b)** January 28, 2026. $659.5 million total volume.

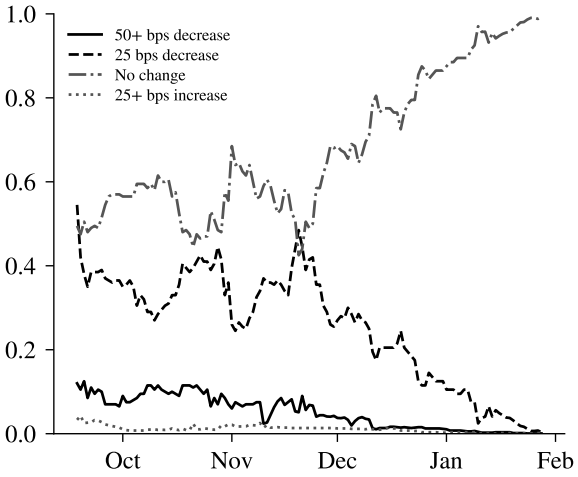

**Figure A.2.** Each panel shows the cleaned time series of probabilities for a single FOMC meeting.

**(a)** June 12, 2024.

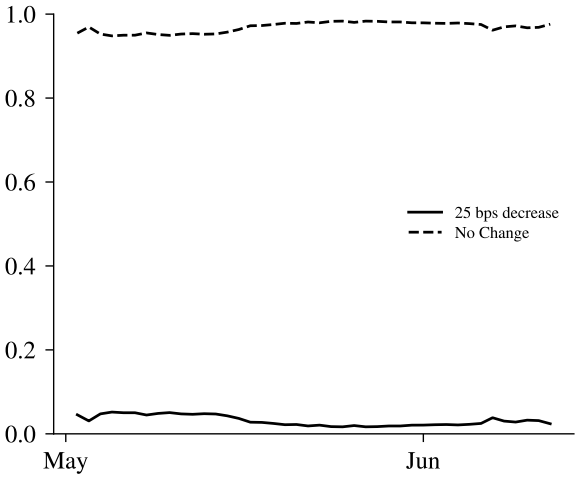


**(b)** July 31, 2024.

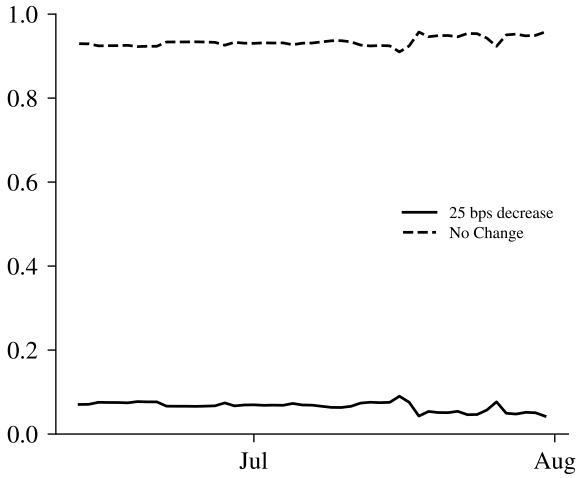


**(c)** September 18, 2024.

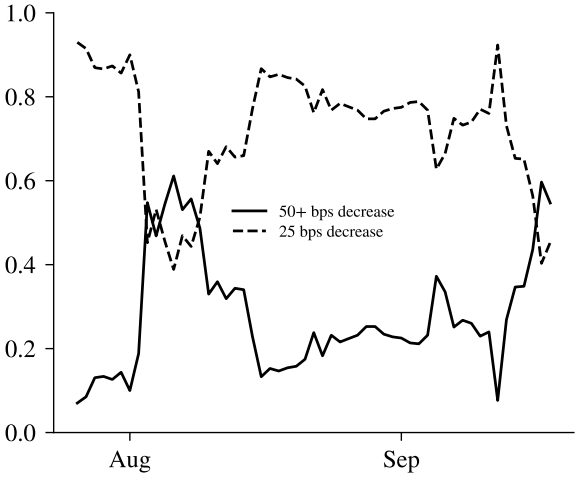


**(d)** November 7, 2024.

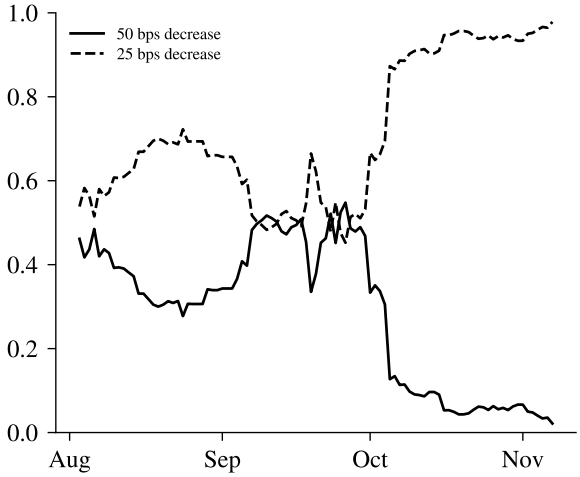


**(e)** December 18, 2024.

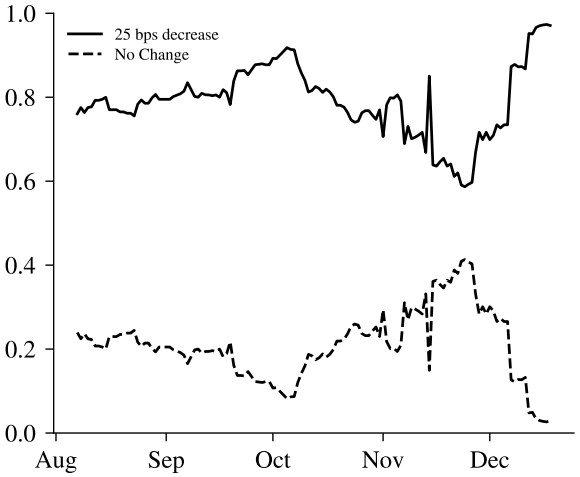


**(f)** January 29, 2025.

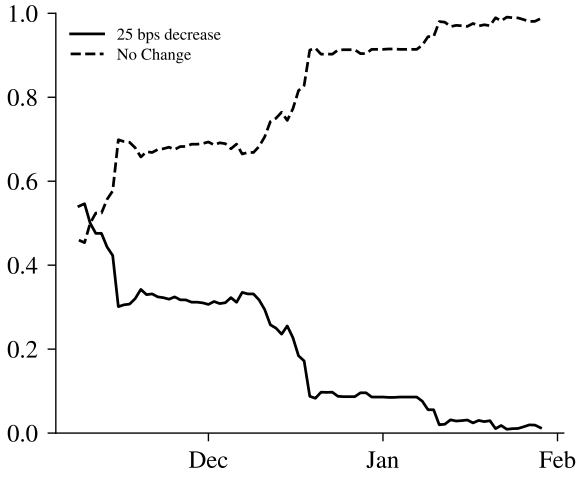

**Figure A.2.** Polymarket FOMC betting quotes, cleaned (continued).

**(a)** March 19, 2025.

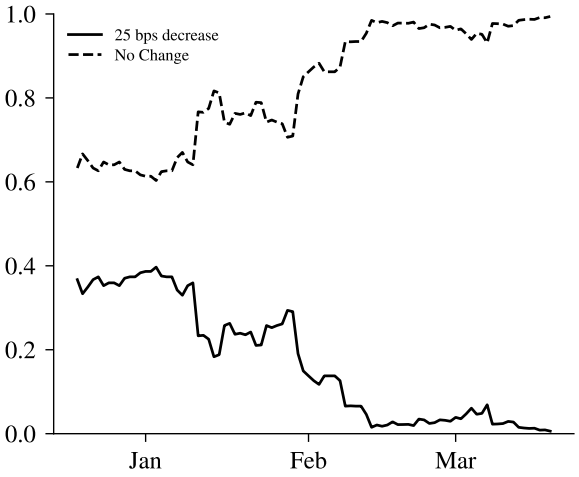


**(b)** May 7, 2025.

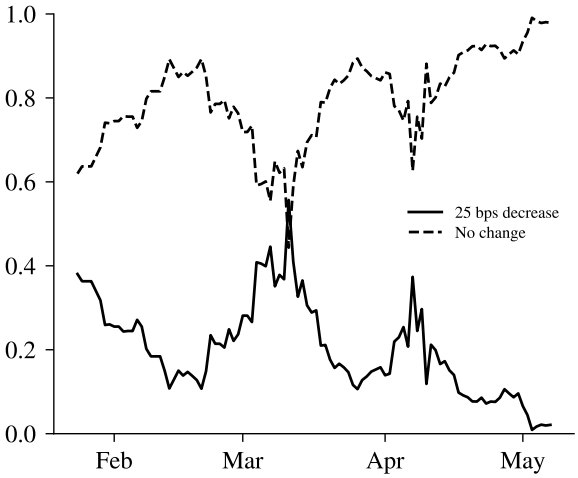


**(c)** June 18, 2025.

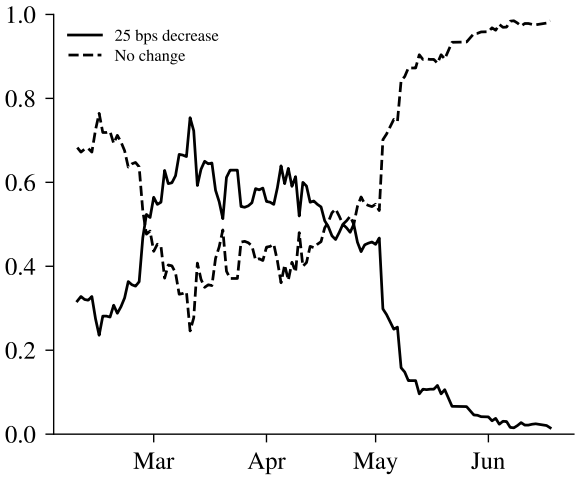


**(d)** July 30, 2025.

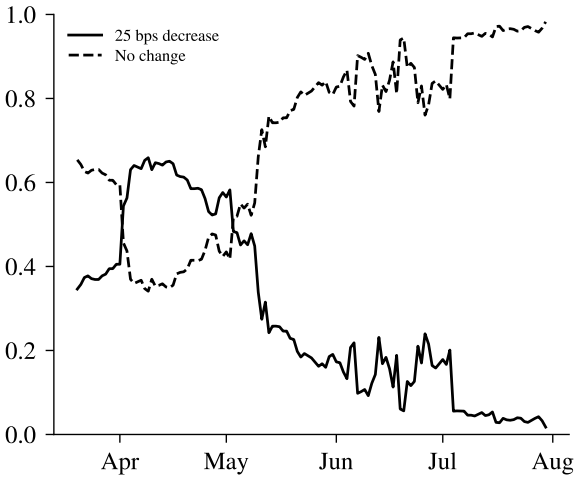


**(e)** September 17, 2025.

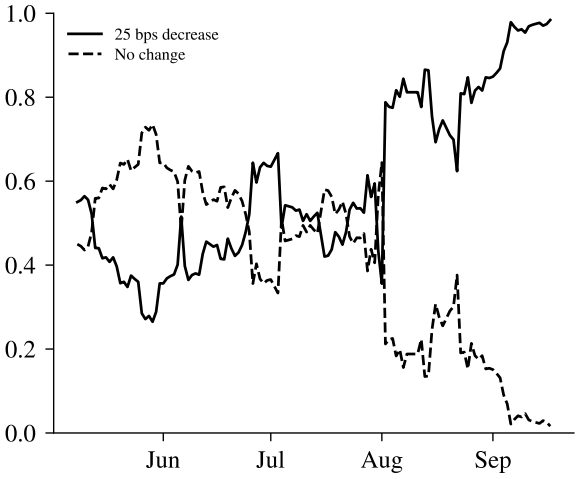


**(f)** October 29, 2025.

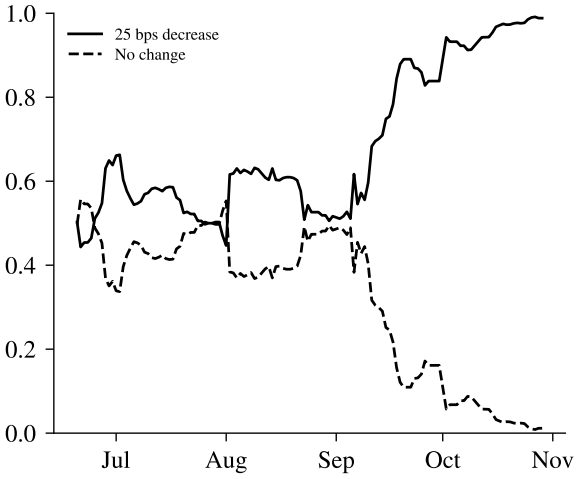

**Figure A.2.** Polymarket FOMC betting quotes, cleaned (continued).

**(a)** December 10, 2025. **(b)** January 28, 2026.

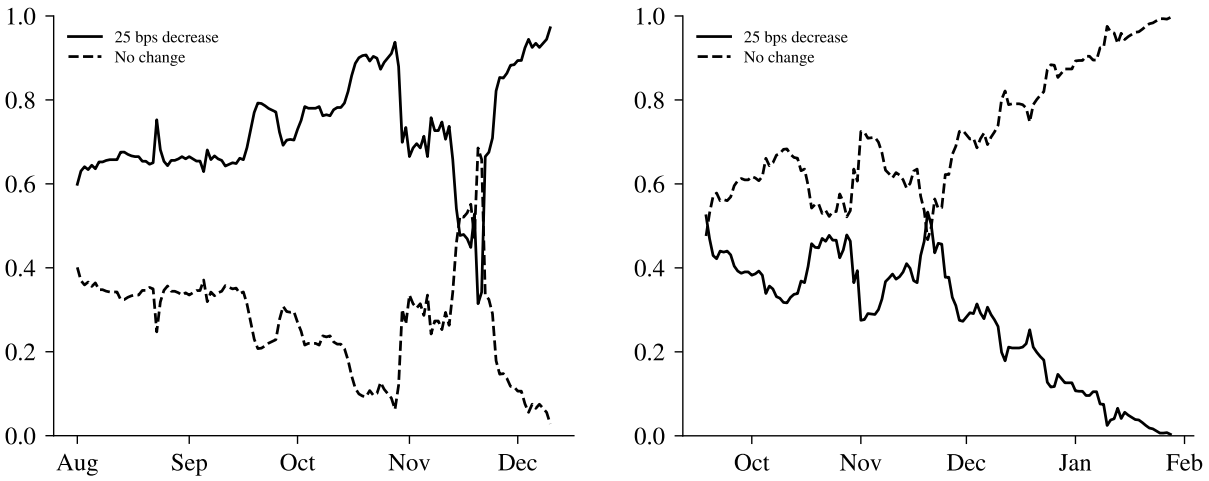


**Figure A.3.** Daily implied volatility smiles across the full sample. Each line represents the annualised Black (1976) implied volatility for each delta quote from Bloomberg. The sample covers June 2024 through March 2026.

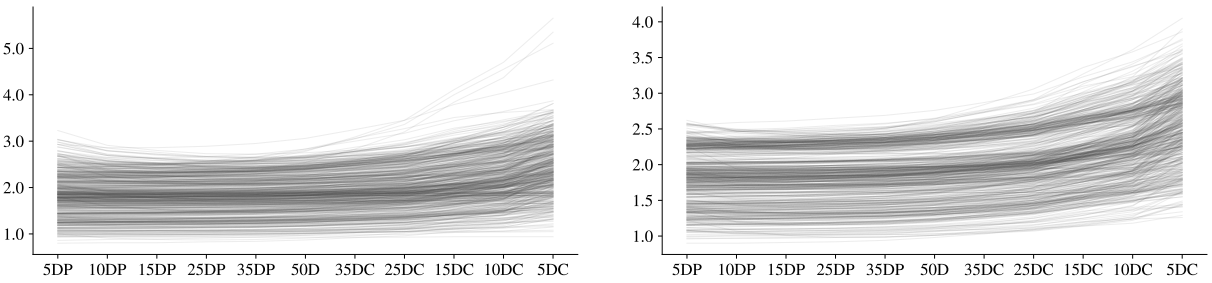


**(a)** 1-month constant maturity **(b)** 2-month constant maturity

**Table A.1.** Maximum relative pricing error summary over the September 2024 pre-meeting window.

Maximum relative pricing errors over all available dates in the 30-day September pre-meeting window, separately for the 1M and 2M option surfaces.

| Surface | Lognormal mean | Lognormal median | Lognormal max | Mixture mean | Mixture median | Mixture max |
|---|---|---|---|---|---|---|
| 1M | 68.3% | 71.3% | 81.1% | 3.05% | 2.65% | 5.98% |
| 2M | 80.8% | 80.6% | 90.7% | 5.05% | 5.13% | 11.13% |

**Table A.2.** Fitted non-parametric moments and parameters

Moments of the fitted non-parametric benchmark after GPD tail completion. Left, central, and right masses are reported as percentages of total probability mass. $\xi_L$ and $\xi_R$ denote the fitted GPD shape parameters in the left and right tails.

| Date | Surface | Mean | Std | Skew | Kurt | Mass | Left mass | Central mass | Right mass | $\xi_L$ | $\xi_R$ |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 2024-08-19 | 1M | 103.631 | 0.836 | 1.655 | 12.281 | 1.000027 | 5.48% | 90.94% | 3.59% | -0.243 | 0.218 |
| 2024-08-27 | 1M | 103.926 | 0.787 | 0.766 | 2.101 | 0.999996 | 5.21% | 90.60% | 4.19% | -0.162 | 0.026 |
| 2024-09-05 | 1M | 104.115 | 0.840 | 0.799 | 1.974 | 0.999995 | 5.10% | 90.99% | 3.92% | -0.171 | -0.011 |
| 2024-09-16 | 1M | 104.452 | 0.770 | 0.667 | 1.364 | 0.999995 | 4.98% | 91.23% | 3.79% | -0.153 | -0.024 |
| 2024-08-19 | 2M | 103.631 | 1.023 | 0.562 | 1.671 | 0.999997 | 5.32% | 92.24% | 2.44% | -0.225 | 0.077 |
| 2024-08-27 | 2M | 103.920 | 1.045 | 1.050 | 6.621 | 1.000022 | 5.36% | 91.79% | 2.84% | -0.220 | 0.221 |
| 2024-09-05 | 2M | 104.121 | 1.144 | 0.972 | 2.282 | 0.999992 | 5.64% | 90.18% | 4.18% | -0.241 | -0.116 |
| 2024-09-16 | 2M | 104.469 | 1.158 | 0.739 | 1.610 | 0.999995 | 3.98% | 92.11% | 3.91% | -0.087 | -0.062 |

**Table A.3.** Maximum relative pricing error - September 2024 estimates.

Maximum relative pricing errors for the selected September 2024 snapshot dates. The columns lognormal and mixture denote the maximum relative pricing error across the option quotes under the single-lognormal benchmark and the two-state mixture model, respectively.

| Date | Surface | $\omega$ | Lognormal | Mixture |
|---|---|---|---|---|
| 2024-08-19 | 1M | 15.4% | 77.5% | 0.62% |
| 2024-08-27 | 1M | 23.2% | 65.6% | 1.23% |
| 2024-09-05 | 1M | 37.2% | 73.8% | 3.20% |
| 2024-09-16 | 1M | 43.6% | 71.3% | 5.47% |
| 2024-08-19 | 2M | 15.4% | 79.5% | 4.16% |
| 2024-08-27 | 2M | 23.2% | 76.2% | 5.13% |
| 2024-09-05 | 2M | 37.2% | 81.6% | 4.24% |
| 2024-09-16 | 2M | 43.6% | 69.4% | 2.09% |

# AI Disclaimer

First, we used ChatGPT and Claude to assist with coding. Second, these AI tools were also used to improve language and clarity. AI mainly improved readability and increased efficiency during the coding process. The main risk of using AI is associated with hallucinations, hence all output was reviewed and cross-checked extensively. The main insight from using AI is that, when used with a sceptical approach, it is excellent at routine tasks that require no judgement, which frees up more time for analytical work.