How Do Inflation Expectations React to CPI Announcements?

Evidence from Minute-Level Kalshi Markets

Mario Diez Martínez

January 14, 2026

Abstract

U.S. CPI releases arrive at fixed times, but beliefs about inflation move minute by minute. This project measures how near-term inflation expectations update in response to CPI announcements using one-minute data from Kalshi CPI month-over-month (MoM) prediction markets. Kalshi lists CPI markets as collections of binary threshold contracts; from minute-level bid/ask quotes across strikes, the analysis recovers an implied probability distribution over CPI outcomes and computes belief moments (mean and variance). The empirical design is a scheduled-announcement event study with a strict pre-release baseline at $t = -5$ minutes and post-release horizons from 5 to 240 minutes. Event-level regressions relate belief revisions to standardized CPI surprises. Diagnostics address anticipation in the pre-release window, placebo pseudo-events, hypothesis tests on the speed and completion of adjustment, and small-sample inference via a block bootstrap by event.

1 Introduction

Inflation expectations enter monetary policy decisions and affect nominal pricing. The empirical difficulty is timing. Survey measures of inflation expectations arrive monthly or quarterly, so many signals can land between survey vintages. That aggregation makes it hard to attribute revisions to a specific release (Bauer and Swanson, 2023a). Market-based measures derived from TIPS move at higher frequency but mix expected inflation with liquidity and risk premia (Gürkaynak et al., 2010).

Kalshi provides an intraday market-based measure of near-term CPI beliefs. Kalshi lists event contracts on CPI releases that trade around announcement times. Work that uses Kalshi data directly motivates this approach. Swanson et al. (2025) study traded macroeconomic event contracts and use Kalshi to measure expectation responses to policy news. Bürgi et al. (2025) study Kalshi pricing and returns and document patterns consistent with favorite-longshot bias in Kalshi prices. This project applies the same data source to CPI announcements and estimates how CPI surprises shift the implied mean and variance of CPI beliefs, and how quickly those updates are completed.

2 Research Question, Relevance, and Hypotheses

2.1 Research question

How do market-implied expectations of CPI MoM inflation update in response to CPI announcements at high frequency?

2.2 Economic and policy relevance

Inflation expectations matter for monetary policy transmission and inflation dynamics. Measuring how CPI releases shift expectations provides direct evidence on belief updating under information frictions (Coibion and Gorodnichenko, 2015). Minute-frequency windows limit confounding relative to low-frequency revisions, following the event-study logic used for scheduled macro announcements in macro-finance (Balduzzi et al., 2001; Andersen et al., 2003). CPI announcement days also matter for hedgers and macro traders because risk concentrates around release times.

2.3 Hypotheses

Let π denote CPI MoM inflation. Let S_i denote the standardized CPI surprise for release i. For post-release horizons $h \in \{5, 10, 15, 30, 60, 90, 120, 240\}$ minutes:

- **H1 (mean updating):** Positive CPI surprises raise the implied mean after the release, so $\beta_h > 0$ in (1).
- **H2 (uncertainty response):** CPI surprises shift the implied variance, so $\gamma_h \neq 0$ in (2).
- **H3 (speed and completion):** The mean response stops changing after a finite horizon. Section 6.4 tests equality across adjacent horizons and tests whether $\beta_{120} = \beta_{240}$.

3 Literature

Inflation expectation formation departs from full-information benchmarks in systematic ways. Coibion and Gorodnichenko (2015) model information rigidity and document empirical patterns consistent with infrequent and incomplete updating. High-frequency event studies exploit scheduled announcements to isolate the surprise component of news; this logic underpins the identification here (Balduzzi et al., 2001; Andersen et al., 2003; Faust et al., 2007). Prediction market prices can deviate from literal probabilities because of frictions and preferences (Manski, 2006). Kalshi-specific evidence motivates measurement using event contracts and cautions against treating prices as probabilities without qualification (Swanson et al., 2025; Bürgi et al., 2025).

4 Data and Construction

4.1 Kalshi CPI MoM markets and event universe

The empirical analysis uses Kalshi CPI month-over-month markets observed at minute frequency around CPI announcement times. Kalshi lists CPI contracts as binary threshold claims across strikes, of the form "Will CPI MoM exceed τ_k?" for thresholds $\{\tau_k\}$. Around each CPI release, the dataset contains minute-level bid/ask quotes for each strike. CPI release dates and the corresponding surprise series are obtained from the CPI announcement calendar and consensus forecasts.

4.2 Cleaning, baseline definition, and sample selection

Raw Kalshi quotes can contain non-informative placeholders and thin-liquidity artifacts. The event-study design uses a strict pre-release baseline at $t = -5$ minutes relative to the CPI release. Placeholder quotes are defined as bid=0 and ask=100.

An event is included at horizon h if (i) it has at least one non-placeholder bracket quote at $t = -5$ and an implied expectation can be computed, and (ii) the post-release minute $t = +h$ is observed. Under this rule, the regression sample has $N = 41$ events at horizons $h \in \{90, 120, 240\}$ and $N = 40$ events at horizons $h \in \{5, 10, 15, 30, 60\}$ due to one event missing post-release quotes at those shorter horizons.

4.3 From threshold prices to moments of beliefs

Let π denote CPI MoM inflation and let $m_{k,t}$ be the mid-price (in cents) for strike τ_k at minute t:

$$m_{k,t} = \frac{\mathrm{bid}_{k,t} + \mathrm{ask}_{k,t}}{2}.$$

Interpreting cents as probabilities, define the implied exceedance probability:

$$p_{k,t} \approx \Pr(\pi > \tau_k \mid t) = \frac{m_{k,t}}{100}.$$

Adjacent strikes define implied bracket probabilities:

$$\Pr(\tau_{k-1} < \pi \le \tau_k \mid t) = p_{k-1,t} - p_{k,t}.$$

Using these brackets (and a standard tail-handling convention), compute the implied mean and variance of beliefs at each minute:

$$\mathbb{E}_t[\pi] = \sum_k c_k \Pr(\tau_{k-1} < \pi \le \tau_k \mid t), \qquad \mathrm{Var}_t[\pi] = \sum_k (c_k - \mathbb{E}_t[\pi])^2 \Pr(\tau_{k-1} < \pi \le \tau_k \mid t),$$

where c_k denotes the representative value (bin center) for bracket k. All outcomes used for estimation are changes relative to baseline $t = -5$.

4.4 Visual evidence: event-level heatmaps

Figure 1 shows two event-level heatmaps that map bracket mid-prices into a time-by-strike view with the extracted implied mean overlaid.

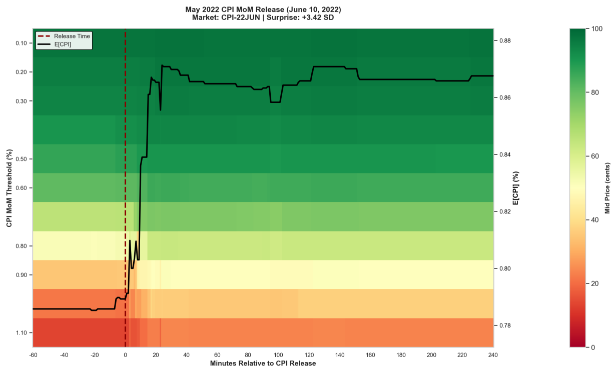

(a) May 2022 CPI (June 10, 2022): positive surprise example.

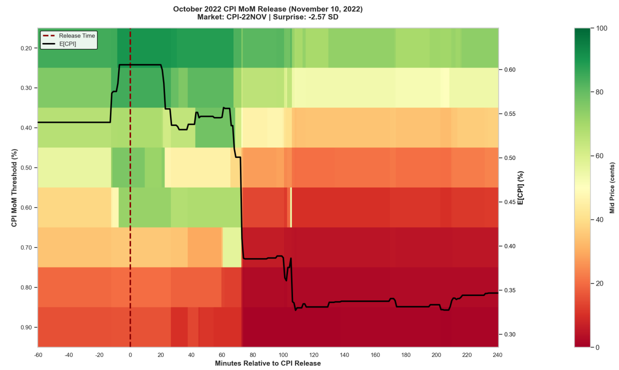

(b) October 2022 CPI (Nov 10, 2022): negative surprise example.

Figure 1: Heatmaps of bracket mid-prices (color) with extracted $\mathbb{E}[\pi]$ overlay (line). The vertical line marks the CPI release time.

5 Identification Strategy

5.1 Event-study logic

The identification strategy is a high-frequency event study around CPI announcements. CPI releases occur at predetermined times. In a narrow window around the announcement, the surprise component of the release constitutes an exogenous information shock affecting inflation expectations (Balduzzi et al., 2001; Andersen et al., 2003; Faust et al., 2007).

5.2 Estimating equations

Index CPI releases by i, define announcement time $t = 0$, and baseline time $t = -5$. For horizon $h \in \{5, 10, 15, 30, 60, 90, 120, 240\}$ minutes, define:

$$\Delta\mathbb{E}_i^{(h)} = \mathbb{E}_{i,+h}[\pi] - \mathbb{E}_{i,-5}[\pi], \qquad \Delta\mathrm{Var}_i^{(h)} = \mathrm{Var}_{i,+h}[\pi] - \mathrm{Var}_{i,-5}[\pi].$$

Let S_i denote the standardized CPI surprise. The main regressions are:

$$\Delta\mathbb{E}_i^{(h)} = \alpha_h + \beta_h S_i + \varepsilon_{i,h}, \tag{1}$$

$$\Delta\mathrm{Var}_i^{(h)} = \delta_h + \gamma_h S_i + u_{i,h}. \tag{2}$$

5.3 Econometric implementation

Equations (1)–(2) are estimated by ordinary least squares separately by horizon h on the event-level sample. Baseline inference uses heteroskedasticity-robust standard errors. Small-sample inference is assessed via a block bootstrap by event (Section 7.3).

6 Empirical Results

6.1 Effects on the implied mean of CPI beliefs

Table 1 reports estimates of β_h in (1). The mean response is economically small and statistically imprecise at very short horizons (5–15 minutes), but becomes statistically significant at the 30-minute horizon and remains significant through 240 minutes. The increase in R^2 from about 0.01 at 5 minutes to about 0.22 by 120 minutes indicates that CPI surprises explain a materially larger share of post-release belief revisions at medium horizons.

Table 1: Event-study estimates: $\Delta\mathbb{E}[\text{CPI}]$ on standardized CPI surprise

Horizon	$\hat{\beta}$	SE	t-stat	p-value	R^2	N
5 min	0.002246	0.003704	0.606493	0.547794	0.009587	40
10 min	0.005710	0.004487	1.272553	0.210908	0.040874	40
15 min	0.005851	0.004450	1.314785	0.196462	0.043512	40
30 min	0.012171	0.004668	2.607554	0.012961	0.151773	40
60 min	0.013334	0.004555	2.927173	0.005749	0.183995	40
90 min	0.016760	0.005402	3.102247	0.003563	0.197926	41
120 min	0.019717	0.005893	3.346034	0.001823	0.223045	41
240 min	0.019454	0.005983	3.251844	0.002368	0.213305	41

Notes: Event-level OLS by horizon. Dependent variable is the change in the Kalshi-implied mean CPI MoM belief from baseline $t = -5$ to $t = +h$ minutes. Regressor is standardized CPI surprise. Heteroskedasticity-robust SEs reported.

6.2 Effects on the implied variance of beliefs

Table 2 reports estimates of γ_h in (2). The variance response is not statistically distinguishable from zero at any horizon, and the point estimates are modestly negative throughout.

Table 2: Event-study estimates: ΔVar[CPI] on standardized CPI surprise

Horizon	$\hat{\gamma}$	SE	t-stat	p-value	R^2	N
5 min	-0.002560	0.002963	-0.863861	0.393085	0.019260	40
10 min	-0.002288	0.003245	-0.705093	0.485052	0.012914	40
15 min	-0.001278	0.002960	-0.431699	0.668400	0.004880	40
30 min	-0.001902	0.002825	-0.673369	0.504787	0.011792	40
60 min	-0.002921	0.002745	-1.063913	0.294082	0.028926	40
90 min	-0.002601	0.002872	-0.905712	0.370650	0.020600	41
120 min	-0.001763	0.002952	-0.597206	0.553822	0.009062	41
240 min	-0.002009	0.002945	-0.682168	0.499164	0.011791	41

Notes: Event-level OLS by horizon. Dependent variable is the change in the Kalshi-implied variance of CPI MoM beliefs from baseline $t = -5$ to $t = +h$ minutes. Regressor is standardized CPI surprise. Heteroskedasticity-robust SEs reported.

6.3 Cross-event visualization

Figure 2 plots $\Delta\mathbb{E}(h)$ against standardized CPI surprises across horizons with fitted OLS lines and confidence bands. Figure 3 provides the analogous plots for $\Delta\text{Var}(h)$.

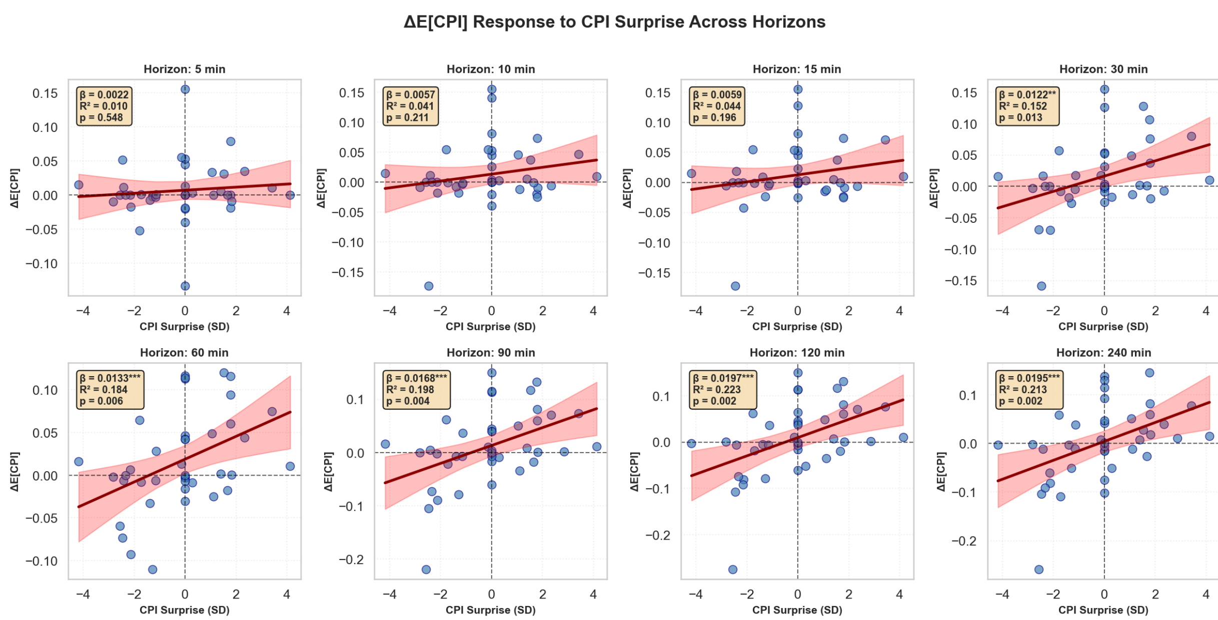

Figure 2: Scatter plots of $\Delta\mathbb{E}(h)$ on standardized CPI surprise across horizons with OLS fit and confidence bands.

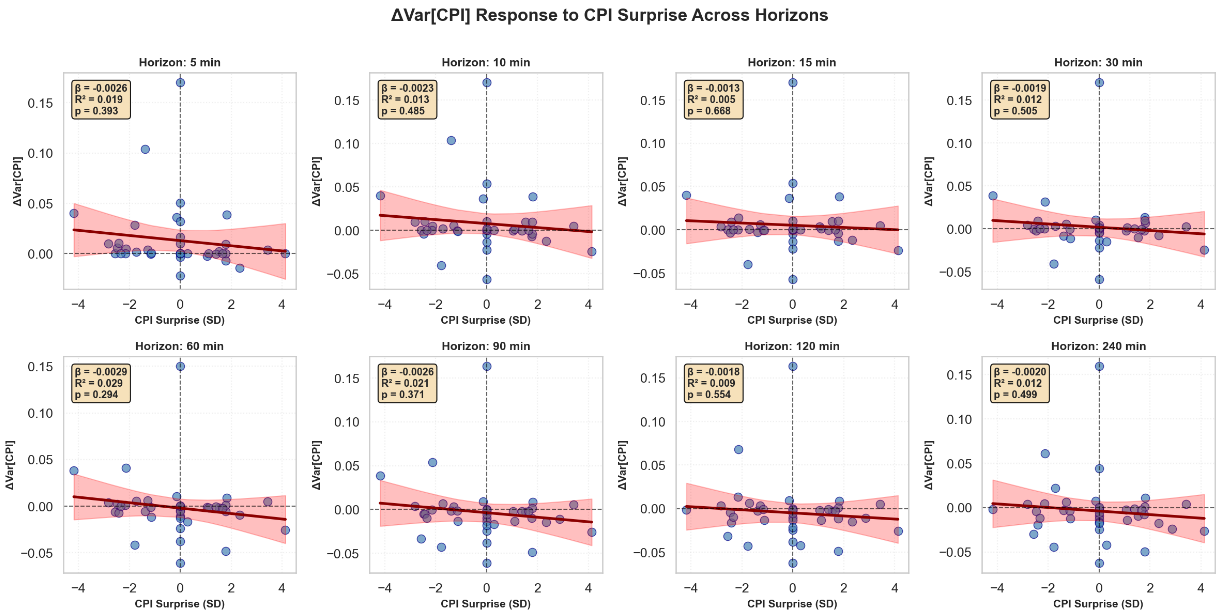

Figure 3: Scatter plots of $\Delta\text{Var}(h)$ on standardized CPI surprise across horizons with OLS fit and confidence bands.

6.4 Speed and completion of adjustment: hypothesis tests

The speed hypothesis is evaluated by tests that compare estimated slope coefficients across horizons.

Test 1 (sequential adjustment): For adjacent horizons (h_1, h_2), test

$$H_0 : \beta_{h_1} = \beta_{h_2} \qquad \text{vs.} \qquad H_1 : \beta_{h_1} \neq \beta_{h_2}.$$

For the mean response, the estimated coefficients rise monotonically from 5 minutes to 120 minutes. The equality tests do not reject at any adjacent pair. The p-values are:

$$\begin{aligned} &(5, 10) : 0.724182, && (10, 15) : 0.508929, \\ &(15, 30) : 0.836455, && (30, 60) : 0.570772, \\ &(60, 90) : 0.686050, && (90, 120) : 0.644290, \\ &(120, 240) : 0.487526. \end{aligned}$$

The late-window comparisons provide no evidence that adjustment continues beyond 90 minutes (90 vs. 120: $p = 0.644$; 120 vs. 240: $p = 0.488$).

For the variance response, the adjacent-horizon equality tests also do not reject. The p-values are:

$$\begin{aligned} &(5,10): 0.524684, && (10,15): 0.590944, \\ &(15,30): 0.439357, && (30,60): 0.397954, \\ &(60,90): 0.532094, && (90,120): 0.580609, \\ &(120,240): 0.476505. \end{aligned}$$

Test 2 (adjustment completion): Test whether the mean response is complete by 120 minutes using

$$H_0 : \beta_{120} \neq \beta_{240} \qquad \text{vs.} \qquad H_1 : \beta_{120} = \beta_{240}.$$

The equality test yields $p = 0.975053$, so the data do not reject equality. The interpretation is that the market fully incorporates the CPI surprise within 120 minutes. For the variance response, the analogous completion test yields $p = 0.953010$.

7 Robustness Checks and Diagnostics

7.1 Pre-trend (anticipation) in the pre-release window

To quantify pre-announcement dynamics, an event-specific linear drift model for the implied mean is estimated in a 240-minute pre-window,

$$\mathbb{E}_{i,t}[\pi] = \alpha_i + \phi_i t + \varepsilon_{i,t}, \qquad t \in [-240, -6],$$

and $H_0 : \phi_i = 0$ is tested. Out of 51 CPI events, 44 have sufficient data to estimate the model, and 37 of these (84.1%) exhibit statistically significant drift at the 5% level. This pattern is consistent with anticipation: beliefs adjust before the official release as traders update on partial information, related macro signals, and informed speculation.

7.2 Placebo pseudo-events (non-CPI timestamps)

To test whether the design is spuriously picking up systematic non-announcement movements, placebo tests use random non-CPI timestamps (excluding CPI windows) and compute 120-minute belief revisions. Table 3 shows that placebo mean changes are not statistically different from zero for either the implied mean or implied variance.

Table 3: Placebo tests (random non-CPI timestamps, 120-minute horizon)

Measure	Mean Δ	SE	t-stat	p-value	N (placebos)
$\Delta\mathbb{E}$[CPI]	-0.000263	0.000422	-0.622556	0.540977	20
ΔVar[CPI]	-0.000156	0.000101	-1.541067	0.139791	20

Notes: One-sample t-tests of $H_0 : \mathbb{E}[\Delta] = 0$ using pseudo-announcement timestamps drawn from non-CPI periods. Horizon fixed at 120 minutes.

7.3 Bootstrap inference

Given modest event counts, heteroskedasticity-robust standard errors are complemented with a block bootstrap by event (resampling events with replacement). Figures 4–5 summarize the bootstrap distributions of horizon-specific slope coefficients.

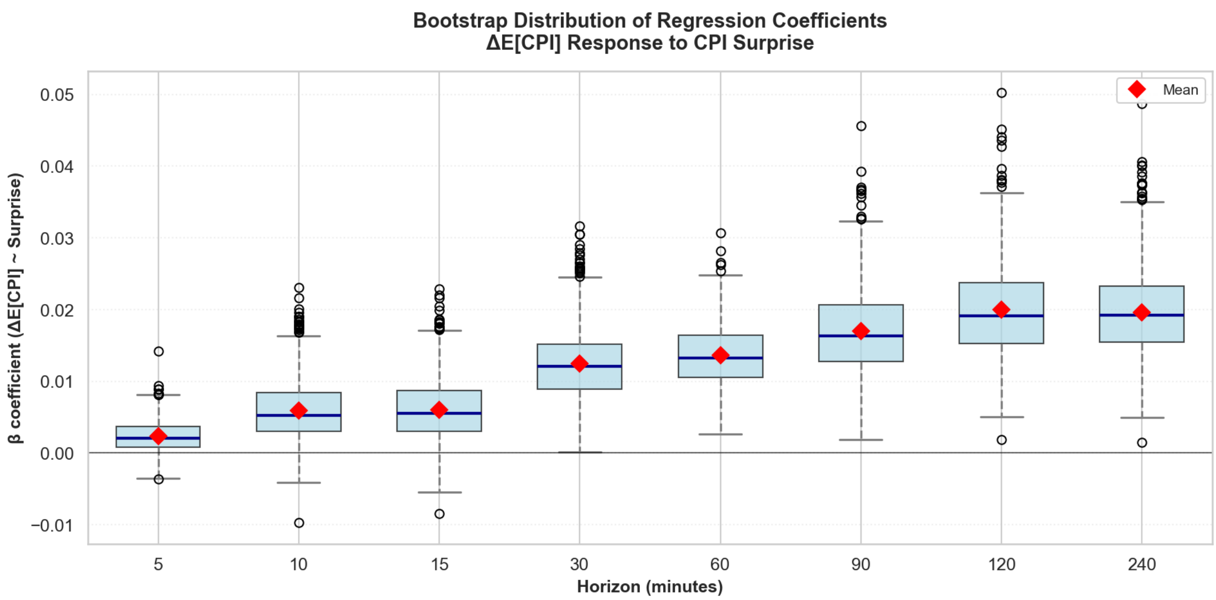

Figure 4: Block-bootstrap distributions of $\hat{\beta}_h$ for the mean-updating regressions across horizons.

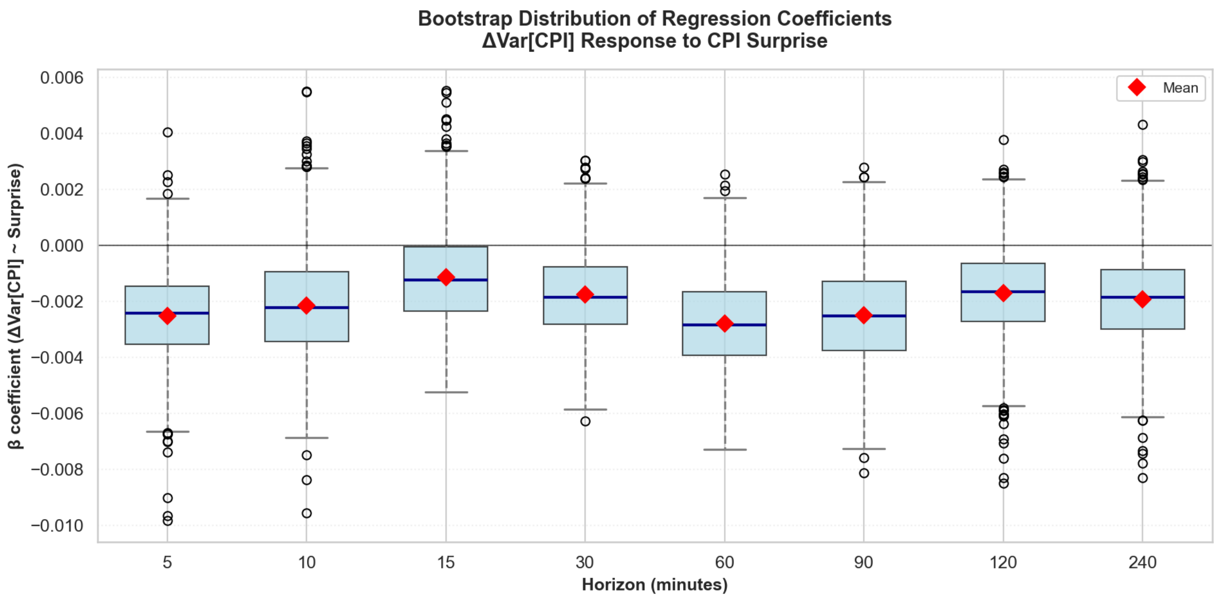

Figure 5: Block-bootstrap distributions of $\hat{\gamma}_h$ for the variance regressions across horizons.

8 Threats to Identification and Limitations

The primary threats are measurement and contamination rather than reverse causality. Thin liquidity and sparse strike coverage can generate noisy moments, particularly variance, which is sensitive to tail assumptions and bracket sparsity. Prediction-market prices may deviate from true probabilities due to risk preferences and trading frictions (Manski, 2006); differencing reduces the relevance of level biases, but short-window changes may still reflect liquidity premia that shift around announcements. Coincident news could contaminate the window; short horizons and placebo diagnostics provide partial protection, but do not eliminate this concern (Faust et al., 2007).

9 Conclusion

Minute-frequency Kalshi CPI markets allow direct measurement of belief updating around CPI announcements. CPI surprises predict economically meaningful revisions in the implied mean of CPI MoM beliefs at horizons of 30 to 240 minutes. Coefficient-equality tests provide no evidence that adjustment continues beyond 90 minutes and imply completion by 120 minutes: β_{120} and β_{240} are statistically indistinguishable ($p = 0.975053$). In contrast, implied variance responses are small and not precisely estimated at any horizon. Placebo tests, anticipation diagnostics, and bootstrap summaries support the interpretation that the estimated mean response reflects announcement-driven updating rather than generic intraday drift.

References

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Vega, C. (2003). Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review*, 93(1):38–62.

Balduzzi, P., Elton, E. J., and Green, T. C. (2001). Economic news and bond prices: Evidence from the US Treasury market. *Journal of Financial and Quantitative Analysis*, 36(4):523–543.

Bauer, M. D. and Swanson, E. T. (2023a). An alternative explanation for the Fed information effect. *American Economic Review*, 113(3):664–700.

Bürgi, C., Deng, W., and Whelan, K. (2025). Makers and takers: The economics of the Kalshi prediction market. Working paper.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.

Faust, J., Rogers, J. H., Wang, S.-Y. B., and Wright, J. H. (2007). The high-frequency response of exchange rates and interest rates to macroeconomic announcements. *Journal of Monetary Economics*, 54(4):1051–1068.

Gürkaynak, R. S., Sack, B., and Wright, J. H. (2010). The TIPS yield curve and inflation compensation. *American Economic Journal: Macroeconomics*, 2(1):70–92.

Manski, C. F. (2006). Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425–429.

Swanson, E. T., Wang, R., and Wu, Y. (2025). The effects of monetary policy on macroeconomic expectations: High-frequency evidence from traded event contracts. Working paper.