

Market-Implied Policy-Outcome Dispersion and Long-Run Equity Volatility

Meeting-Level Evidence from Kalshi FOMC Contracts

Fuli Yang | GARCH Institute | Cambridge Massachusetts | dr@garch.ac

September 2026

Abstract

This paper tests whether dispersion in market-implied Federal Open Market Committee outcomes contains information about the long-run component of U.S. equity volatility. For 38 meetings, I construct the cross-outcome variance of Kalshi probabilities over days -30 to -3 and embed it in a fixed ($K=4$), uniform-weight GARCH-MIDAS model for 824 S&P 500 trading days. The raw loading is 0.3621 (LR = 7.9637; asymptotic $p = 0.00477$; 999-draw meeting randomization $p = 0.027$). Residualizing on the implied policy mean and log volume gives 0.3570 (LR = 8.5072; $p = 0.00354$; randomization $p = 0.027$). A stronger support audit conditions on the number and span of outcome bins that the exchange lists for each meeting. Its loading is 0.3675 (LR = 9.5400; $p = 0.00201$), and a conditional meeting randomization that re-estimates the support adjustment in every draw gives $p = 0.018$. Activating the factor three business days before, on, or one business day after the meeting yields positive loadings with p -values of 0.0041, 0.0020, and 0.0197. Entropy is insignificant, showing that economic distance across rate outcomes—not probability fragmentation alone—carries the signal. All leave-one-meeting-out variance loadings remain positive. The evidence is an in-sample relation between pre-decision policy-outcome dispersion and the fitted long-run volatility state; a short 2026 holdout does not improve forecasts.

Keywords: prediction markets, monetary policy, GARCH-MIDAS, volatility, event contracts

JEL codes: C22, E44, E52, G14

1 Introduction

Expectations about monetary policy are usually summarized by a mean, a modal action, or an announcement surprise. Those summaries discard the shape of the expected outcome distribution. Two meetings can have the same expected policy rate while assigning very different probability mass to separated outcomes. The second distribution contains more policy-outcome dispersion even though its first moment is unchanged.

This paper asks whether that dispersion enters the slow-moving component of U.S. equity volatility. Kalshi lists event contracts tied to specified FOMC decisions. Prices across outcome bins can be assembled into a market-implied probability distribution before each meeting. I measure its cross-outcome variance over a fixed pre-decision window and use that meeting-frequency series as a long-run state variable in a GARCH-MIDAS model.

The central empirical issue is effective sample size. Daily S&P 500 returns provide 824 observations for estimating conditional variance dynamics, but the dispersion factor changes at only 38 meetings. The paper therefore distinguishes the estimation sample from the identifying sample. Likelihood-based daily inference is reported, but the primary robustness test randomizes factor values across meeting labels. This preserves the number and timing of meetings while breaking the mapping between a meeting's policy distribution and subsequent volatility state.

The raw Kalshi-implied dispersion loading is 0.3621, with $LR = 7.9637$ and asymptotic $p = 0.00477$. The meeting-label p -value is 0.027. A second specification removes the linear meeting-level association of dispersion with the implied policy mean and log contract volume before it enters the MIDAS filter. The residualized loading is 0.3570, $LR = 8.5072$, asymptotic $p = 0.00354$, and meeting-label $p = 0.027$. The residualized and raw factors correlate 0.982, indicating that the result is not driven by those two controls.

Influence analysis reinforces the finding. Each of 38 meetings is omitted in turn and every estimated loading remains positive. Thirty-seven residualized deletions retain asymptotic significance at five percent. The influential exception is 18 December 2024: omitting that meeting lowers the loading but does not reverse it. The economically important diagnostic is sign and magnitude stability; the deletion p -values remain model-based rather than a second set of meeting-randomization tests.

The paper's contribution is to show that a full market-implied policy distribution contains information not summarized by its mean. The factor is deliberately called **Kalshi-implied policy-outcome dispersion**. Event-contract prices may reflect beliefs, risk preferences, liquidity, and market microstructure; the empirical object is the dispersion of the priced outcome distribution. Residualization against volume directly tests one important microstructure channel.

This contribution connects three literatures. Prediction-market research asks when contract prices can be interpreted as probabilities, how well they are calibrated, and what information aggregation they provide (Wolfer and Zitzewitz 2004; Manski 2006; Page and Clemen 2013). Monetary-policy event studies measure asset-price responses to unexpected decisions and to the information conveyed by central-bank communication (Bernanke and Kuttner 2005; Gürkaynak, Sack, and Swanson 2005). GARCH-MIDAS models separate fast return dynamics from a persistent variance component linked to lower-frequency macroeconomic conditions (Engle, Ghysels, and Sohn 2013). The present paper links a complete pre-meeting policy distribution to that persistent component.

The closest measurement contribution is Goel (2026), who develops market-implied dispersion from Kalshi as a measure of macroeconomic surprise magnitude. I apply the distributional idea to recurring FOMC outcome contracts and a different dependent object: the persistent component of equity volatility. Relative to broad text-based policy uncertainty measures such as Baker, Bloom, and Davis (2016), the factor is meeting-specific, forward-looking, and defined on an explicit policy-outcome support. Vähämaa (2011) links Federal Reserve decisions to implied volatility; this paper instead asks whether dispersion visible before the decision helps explain a lower-frequency realized-return volatility state.

The distinction from an announcement-surprise design is central. An announcement surprise is realized at the decision and compares the action with a pre-decision expectation. My factor is known before the decision and measures how dispersed the priced alternatives remain. It can therefore be used as a state variable without importing the announcement outcome into the predictor. The distinction from a generic uncertainty index is equally important: the factor is tied to an enumerated set of mutually exclusive policy outcomes for a specific meeting.

The empirical claim is strong but tightly defined. Kalshi-implied variance has a positive loading in a fixed-architecture long-run volatility model; the result survives meeting-level inference, removal of each meeting, residualization against the first moment and volume, control for the exchange's listed outcome menu, and alternative activation dates. Entropy does not carry the same signal. This contrast distinguishes distance-sensitive policy uncertainty from generic fragmentation of probability mass. The short holdout answers a separate forecasting question with few new meetings. Sections 2 and 3 develop the measurement and inference. Section 4 reports the estimates and robustness results. Section 5 gives the economic interpretation, and Section 6 concludes.

2 Institutional Setting and Measurement

2.1 FOMC event contracts

Kalshi's FEDDECISION contract concerns whether the FOMC changes or maintains the target federal funds rate range at a specified meeting. The exchange lists recurring meeting-specific iterations. Contracts have a one-dollar settlement value, one-cent minimum tick, and stop trading shortly before the decision. Settlement follows the Federal Reserve's documented action.

Each outcome security pays one dollar if its stated policy bin occurs and zero otherwise. For a set of exhaustive, mutually exclusive outcomes, prices provide a directly observable cross-section from which to construct a meeting distribution. Normalization absorbs small departures from adding to one that can arise from bid-ask spreads, asynchronous observations, or market frictions. The resulting probabilities are therefore a market-implied distribution over the listed outcomes, not a survey transformation.

The distribution contains at least four distinct objects. Its mean locates the expected policy level, its modal bin identifies the most heavily priced outcome, its variance measures probability-weighted distance from the mean, and entropy measures fragmentation without regard to the distance between bins. A meeting can have an unchanged mean but a larger variance if mass shifts from the center toward opposing tails. Two meetings can also have similar entropy while placing mass on outcomes that are very different distances apart. This distinction motivates the variance-versus-entropy audit.

For a given meeting and trading date, mutually exclusive outcome-contract prices are mapped into probabilities p_{jmt} over policy outcomes x_{jm} . Daily probabilities are normalized to sum to one across the available outcome bins. The daily implied mean and variance are

$$\mu_{mt} = \sum_j p_{jmt} x_{jm}, v_{mt} = \sum_j p_{jmt} (x_{jm} - \mu_{mt})^2.$$

The meeting factor is the average of v_{mt} from calendar day -30 through day -3 relative to the FOMC decision. Every meeting contributes 28 daily observations. Excluding days -2 through 0 limits mechanical convergence as settlement approaches and fixes the information window before the announcement. The factor is standardized across meetings.

The support audit separates the exchange's *listed menu* from the distribution's *active probability mass*. Listed-bin count and listed support span describe how many outcomes the contract menu makes available and how far that menu extends. Active-bin count and active span are outcomes receiving positive probability; they are manifestations of dispersion itself and are not treated as nuisance controls. The support-adjusted factor residualizes variance on the policy mean, log volume, listed-bin count, and listed span. Entropy is standardized as a separate competing factor.

2.2 Sample and quality controls

The meeting panel contains 38 FOMC decisions. The corresponding S&P 500 return sample runs from 14 June 2023 through 25 September 2026 and contains 824 trading days after return construction. Table 1 separates the two sample dimensions.

Table 1 Design and data construction

Item	Definition
Identifying sample	38 FOMC meetings
Conditional-variance sample	824 S&P 500 trading days
Factor window	Calendar days -30 to -3; 28 observations per meeting
Raw factor	Standardized mean cross-outcome variance of Kalshi-implied probabilities
Base residualized factor	Raw variance residualized on policy mean and $\log(1 + \text{contract volume})$, then standardized
Support-adjusted factor	Variance additionally residualized on listed-bin count and listed support span
Alternative shape factor	Standardized Shannon entropy of normalized outcome probabilities
MIDAS architecture	$K = 4, \omega = 1$; fixed before the v2.1 inference audit
Information timing	Kalshi factor active on the meeting date; statement text active next trading day
Meeting inference	999 meeting-label permutations; leave one meeting out

Contract volume is averaged over the same -30 to -3 window. The residualized factor is constructed by regressing standardized meeting dispersion on an intercept, the meeting’s average implied policy mean, and log one plus average contract volume. Its correlation with the raw factor is 0.982. Thus, the second specification tests dispersion that is orthogonal to both the first moment of policy expectations and a direct trading-activity measure.

The 28-observation window is identical across meetings. This balances two requirements: enough pre-meeting observations to average noisy contract prices, and enough distance from settlement to avoid making the factor a measure of last-minute convergence. Every meeting passes the same observation-count rule. The audit table records the contract identifier, meeting date, raw dispersion, standardized dispersion, mean, volume, residualized dispersion, and source-window counts.

The sample spans tightening, pause, and easing environments. Consequently, dispersion is not mechanically a “high-rate” factor. In the meeting cross section, raw standardized dispersion correlates -0.131 with the implied policy mean and only 0.017 with average contract volume. The residualization is therefore a demanding specification check rather than a wholesale redefinition of the factor.

3 Empirical Design

Daily percentage log returns follow

$$r_t = \mu + \sqrt{\tau_t} g_t z_t, z_t \sim N(0, 1),$$

where g_t is the short-run GARCH component and τ_t is long-run variance. The long-run component is

$$\tau_t = \exp\left(m + \theta X_t^{(4,1)}\right).$$

$X_t^{(4,1)}$ is the uniform-weight average of the four most recent meeting factors. $K = 4$ and $\omega = 1$ are held fixed throughout the present audit; no grid search is used to select the reported model. Gaussian quasi-maximum likelihood estimates the restricted model and the one-factor alternative on the same daily sample.

The four-meeting window corresponds approximately to half a year of scheduled FOMC decisions. Uniform weights make the state transparent: each of the four most recent meeting distributions contributes one quarter of the slow-moving signal. Fixing both the horizon and weight shape before the inference audit prevents the reported likelihood ratio from benefiting from a search over many MIDAS kernels.

The restricted and factor models share the same return data and short-run variance structure. The likelihood-ratio statistic is twice the difference in maximized log likelihood. The factor loading is the only added long-run state coefficient under the reported comparison. Because the factor is standardized at the meeting level, the loading has a stable interpretation across the raw and residualized specifications.

Because the factor is standardized, a one-standard-deviation increase changes long-run variance by $100[\exp(\theta) - 1]$ percent and long-run volatility by $100[\exp(\theta/2) - 1]$ percent. These transformations describe the fitted long-run component; they are not daily return effects.

3.1 Meeting-label randomization

For each of 999 draws, I permute the 38 factor values across the fixed meeting dates, reconstruct the MIDAS state, re-estimate the model, and store its LR statistic. The randomization p-value is

$$p = \frac{1 + \sum_{b=1}^{999} 1(LR_b \geq LR_{obs})}{1000}.$$

The test uses meetings—not days—as the exchangeable units. It directly asks whether the observed assignment of dispersion states to FOMC meetings produces an LR statistic unusually large relative to arbitrary assignments.

This randomization retains the entire daily return history, the spacing of FOMC meetings, and the persistence induced by the four-meeting MIDAS filter. It changes only which dispersion realization is attached to which meeting date. A p-value of 0.027 means that the observed alignment between meeting dispersion and the long-run volatility state is exceeded in 26 of the 999 randomized assignments under the plus-one convention.

3.2 Leave-one-meeting-out influence

The model is re-estimated after omitting each meeting's factor value in turn. The resulting distribution of (), LR statistics, and p-values identifies whether one meeting creates the relation. This is especially relevant with 38 low-frequency observations.

3.3 Support-conditional randomization and timing

The support audit recognizes that contract menus vary across meetings. For every one of 999 permutations, I reassign raw meeting variance, then re-estimate the residualization on policy mean, log volume, listed-bin count, and listed span before rebuilding the MIDAS state. This conditional procedure propagates estimation of the nuisance projection into the null distribution rather than treating the observed residuals as fixed.

Information timing is assessed by activating the same support-adjusted meeting state three business days before the decision, on the meeting date, and one business day afterward. The first convention is fully pre-decision, the second matches the main state dating, and the third tests whether a one-session dating choice drives the result.

4 Results

4.1 Raw and residualized dispersion

Table 2 reports the main results. Raw Kalshi-implied policy-outcome dispersion has a loading of 0.3621. Relative to the restricted model, log likelihood rises from -1015.8786 to -1011.8968, giving LR = 7.9637. The asymptotic p-value is 0.00477 and the meeting-label randomization p-value is 0.027.

The residualized factor produces a loading of 0.3570. Log likelihood rises to -1011.6251, giving LR = 8.5072, asymptotic p = 0.00354, and meeting-label p = 0.027. A one-standard-deviation increase in this factor is associated with a 42.9 percent increase in the fitted long-run variance component, or a 19.5 percent increase in its volatility scale.

Table 2 Fixed-architecture GARCH-MIDAS results and measurement audit

Factor	Days	Meetings	Loading	LR	Asymptotic p	Meeting permutation p
Raw Kalshi dispersion	824	38	0.3621	7.9637	0.00477	0.027
Mean-and-volume residualized variance	824	38	0.3570	8.5072	0.00354	0.027
Listed-support-adjusted variance	824	38	0.3675	9.5400	0.00201	0.018
Entropy	824	38	0.1211	0.6867	0.407	0.553

The 95th percentile of the residualized-factor permutation null is 6.6386, below the observed LR of 8.5072. The 99th percentile is 10.6692. The observed statistic therefore lies in the upper tail without relying on 824 days as independent factor observations.

The listed-support audit strengthens the result. Raw variance has only a 0.196 correlation with listed menu width, and conditioning on that menu produces a loading of 0.3675 and LR of 9.5400. The conditional randomization p-value is 0.018; the 95th percentile of its LR null is 7.1086. Variation in which bins receive positive probability is not controlled away because it is precisely how a priced distribution becomes dispersed.

The raw and residualized estimates are unusually close. Residualization lowers the loading by only 0.0051, or 1.4 percent of the raw estimate, while the LR statistic rises by 0.5435. The conclusion is therefore not an

artifact of choosing between two related factor definitions. Both specifications say that meetings with wider priced policy distributions are followed by a higher fitted long-run equity-volatility state.

The exponential link gives the loading an economically interpretable scale. For the raw factor, one standard deviation corresponds to a 43.6 percent increase in fitted long-run variance and a 19.9 percent increase in fitted long-run volatility. For the residualized factor, the corresponding figures are 42.9 and 19.5 percent. These are changes in the persistent component conditional on the GARCH-MIDAS model, not predictions of a same-day absolute return.




Figure 1. Panel A compares the residualized factor's observed LR statistic with 999 meeting-label permutations. Panel B reports the loading after omitting each meeting. Blue points are significant at five percent; the orange point omits 18 December 2024.

4.2 Influence and state dependence

All 38 leave-one-meeting-out loadings are positive. For the residualized model, 37 deletions retain $p < 0.05$. Omitting 18 December 2024 gives the largest p-value, approximately 0.175, and the smallest loading, approximately 0.226. The meeting is influential in magnitude, but its removal does not change the sign.

This is the expected form of influence for a state variable. A meeting associated with an unusually informative policy distribution should contribute meaningfully to the estimated state transition. The relevant diagnostic is whether it single-handedly reverses the result. It does not. The remaining 37 meeting deletions produce significant positive loadings, including meetings before and after the influential episode.

The 18 December 2024 meeting has the largest raw dispersion in the sample, 2.727 standard deviations above the meeting mean. Its average contract volume is not the sample maximum, which helps explain why residualizing on volume does not remove it. Omitting the meeting reduces the residualized loading from 0.357 to 0.227 but leaves it positive. Across all omissions, residualized loadings range from 0.227 to 0.436; 37 of 38 have asymptotic p-values below 0.05.

The raw-factor leave-one-out exercise gives the same sign stability. Its loadings range from 0.203 to 0.432 and all are positive. Thirty-six of the 38 raw deletions remain significant at five percent. The residualized specification improves the deletion profile, which is consistent with mean and volume controls removing noise rather than creating the relation.

4.3 Information beyond the policy mean and trading volume

The residualization test changes the factor only modestly. Its 0.982 correlation with raw dispersion and slightly larger LR statistic show that the relation is not generated by the direction of expected policy or by meetings with greater contract volume. The empirical content lies in the second moment of the priced outcome distribution.

The entropy comparison identifies what kind of second-moment information matters. Entropy's loading is 0.1211, its LR statistic is 0.6867, and its asymptotic and meeting-randomization p-values are 0.407 and 0.553. Entropy rises when probability mass is spread across more bins but does not encode whether those bins are adjacent or far apart. The positive variance result alongside a null entropy result therefore points to economically separated rate outcomes rather than generic probability fragmentation.

The support-adjusted result is also stable to factor dating. Activating it three business days before the meeting gives a loading of 0.3447 (LR = 8.2388; $p = 0.00410$). Meeting-date activation gives 0.3675 (LR = 9.5400; $p = 0.00201$), and activation one business day after gives 0.3138 (LR = 5.4424; $p = 0.01965$). No single convention determines the sign or statistical conclusion.

Earlier common-sample analysis also compares the factor with FOMC statement-language uncertainty under actual information timing. On 36 overlapping meetings, Kalshi dispersion remains incrementally significant conditional on text ($p = 0.0367$), while the text factor's incremental p-value conditional on Kalshi is 0.2384. The market factor enters on the meeting date because it is observed beforehand; statement language enters on the next trading day because it is released with the decision.

This timing discipline matters. Assigning statement text to the meeting date would allow information released with the decision to explain a state dated before market participants could observe the text. By contrast, the Kalshi window closes three days before the decision, so its information set is unambiguously pre-announcement. The common-sample comparison therefore favors neither factor through look-ahead timing.

4.4 Forecasting as a separate application

A strict 2026 holdout contains 106 test days. Its restricted-model QLIKE is 0.8688, compared with 0.8724 for the Kalshi variance model; the HAC p-value for improvement is 0.7506. The joint text-plus-Kalshi model has QLIKE 0.8697. The current evidence therefore establishes in-sample information content for the fitted long-run component, not short-horizon forecast superiority. This distinction is substantive: the meeting-level test concerns whether policy-outcome dispersion maps into the estimated volatility state, whereas the holdout asks whether a small number of later meetings improves a noisy daily loss measure.

The holdout result also clarifies what additional data can resolve. Forecast evaluation at the factor's natural frequency requires many new FOMC decisions, not merely more trading days between a small number of meetings. As the event-contract history lengthens, a rolling or expanding-window design can estimate the factor loading using one set of meetings and evaluate long-run variance forecasts over a genuinely separate set of policy cycles.

5 Interpretation

Cross-outcome variance measures how broadly the event market prices separated policy actions. It can rise with the expected policy rate unchanged, and it can remain elevated even when one action is modal. That feature differentiates it from a policy surprise and from the distribution mean.

The positive loading is consistent with a macro-financial state in which economically separated policy outcomes remain priced before the decision. Such states can coincide with greater uncertainty about inflation, growth, and the policy reaction function, all of which can propagate into persistent equity volatility. The residualized specifications establish that the relation is not reducible to mean policy expectations, contract volume, or the width of the listed contract menu. The entropy comparison further shows that distance across outcomes is essential.

The result should be read at the frequency at which it is identified. Thirty-eight meetings supply the low-frequency state changes. The 824 return observations discipline the GARCH dynamics around those states.

Meeting-label inference, leave-one-meeting-out stability, and a fixed MIDAS architecture make that separation explicit.

The factor also has a practical advantage over a single policy-probability series. Contract menus can change as the relevant range of policy actions changes, but the first two moments remain comparable after outcomes are mapped to a common policy scale. A dispersion factor can therefore summarize uncertainty when the modal outcome shifts from hikes to holds or cuts. The construction uses the full cross section rather than selecting one contract whose economic meaning changes across cycles.

For risk monitoring, the result suggests treating the priced policy distribution as a macro-financial state input. A high-dispersion meeting does not imply that volatility must jump on announcement day. It indicates that the persistent variance component is estimated to be higher when multiple separated policy outcomes remain priced in advance. This distinction is useful for horizon allocation: short-run GARCH dynamics absorb immediate return shocks, while the meeting factor shifts the slower variance baseline.

The evidence also disciplines interpretation of prediction-market probabilities. The paper does not require prices to equal physical probabilities exactly. Any stable transformation involving risk premia, hedging demand, or liquidity can still produce a useful priced-distribution factor. What the residualization test establishes is narrower and observable: the second-moment signal is not a disguised linear function of the policy mean or average trading volume.

6 Conclusion

Kalshi-implied policy-outcome variance is positively related to the long-run component of U.S. equity volatility. Raw, mean-and-volume residualized, and listed-support-adjusted factors are significant under likelihood-based and meeting-randomization inference. The strongest audit produces $LR = 9.5400$ and conditional randomization $p = 0.018$. All leave-one-meeting-out loadings remain positive, and the result survives pre-meeting, meeting-date, and post-meeting activation conventions.

The evidence supports a precise claim: pre-decision dispersion in market-implied FOMC outcomes contains in-sample information about the estimated long-run volatility state. It does not depend on counting 824 daily returns as 824 independent policy observations, and it does not rest on a selected weighting grid. The next empirical frontier is broader out-of-sample validation across additional meetings and volatility environments.

The central result is statistically and economically material. A one-standard-deviation mean-and-volume residualized factor maps into approximately 19.5 percent higher fitted long-run volatility. Entropy is null, so the finding is not that any diffuse probability distribution predicts volatility. It identifies the probability-weighted economic distance among priced policy outcomes as a distinct state variable for equity volatility.

Declarations

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflicts of interest

The author declares no competing financial or non-financial interests.

Data and code availability

The replication package contains the meeting-level factor audit, processed S&P 500 return inputs, permutation draws, leave-one-meeting-out estimates, model summaries, and estimation code. Kalshi contract records and market data remain attributable to their original sources and applicable access terms; the package provides the processed factors, contract identifiers, dates, observation counts, and source manifest needed to audit the reported construction.

Use of generative AI

Generative AI tools were used for limited coding assistance and minor editorial support. All research design, analysis, interpretation, and final content were determined and verified by the author.

References

- Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. “Measuring Economic Policy Uncertainty.” *Quarterly Journal of Economics* 131 (4): 1593–1636. <https://doi.org/10.1093/qje/qjw024>
- Bernanke, Ben S., and Kenneth N. Kuttner. 2005. “What Explains the Stock Market’s Reaction to Federal Reserve Policy?” *Journal of Finance* 60 (3): 1221–1257. <https://doi.org/10.1111/j.1540-6261.2005.00760.x>
- Engle, Robert F., Eric Ghysels, and Bumjean Sohn. 2013. “Stock Market Volatility and Macroeconomic Fundamentals.” *Review of Economics and Statistics* 95 (3): 776–797. https://doi.org/10.1162/REST_a_00300
- Goel, Tanuj. 2026. “Measuring Macroeconomic Surprise Magnitude with Prediction Markets: Market-Implied Dispersion from Kalshi.” *International Review of Economics & Finance* 110: 105577. <https://doi.org/10.1016/j.iref.2026.105577>
- Gürkaynak, Refet S., Brian Sack, and Eric Swanson. 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. <https://www.ijcb.org/journal/v1n1/do-actions-speak-louder-words-response-asset-prices-monetary-policy-actions-and>
- KalshiEX LLC. 2023. “FEDDECISION Contract Certification and Rulebook.” Filing under CFTC Regulation 40.2(a). <https://assets.kalshi.com/regulatory/product-certifications/FEDDECISION.pdf>
- Manski, Charles F. 2006. “Interpreting the Predictions of Prediction Markets.” *Economics Letters* 91 (3): 425–429. <https://doi.org/10.1016/j.econlet.2006.01.004>
- Page, Lionel, and Robert T. Clemen. 2013. “Do Prediction Markets Produce Well-Calibrated Probability Forecasts?” *Economic Journal* 123 (568): 491–513. <https://doi.org/10.1111/j.1468-0297.2012.02561.x>
- Vähämaa, Sami. 2011. “The Fed’s Policy Decisions and Implied Volatility.” *Journal of Futures Markets* 31 (10): 995–1009. <https://doi.org/10.1002/fut.20503>
- Wolfers, Justin, and Eric Zitzewitz. 2004. “Prediction Markets.” *Journal of Economic Perspectives* 18 (2): 107–126. <https://doi.org/10.1257/0895330041371321>