When Prediction Markets Disagree: Cross-Platform Dispersion and Equity Volatility in the 2024 U.S. Election

Emrehan Aktuğ
Sabancı University

Orhan Torul
Boğaziçi University

July 28, 2026

Abstract

We ask whether disagreement across prediction markets helps predict the volatility of politically sensitive equities during the 2024 U.S. presidential election. Using 5-minute odds from seven betting platforms and intraday prices for 71 politically exposed firms, we construct hourly measures of cross-platform dispersion and realized volatility. Higher disagreement predicts higher volatility in a partisan long–short portfolio: after controlling for broad realized and implied market volatility, a one-standard-deviation increase raises spread volatility by between 4% and 11%, depending on the intraday price source, with the range reflecting vendor differences in quote coverage and timestamp conventions. A formal selectivity test shows that the relationship is significantly stronger for politically exposed equities than for the broad market or a clean non-political (utilities and consumer-staples) benchmark. Consistent with a predominantly second-moment, uncertainty-based channel, disagreement predicts absolute and squared returns far more strongly than signed returns, and its effect persists for roughly one trading day. We interpret the evidence as selective predictive content specific to the 2024 episode rather than as a causal estimate or a universal pattern.

Keywords: Prediction Markets; Political Uncertainty; Equity Volatility; Disagreement; Heterogeneous Beliefs; Price Discovery
JEL Classifications: G14; G12; D84

1 Introduction

During the 2024 U.S. presidential election, at least seven global platforms simultaneously quoted odds on the outcome. These platforms often disagreed. The cross-sectional standard deviation of implied Trump-win probabilities averaged 2.6 percentage points over our sample and more than doubled around the largest campaign shocks, exceeding 6 percentage points after the July assassination attempt and after President Biden's withdrawal. This paper asks whether such cross-platform disagreement matters for equity markets.

A large literature shows that political uncertainty affects asset prices (Pástor and Veronesi, 2012, 2013; Kelly et al., 2016; Baker et al., 2016). A smaller literature studies information aggregation in prediction markets (Arrow et al., 2008; Wolfers and Zitzewitz, 2004). Yet the link between prediction-market *disagreement* and equity-market *volatility* remains unexplored. Our motivation draws on the heterogeneous-beliefs tradition (Miller, 1977; Diether et al., 2002; Hong and Stein, 2007; Banerjee and Kremer, 2010): when even active markets fail to converge on the likely winner, the range of policy scenarios that investors must price

widens, raising the volatility of firms whose fortunes depend on the electoral outcome. Cross-platform dispersion is distinctive because it is measured from traded prices on the *same* event across venues at high frequency—unlike the consensus probability (a first moment), polling dispersion (low-frequency, survey-based), or option-implied volatility (which bundles election risk with macroeconomic and risk-premium components). This dispersion is a signal of a latent ambiguity state that raises the conditional variance of future probability revisions.

We study this link with a selectivity design. If cross-platform disagreement captures election-specific uncertainty, it should predict the volatility of firms exposed to the electoral outcome more strongly than it predicts the volatility of the broad market. This is what we find in the data. Disagreement strongly predicts the volatility of a partisan long–short equity spread, and a formal interaction test confirms that this difference is statistically significant. By contrast, its relationship with broad market volatility is much weaker once standard controls are included. Local projections show that this predictive effect persists for roughly one trading day, while the corresponding market-wide effect is small and imprecisely estimated.

Our contribution differs from work that links the level or change in a single platform's probability to equity returns (Addoum and Kumar, 2016; Aktuğ and Torul, 2026a). That channel asks *which way* the election is going. We instead study a second-moment object: *how much* markets disagree, and whether that disagreement predicts the volatility of politically exposed equities. The selectivity design helps separate election-specific uncertainty from broader confounds. We sharpen it along three lines: we benchmark partisan-spread volatility against a clean non-political portfolio of utilities and consumer-staples firms; we separate the second-moment channel from a directional one by testing absolute against signed returns; and we validate the venue classification with platform-level price-discovery and quote-staleness diagnostics. If disagreement merely proxied for general news intensity or stale-quote dispersion, one would expect it to predict broad market volatility and partisan-spread volatility more similarly than we observe in the data. Accordingly, we interpret the evidence as documenting a selective predictive relation rather than as establishing a fully causal mechanism.[1]

2 Conceptual framework

The framework below is deliberately parsimonious; the Online Supplement provides its complete derivation. Let π_t be the latent probability that the Republican candidate wins. Platform j posts $p_{j,t} = \pi_t + u_{j,t}$, where the venue-specific term $u_{j,t}$ captures differences in beliefs and information processing across each platform's clientele—not mechanical microstructure noise—and has conditional variance equal to a latent ambiguity state A_t: when the electoral signal is harder to interpret, distinct clienteles diverge more and A_t rises. Cross-platform disagreement D_t—the dispersion of the $p_{j,t}$—then satisfies $\mathbb{E}(D_t^2 \mid A_t) = A_t$, so it is an observable, high-frequency signal of A_t. A single platform's quote cannot deliver this: one quote mixes π_t with venue noise, whereas dispersion across venues identifies a second-moment object.

Equity prices respond to *revisions* in π_t. If the conditional variance of the next revision is $\sigma^2_{\pi,0} + \kappa A_t$ with $\kappa > 0$, and firm i has political exposure λ_i, then

$$\mathrm{Var}\big(r_{i,t+1} \mid A_t\big) = \beta_i^2 \sigma^2_{m,t} + \lambda_i^2 \big(\sigma^2_{\pi,0} + \kappa A_t\big) + \sigma^2_{\varepsilon,i}. \tag{1}$$

Three implications guide the empirics. (i) Disagreement predicts volatility beyond the consensus probability and broad implied volatility, because it proxies for A_t. (ii) The effect is *selective*: it scales with λ_i^2, so it is larger for the partisan spread (large $|\lambda|$) than for the broad market or a non-political benchmark

[1] Our empirical design is informed by the market microstructure literature on price discovery across trading venues (Hasbrouck, 1995; Chakravarty et al., 2004).

($\lambda \approx 0$). (iii) Because revisions have zero conditional mean, the channel works predominantly through second moments—absolute and squared returns—rather than signed returns. We treat (1) as an interpretive lens rather than as a causal claim, and we test each implication below.

3 Data

We combine prediction-market and equity-market data from January 22 to November 5, 2024.[2] Prediction-market quotes and equity prices are sampled at 5-minute intervals; realized volatility is computed from these returns within each trading hour. Baseline regressions use the hourly frequency to reduce microstructure noise and mitigate the effects of stale quotes.

Our prediction-market sample contains implied Trump-win probabilities from seven platforms: PolyMarket, Betfair, Bovada, BetOnline, William Hill, Unibet, and Everygame. To distinguish between leading and lagging venues, we use Gonzalo–Granger (GG) weights from a companion price-discovery exercise estimated at the 5-minute frequency (Gonzalo and Granger, 1995; Aktuğ and Torul, 2026b; Bossaerts et al., 2024). GG weights decompose the common efficient price into each venue's contribution: a platform with a large positive weight contributes more to price discovery, while one with a near-zero weight tends to adjust with a lag.[3] PolyMarket, Betfair, and BetOnline receive positive and relatively large GG weights. Bovada, William Hill, Unibet, and Everygame receive near-zero or negative weights. We refer to these two groups as "price leaders" and "lagging" (lower-price-discovery) platforms, recognizing that the classification rests on a statistical decomposition rather than direct observation of trader types. Table 5 reports the GG weights alongside each platform's quote coverage, update frequency, and staleness: the price leaders update far more often and are markedly less stale than the lagging, bookmaker-style venues, supporting the classification on an independent microstructure dimension.

On the equity side, we use 5-minute prices for 71 publicly traded U.S. firms classified as Republican-leaning (36 firms) or Democrat-leaning (35 firms) by the direction of their overall partisan giving—contributions to candidates *and* to ideological groups—as compiled by OpenSecrets. The sample is constructed from OpenSecrets' industry-level rankings of the largest organization contributors to federal campaigns over the 2020–2024 election cycles: working through each industry's top-contributor list, we retain the publicly traded firms—many of the largest contributors are privately held—with complete high-frequency trading data throughout the sample period.[4] We construct equal-weighted RED and BLUE portfolios and define *SPREAD* as the RED-minus-BLUE return and *MKT* as the equal-weighted average across all 71 firms. As a clean placebo we also build a *non-political* portfolio: an equal-weighted basket of sixteen large utilities and consumer-staples firms with no partisan tilt, disjoint from the 71 exposed firms, using institutional-grade intraday prices from DataBento.[5] Our benchmark implied-volatility control is VIXY, an exchange-traded

[2] The sample begins on January 22—the first trading day after Ron DeSantis withdrew from the Republican primary—to ensure that prediction-market prices reflect general-election uncertainty rather than primary-nomination dynamics.

[3] Because disagreement is measured as cross-sectional dispersion rather than deviations from a single platform's quote, our main results do not depend on the GG classification. The classification is used only in the supplementary decomposition in Table 4.

[4] The Online Supplement lists every firm with its donation totals and the resulting Republican share; the share reproduces the binary lean for all but two firms (one with balanced giving and one, Tesla, with no PAC record), both assigned to the Republican group. The result is robust to a stricter tercile split that drops the middle third (Table 8) and to dropping the two judgment cases (Online Supplement).

[5] The non-political series spans slightly fewer hours than the political sample (1,368 versus 1,401) because of occasional intraday gaps in the DataBento quotes for individual constituents; the Online Supplement lists the constituents.

product linked to short-term VIX futures that is available at intraday frequency.[6]

After merging and restricting to U.S. trading hours (13:30–20:00 UTC), the sample contains more than 15,000 five-minute observations, 1,402 hourly observations, and 201 daily observations. Summary statistics are in Table A.1 in the Appendix.

4 Methodology

At each 5-minute timestamp we measure disagreement as the cross-sectional standard deviation of Trump-win probabilities across the seven platforms:

$$\text{DISAGREE}_t = \text{SD}\big(p_{1,t}, \ldots, p_{7,t}\big). \tag{2}$$

We aggregate this measure to hourly frequency and relate it to hourly realized volatility. For each trading hour we compute the standard deviation of the twelve 5-minute log returns within that hour, separately for the partisan spread and the broad market portfolio.[7]

Our baseline specification is:

$$\ln\big(\text{SPREAD vol}_t\big) = \alpha + \beta\,\text{DISAGREE}_t + \gamma \ln\big(\text{MKT vol}_t\big) + \delta \ln\big(\text{VIXY}_t\big) + \varepsilon_t, \tag{3}$$

where standard errors use the Newey–West HAC estimator with the Andrews bandwidth (Newey and West, 1987). Including both realized market volatility and VIXY absorbs broad volatility conditions, so that β captures whether cross-platform dispersion contains incremental information about politically exposed stocks beyond broad market volatility.

To trace the dynamics we estimate local projections following Jordà (2005):

$$\overline{y}_{t,t+h} = \alpha_h + \beta_h\,\text{DISAGREE_SHOCK}_t + \boldsymbol{\gamma}_h'\mathbf{Z}_t + \varepsilon_{t+h}, \tag{4}$$

where $\overline{y}_{t,t+h}$ is average log volatility from t to $t+h$, the shock is the standardized disagreement measure, and $\mathbf{Z}_t$ includes lagged volatility and lagged disagreement.[8]

To test the selectivity claim formally, we stack SPREAD and MKT volatility into a single panel and estimate:

$$\ln\big(\text{vol}_{i,t}\big) = \alpha + \beta_1\,\text{DISAGREE}_t + \beta_2\,D_i + \beta_3\,\text{DISAGREE}_t \times D_i + \boldsymbol{\gamma}'\mathbf{X}_{i,t} + \varepsilon_{i,t}, \tag{5}$$

where $D_i = 1$ for the partisan spread. In the controlled version, $\mathbf{X}_{i,t}$ includes $\ln(\text{VIXY}_t)$ interacted with D_i as a common implied-volatility control. Realized market volatility cannot serve as a common control in this stacked specification because it is itself the dependent variable for the $D_i = 0$ observations. The coefficient β_3 tests whether disagreement predicts partisan-spread volatility significantly more strongly than broad market volatility.

To probe the mechanism and identification, we run three further sets of regressions. First, we replace realized volatility on the left of (3) with the hourly *signed*, *absolute*, and *squared* partisan-spread return; the second-moment channel of Section 2 predicts a far stronger relation for absolute and squared returns than for signed returns. Second, we re-estimate the stacked test (5) with the non-political portfolio in place of the

[6] Results are materially unchanged when we replace VIXY with the standard CBOE VIX close-to-close change (Table 8). VIXY tracks expected volatility via VIX futures, so it may respond with a lag relative to spot-VIX shocks; this makes it a conservative control.

[7] Table 8 shows that conclusions are similar under alternative volatility proxies, including realized variance and mean absolute returns.

[8] The HAC bandwidth is set to $\max(h, \lfloor 4(T/100)^{2/9} \rfloor)$.

broad market, providing a clean placebo. Third, alongside the local projections we estimate a bivariate VAR in disagreement and spread volatility with a recursive (Cholesky) ordering that places disagreement first, and we examine event windows around the four largest campaign shocks of 2024. The Online Supplement adds variance-inflation diagnostics, a continuous Gonzalo–Granger-weighted disagreement measure, and cross-sectional dispersion tests.

5 Results

5.1 Descriptive evidence

Figure 1 plots the disagreement measure over time. Cross-platform dispersion rises around major campaign events: the Trump assassination attempt, President Biden's withdrawal, and Vice President Harris's emergence as the Democratic nominee. The series clearly co-moves with the political environment.

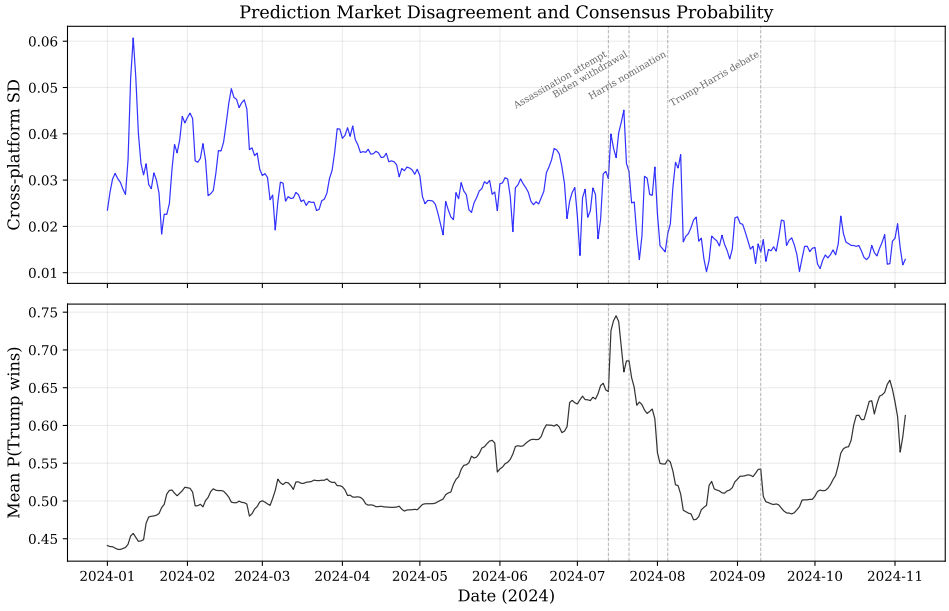


Figure 1: Prediction Market Disagreement and Consensus Probability

Notes: The top panel shows the daily mean cross-platform standard deviation of Trump-win probabilities. The bottom panel shows the daily mean probability across seven platforms. Dashed lines mark key events.

5.2 Core selectivity

Table 1 presents the main results. Column (1) shows that higher disagreement predicts higher volatility in the partisan spread. Column (2) reports a much weaker relationship for broad market volatility. In Columns (3) and (4), the disagreement coefficient remains economically large and statistically significant after controlling for both realized market-wide volatility and implied volatility. Adding a control for the mechanical level of probability uncertainty—which peaks when probabilities are near 50%—does not change the result.

The selectivity test in Equation (5) confirms this pattern formally. The interaction coefficient $\hat{\beta}_3$ is significant in both uncontrolled ($t = 3.76$) and controlled ($t = 3.85$) specifications: disagreement predicts partisan-spread volatility significantly more than broad market volatility.

In terms of magnitude, a one-standard-deviation increase in disagreement (0.93 percentage points) is associated with a 17% increase in partisan-spread volatility in the uncontrolled specification and an 11% increase after absorbing broad market volatility.[9]

This magnitude, however, depends materially on the intraday price source, and we report the range rather than a single number. Estimating the controlled specification (3) on our baseline series (Twelve Data) gives $\hat{\beta} = 12.11$ ($t = 6.69$), a one-standard-deviation response of about 11%; estimating the identical specification on institutional-grade DataBento prices for the same firms gives $\hat{\beta} = 4.19$ ($t = 2.53$), a response of about 4% (Table 8, Panel I). The sign and the statistical significance are common to both vendors, but the point estimate is not, and the economic magnitude should be read as the 4–11% range rather than as either endpoint.

The gap is not a firm-universe artifact: restricting the baseline to the 69 of our 71 firms that DataBento covers leaves it essentially unchanged ($\hat{\beta} = 11.51$, $t = 6.34$). It reflects the two vendors' intraday grids. Over U.S. trading hours about 95% of DataBento's 5-minute firm-bars are populated, against about 70% for Twelve Data, whose coverage rises from roughly 20% in January 2024 to above 99% from August onward; the baseline construction forward-fills the remainder within the trading day, so hourly realized volatility is built from a varying number of genuine 5-minute returns. The resulting hourly log-volatility series are correlated at 0.78. Reassuringly, within a trading day bar coverage is essentially uncorrelated with disagreement (correlation -0.06), so the difference does not arise from a mechanical high-frequency link between missing quotes and cross-platform dispersion. The two vendors also label 5-minute bars differently: the cross-correlation of their 5-minute market returns peaks at a one-bar displacement rather than at zero (0.34 against 0.03 at zero lag). Re-aligning the DataBento series to the baseline's convention raises its controlled estimate only from $\hat{\beta} = 4.19$ to $\hat{\beta} = 4.88$ ($t = 2.98$), so the labelling difference is real but accounts for little of the gap. We retain Twelve Data as the baseline because it covers all 71 sample firms, whereas DataBento covers 69, and we report the DataBento estimates alongside it rather than in place of it.

[9] With ln(vol) as the dependent variable, the percentage effect is $\exp(\hat{\beta} \times \sigma_{\text{disagree}}) - 1$, or approximately $\hat{\beta} \times \sigma_{\text{disagree}}$ for small values.

Table 1: Prediction Market Disagreement and Equity Volatility

	(1) SPREAD vol	(2) MKT vol	(3) SPREAD vol	(4) SPREAD vol
Disagreement	17.990***	6.347***	12.107***	12.466***
	(8.93)	(2.99)	(6.69)	(6.78)
ln(MKT vol)			0.711***	0.716***
			(26.29)	(26.38)
Prob. uncertainty				3.563**
				(2.19)
ln(VIXY)			0.453**	0.294
			(2.49)	(1.43)
Observations	1401	1401	1401	1401
R^2	0.065	0.011	0.452	0.454
Selectivity test: H_0: $\beta_{SPREAD} = \beta_{MKT}$				
Uncontrolled: $\hat{\beta}_{\text{interact}} = 11.644$, $t = 3.76$***				
Controlled: $\hat{\beta}_{\text{interact}} = 11.360$, $t = 3.85$***				

Notes: Dependent variable is ln(realized volatility) at hourly frequency. Disagreement is the cross-sectional standard deviation of probability levels across 7 prediction market platforms. Prob. uncertainty = $0.25 - (\bar{p} - 0.5)^2$. HAC (Newey–West) standard errors; t-statistics in parentheses. The selectivity test reports the coefficient on Disagree × SPREAD from a stacked regression of ln(vol) on disagreement, a SPREAD indicator, and their interaction. The controlled version adds ln(VIXY) interacted with the indicator as a common implied-volatility control. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The broad-market control in Table 1 is itself an average of politically exposed firms. Table 2 therefore adds a clean placebo: an equal-weighted portfolio of utilities and consumer-staples firms with no partisan tilt. Disagreement predicts the volatility of this non-political basket far less strongly than the partisan spread—about a 5% versus an 11% response to a one-standard-deviation shock—and the formal selectivity interaction remains significant against both the broad-market and the non-political benchmark. Disagreement is not orthogonal to non-political volatility: it carries a small market-wide component that is largely absorbed by implied volatility. But its predictive content is concentrated, as the framework implies, where political exposure is largest.

Table 2: Selectivity against a non-political benchmark

	(1) SPREAD vol	(2) MKT vol	(3) Non-pol. vol
Disagreement	17.990***	6.347***	5.629***
	(8.93)	(2.99)	(3.21)
Observations	1401	1401	1368
R^2	0.065	0.011	0.008
Selectivity ($\hat{\beta}$ on Disagree×SPREAD, controlled):			
SPREAD vs. MKT: 11.360, $t = 3.85$***			
SPREAD vs. non-political: 8.780, $t = 3.31$***			

Notes: Dependent variable is ln(hourly realized volatility). The non-political portfolio is an equal-weighted basket of 16 large utilities and consumer-staples firms (DataBento intraday prices) with no partisan tilt. Disagreement predicts non-political volatility far less strongly than partisan-spread volatility, and the selectivity interaction is significant against both the broad-market and the clean non-political benchmark. HAC t-statistics. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5.3 Mechanism: second moments, not direction

Table 3 tests the framework's sharpest implication. Disagreement predicts the *absolute* and *squared* partisan-spread return strongly but the *signed* return only weakly. The signed-return relation is an order of magnitude smaller in explanatory power and survives the inclusion of the consensus-probability change; a small, statistically significant positive signed component remains, reflecting that in 2024 cross-platform dispersion rose alongside Republican-favorable news, and perhaps a thin risk premium or skewness not captured by the consensus mean. We therefore read the channel as predominantly—though not exclusively—second-moment, consistent with an uncertainty rather than a directional interpretation.

Table 3: Mechanism: disagreement predicts second moments more than direction

	(1) Signed return	(2) Signed return + Δconsensus	(3) |Return|	(4) Return^2
Disagreement	1.255**	1.270**	3.079***	1.601***
	(2.52)	(2.56)	(7.86)	(6.69)
Δconsensus		0.790		
		(0.58)		
Observations	1400	1400	1400	1400
R^2	0.004	0.004	0.042	0.026

Notes: Hourly partisan-spread outcomes regressed on cross-platform disagreement. The dependent variable is the signed hourly return in Columns (1) and (2), its absolute value in Column (3), and its square in Column (4); Columns (1) and (2) therefore share the same dependent variable, and Column (2) differs from Column (1) only by adding the contemporaneous change in the consensus probability. Disagreement predicts absolute and squared returns far more strongly (an order of magnitude larger R^2) than signed returns, though a small positive signed-return relation remains. HAC t-statistics in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.10$.

5.4 Price leaders vs. lagging platforms

Table 4 decomposes disagreement by platform type. In a joint regression, disagreement among lagging, lower-price-discovery platforms is more strongly associated with partisan volatility than disagreement among price leaders. One interpretation is that slower-updating, retail-oriented venues reflect broader public uncertainty that is more closely aligned with equity-relevant risk—they drift apart precisely when information is hard to interpret—whereas price leaders incorporate information more quickly and contribute less marginal dispersion. We view this decomposition as suggestive rather than definitive: it does not rule out alternative explanations related to differences in liquidity, clientele, or updating frequency across platform types.

Table 4: Price-Leader vs. Lagging Platform Disagreement

	(1) Joint	(2) Leaders only	(3) Lagging only
Price-leader disagreement	5.133***	9.060***	
	(3.92)	(6.31)	
Lagging-platform disagreement	13.555***		15.759***
	(7.66)		(9.53)
Observations	1379	1379	1379
R^2	0.089	0.038	0.078

Notes: Dependent variable is ln(SPREAD vol) at hourly frequency. Price leaders (GG weight $\psi > 0.05$): PolyMarket, Betfair, BetOnline. Lagging, lower-price-discovery platforms ($\psi \leq 0.05$): Bovada, William Hill, Unibet, Everygame. HAC standard errors; t-statistics in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 5 supports the venue classification on an independent dimension. Betfair and PolyMarket—the two genuine markets—update far more frequently and are far less stale than the bookmaker-style venues; BetOnline is classified as a price leader by its Gonzalo–Granger weight despite slower quote updating. A continuous Gonzalo–Granger-weighted disagreement measure (Online Supplement) yields the same ranking—dispersion weighted toward lower-price-discovery venues carries a somewhat larger univariate loading—though the leader- and lagger-weighted measures are highly correlated, so we treat the decomposition as suggestive.

Table 5: Prediction-market platforms: price discovery and activity

Platform	GG weight ψ	Coverage	Updates/hr	Staleness	Group
Betfair	+1.171	1.00	0.64	19	Leader
PolyMarket	+1.061	1.00	0.13	93	Leader
BetOnline	+0.295	0.80	0.03	346	Leader
Bovada	+0.006	1.00	0.03	419	Lagging
William Hill	-0.007	0.61	0.01	920	Lagging
Everygame	-0.686	0.53	0.00	3571	Lagging
Unibet	-0.841	0.96	0.02	750	Lagging

Notes: Gonzalo–Granger weights ψ are normalized to sum to one across platforms and may fall outside $[0,1]$. Coverage is the fraction of 5-minute intervals with a valid quote; Updates/hr is the average number of quote changes per hour; Staleness is the mean run length (in 5-minute intervals) of unchanged quotes. Leaders ($\psi > 0.05$): PolyMarket, Betfair, BetOnline.

5.5 Dynamic effects

Figure 2 displays local-projection impulse responses. Panel A shows that a disagreement shock is followed by higher volatility in the partisan spread over roughly the next trading day. This persistence is consistent with the slow-moving nature of disagreement itself (half-life $\approx$ 33 hours): because platforms do not converge immediately, the information captured by cross-platform dispersion remains elevated over multiple hours. Panel B shows a much smaller and more imprecisely estimated response for broad market volatility. The contrast reinforces the selectivity finding: cross-platform dispersion is more tightly linked to election-sensitive assets than to market-wide conditions.

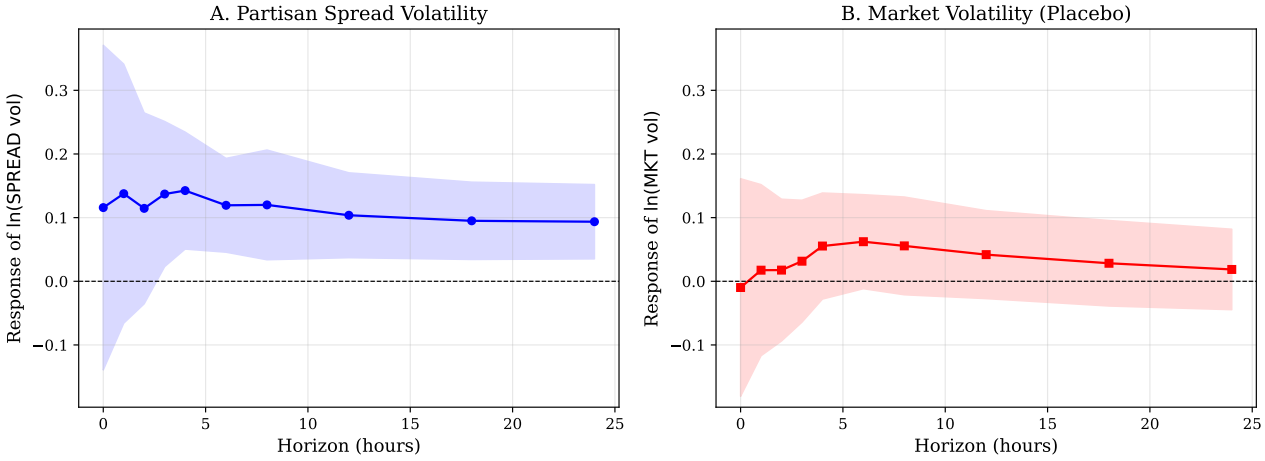


Figure 2: Local Projection IRFs: Disagreement Shock → Volatility

Notes: Each panel plots the local-projection impulse response (horizon in hours) to a one-standard-deviation cross-platform disagreement shock. Panel A is the response of partisan-spread realized volatility; Panel B is the response of broad-market realized volatility (MKT, the equal-weighted average of all 71 firms) and serves as a placebo—it is far smaller and statistically imprecise, underscoring that the effect is concentrated in the partisan spread. Controls include lagged log volatility and lagged disagreement; shaded areas are 95% confidence intervals based on HAC standard errors.

A recursive (Cholesky-ordered) bivariate VAR is directionally consistent, though less precisely estimated. The orthogonalized impulse response of spread volatility to a one-standard-deviation disagreement shock is positive and builds over the first trading day, although its confidence band includes zero at several horizons (Figure 3). The VAR's marginal contribution over the local projections is the joint test it permits: disagreement Granger-causes spread volatility (F-test $p = 0.03$) while the reverse does not ($p = 0.15$). The ordering is a timing convention, not a structural identification.

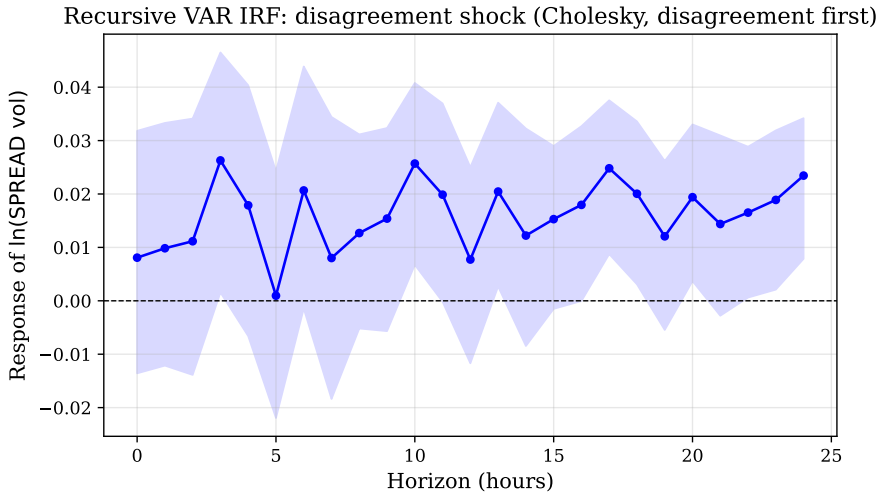


Figure 3: Recursive (Cholesky-ordered) VAR impulse response: disagreement shock → spread volatility

Notes: Orthogonalized (Cholesky) impulse response of ln(SPREAD vol) to a one-standard-deviation disagreement shock, from a bivariate VAR with disagreement ordered first. Lag length selected by AIC. The shaded area is a 95% bootstrap confidence band; it includes zero at several horizons. Horizons in hours.

5.6 Predictive ordering and persistence

Table 6 reports Granger-causality tests. Lagged disagreement helps predict partisan-spread volatility, whereas lagged volatility does not predict disagreement. We view this as evidence on temporal ordering rather than as a stand-alone causal design. An AR(1) specification further indicates that disagreement is highly persistent, with an implied half-life of approximately 33 hours.

Table 6: Granger Causality Tests (Hourly)

	Disagreement → SPREAD vol		SPREAD vol → Disagreement	
Lags	F-stat	p-value	F-stat	p-value
1	48.05***	0.0000	0.70	0.4016
2	30.11***	0.0000	0.32	0.7231
3	25.73***	0.0000	0.77	0.5134
4	21.29***	0.0000	0.78	0.5384
5	15.19***	0.0000	0.63	0.6795
6	7.65***	0.0000	1.57	0.1521

Notes: F-tests from hourly VAR. Null: lagged values of the row variable do not Granger-cause the column variable. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5.7 Event windows

Table 7 and Figure 4 summarize the four largest campaign shocks of 2024. Around the two biggest, disagreement and partisan-spread volatility move in the same direction: both rose sharply after the July 13 assassination attempt (disagreement +20%, spread volatility +48% across the surrounding 96-hour windows) and both fell after President Biden's withdrawal resolved a key source of uncertainty (disagreement −55%, spread volatility −8%). Around the two debates the comovement is weaker. As a formal check, adding an indicator for all four event days leaves the disagreement coefficient essentially unchanged ($\hat{\beta} = 11.9$, $t = 6.6$), so the baseline relation is not driven by a handful of salient days.

Table 7: Event windows around the four largest 2024 campaign shocks

		Disagreement			SPREAD vol		
Event	Date	Pre	Post	%Δ	Pre	Post	%Δ
Assassination attempt	Jul. 13	0.0299	0.0358	+20	0.0410	0.0607	+48
Biden withdrawal	Jul. 21	0.0407	0.0184	-55	0.0565	0.0520	-8
Biden–Trump debate	Jun. 27	0.0302	0.0212	-30	0.0382	0.0432	+13
Harris–Trump debate	Sep. 10	0.0167	0.0136	-19	0.0463	0.0417	-10

Notes: Mean cross-platform disagreement and partisan-spread realized volatility in the 96-hour windows before and after each event, split at the actual event time (UTC; the events occurred in the evening, U.S. time). Disagreement and spread volatility move sharply in the same direction around the assassination attempt and President Biden's withdrawal; the comovement is weaker around the debates. Adding an event-day indicator to the baseline does not displace the disagreement coefficient.

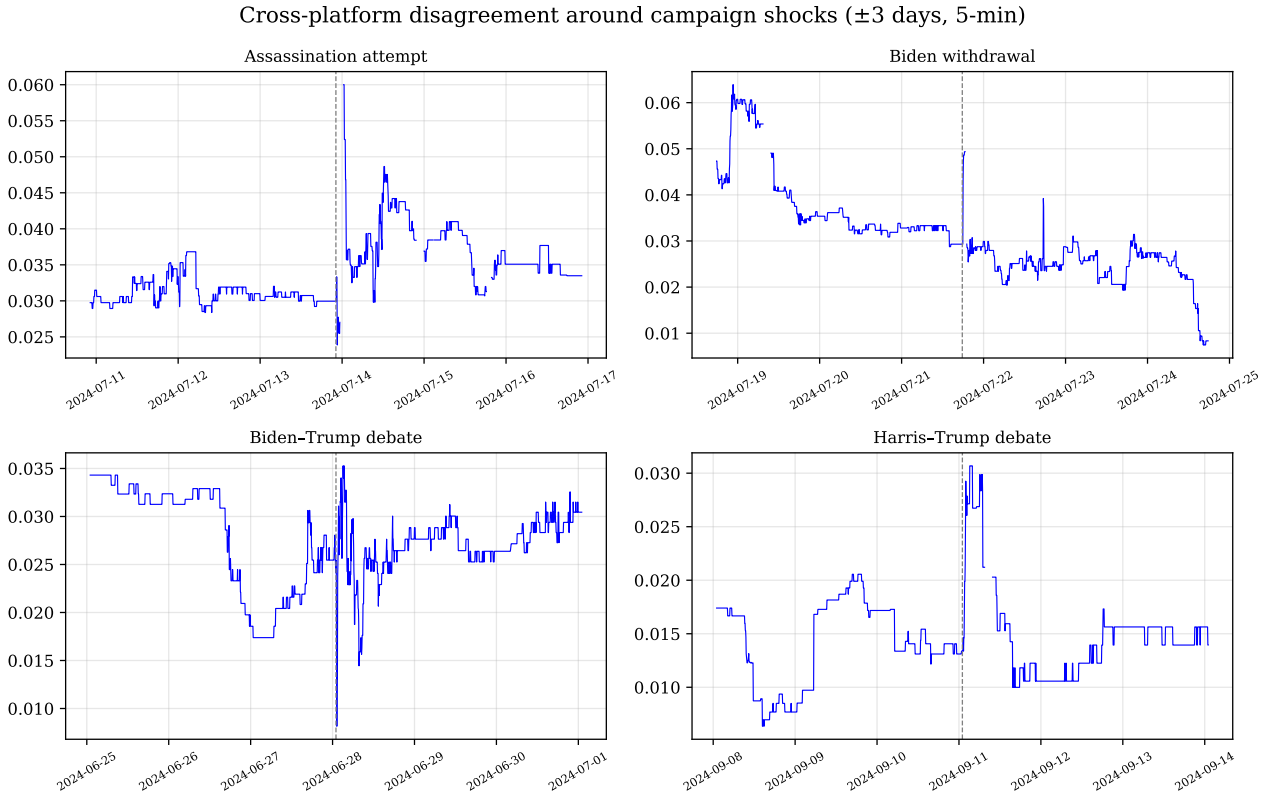


Figure 4: Cross-platform disagreement around campaign shocks

Notes: Cross-platform disagreement (5-minute) in ±3-day windows around the assassination attempt (July 13, 22:11 UTC), President Biden's withdrawal (July 21, 17:46 UTC), and the two presidential debates (the evenings of June 27 and September 10, each 01:00 UTC the following day). Dashed lines mark the actual event times; because the events occurred in the U.S. evening, the immediate reaction appears at the line rather than a day later.

5.8 Economic significance

Disagreement adds incrementally to standard controls, raising the R^2 of the hourly volatility regression from 0.42 to 0.45—a meaningful gain for an intraday volatility forecast. The clearest reading is in forecasting and risk-management terms: a volatility-targeting investor would trim a partisan position by roughly 4 to 10% following a one-standard-deviation disagreement shock, the range reflecting the two vendor elasticities reported above. As a back-of-the-envelope illustration only—no option trades on the partisan spread—a 4 to 11% rise in volatility maps to roughly a 4 to 11% richer at-the-money straddle premium[10] on a hypothetical claim on the spread, since premia scale approximately linearly in volatility over short horizons. For external context, the S&P 500 options market itself priced a 2.4% move over election week (a four-day at-the-money straddle; DataBento); this is a different object, but it shows the period's election-driven option-implied uncertainty.

[10] An at-the-money straddle is a long call and a long put struck at the current price; its premium—the combined cost of the two options—rises with expected volatility, which is why it is the standard instrument for pricing or trading volatility.

6 Robustness

Table 8 consolidates the robustness checks. The results survive changes in how disagreement and volatility are measured. Replacing the cross-sectional standard deviation with the range, interquartile range, coefficient of variation, or Gini coefficient does not alter the qualitative finding. Nor does replacing realized volatility with realized variance or mean absolute returns.

The conclusions hold at the daily frequency. Disagreement predicts partisan-spread volatility and remains informative after controlling for market-wide volatility. The results are similar whether the daily implied-volatility control is VIXY or the standard CBOE VIX close.

In subsamples, disagreement positively predicts partisan-spread volatility both before and after the assassination attempt of July 13, though the coefficient is larger in the first half of the sample when cross-platform variation in disagreement was roughly twice as high.

The estimates are not driven by any single platform. Leave-one-out exercises preserve the main result. A block bootstrap delivers confidence intervals that exclude zero, and a time-series permutation test that randomly shuffles disagreement across hours yields $p < 0.001$ in both uncontrolled and controlled specifications. Newey–West inference is stable across bandwidth choices. We find no strong asymmetry depending on whether Trump is leading or trailing, and the relationship holds across different intraday windows. A nine-market specification that adds Kalshi and Pinnacle is considerably noisier, which is consistent with its much smaller usable sample.

Replacing contemporaneous disagreement with its one-hour lag yields nearly identical point estimates, confirming that the relationship is predictive rather than merely contemporaneous. Including a lagged dependent variable to absorb volatility persistence reduces the coefficient modestly ($\hat{\beta} = 10.07$, $t = 5.99$), but it remains highly significant. Hour-of-day and day-of-week fixed effects produce similar results ($\hat{\beta} \approx 13$, $t > 8$), ruling out intraday seasonality. Two-way clustering by date and hour of day yields a t-statistic of 8.15, larger than the HAC baseline. A kitchen-sink specification that includes all controls, both sets of fixed effects, and a lagged dependent variable simultaneously renders the coefficient insignificant. We interpret this as an over-controlled specification with severe multicollinearity: the hourly autocorrelation of disagreement exceeds 0.97, so contemporaneous and lagged disagreement are nearly indistinguishable. Variance-inflation factors confirm this: contemporaneous and lagged disagreement have factors near 24, while every other regressor is below two. The insignificance is therefore specific to entering disagreement together with its near-identical lag; partialling the *standard controls* (not the lag) out of both sides leaves the disagreement coefficient unchanged ($\hat{\beta} = 12.5$, $t = 6.8$), and each timing enters significantly on its own in the specifications above.

The result does not hinge on the firm classification, and the sign and significance of the baseline coefficient do not hinge on the data vendor, though its magnitude does. A stricter tercile split that forms the spread from the most Republican and most Democratic thirds of firms yields a slightly larger coefficient ($\hat{\beta} = 16.4$, $t = 8.2$), and dropping the two firms whose classification requires judgment leaves it essentially unchanged ($\hat{\beta} = 15.8$, $t = 7.8$; Online Supplement). Panel I collects the vendor comparison discussed in Section 5. Re-estimating the baseline on independent, institutional-grade intraday prices from DataBento preserves a positive and significant relation, uncontrolled ($\hat{\beta} = 8.06$, $t = 3.77$) and controlled ($\hat{\beta} = 4.19$, $t = 2.53$), but at roughly a third of the baseline magnitude—about a 4% rather than an 11% one-standard-deviation response. Because the two vendors populate the 5-minute grid at very different rates (about 95% of firm-bars for DataBento against about 70% for Twelve Data), and because hourly realized volatility built from twelve 5-minute returns is sensitive to how many of those returns are genuine rather than forward-filled, we report the 4–11% range rather than defend a single elasticity.

Taken together, the lagged-disagreement, fixed-effects, alternative-measure, leave-one-out, and permutation specifications indicate that the core finding is stable across timing assumptions, control sets,

and measurement choices.

Table 8: Robustness Summary

Specification	$\hat{\beta}$	t	p	N	R^2
Panel A: Subsamples					
Pre-assassination (Jan–Jul 12)	19.899***	5.91	0.000	834	0.045
Post-assassination (Jul 13+)	7.455**	2.56	0.010	567	0.008
Pre-Biden (Jan–Jul 20)	18.693***	5.72	0.000	869	0.040
Post-Biden (Jul 22+)	6.001	1.28	0.202	532	0.003
Early campaign (Jan–Jun)	17.486***	4.82	0.000	771	0.032
Late campaign (Sep–Nov)	6.242	0.64	0.523	322	0.001
Panel B: Alternative disagreement measures					
Range	5.872***	7.53	0.000	1401	0.047
IQR	10.937***	7.32	0.000	1401	0.055
Coeff. of variation	10.050***	10.07	0.000	1401	0.082
Gini	22.222***	10.79	0.000	1401	0.086
Log disagreement	0.408***	8.32	0.000	1401	0.059
Panel C: Alternative volatility measures					
ln(SD) [baseline]	17.990***	8.93	0.000	1401	0.065
ln(RV)	36.460***	9.61	0.000	1401	0.085
ln(|ret|)	17.366***	9.06	0.000	1401	0.065
Panel D: Daily frequency					
SPREAD vol	22.816***	5.97	0.000	201	0.237
Controlled (VIXY)	16.158***	5.22	0.000	201	0.468
Controlled (VIXCLS)	12.794***	3.26	0.001	201	0.412
Panel E: Predictive specifications					
Lagged disagree	18.071***	9.04	0.000	1400	0.066
Lagged + controls	12.041***	6.64	0.000	1400	0.450
Lagged + controls + LDV	10.072***	5.99	0.000	1400	0.468
Panel F: Fixed effects & clustering					
Hour-of-day FE + controls	13.250***	8.53	0.000	1401	0.681
Day-of-week FE + controls	12.343***	6.93	0.000	1401	0.457
Both FE + controls	13.352***	8.70	0.000	1401	0.684
Kitchen sink (all + LDV + FE)	5.533	1.08	0.281	1400	0.704
Two-way clustered SEs	12.107***	8.15	0.000	1401	0.452
Panel G: Inference & sensitivity					
Block bootstrap	18.528***	6.48	0.000	1401	–
Asymmetry: interaction	−6.223	−1.53	0.125	1401	0.071
LOO: drop PolyMarket	16.185***	9.14	0.000	1402	0.066
LOO: drop Betfair	12.563***	4.47	0.000	1402	0.024
LOO: drop Bovada	14.985***	8.07	0.000	1402	0.057
LOO: drop BetOnline	16.780***	9.31	0.000	1402	0.075
LOO: drop William Hill	16.048***	7.74	0.000	1402	0.053
LOO: drop Unibet	15.440***	7.33	0.000	1402	0.052
LOO: drop Everygame	17.640***	8.76	0.000	1402	0.064
9-market spec	−17.220	−1.38	0.166	161	0.013
Panel H: Trading hours					
Full (9:30–16:00)	17.859***	8.79	0.000	1402	0.063
Excl. open/close	18.632***	8.50	0.000	1205	0.101
Morning (9:30–12:30)	20.113***	6.97	0.000	804	0.074
Afternoon (12:30–16:00)	15.173***	7.49	0.000	603	0.138
Panel I: Equity price vendor (controlled unless noted)					
Twelve Data, 71 firms [baseline]	12.107***	6.69	0.000	1401	0.452
Twelve Data, 69-firm universe	11.509***	6.34	0.000	1401	0.450
DataBento, 69 firms, uncontrolled	8.064***	3.77	0.000	1392	0.011
DataBento, 69 firms	4.192**	2.53	0.011	1392	0.522

Notes: Each row reports the coefficient on the disagreement measure in a regression of ln(SPREAD vol) on disagreement. HAC standard errors unless otherwise noted. "Predictive" rows use lagged disagreement; "FE" rows include fixed effects with full controls (MKT vol + VIXY); "Kitchen sink" adds all controls, both FE sets, and a lagged dependent variable; "Two-way clustered" uses Cameron–Gelbach–Miller (2011) SEs clustered by date and hour of day. Panel I compares equity price vendors: the 69-firm universe is the subset of the 71 sample firms DataBento covers. In the controlled rows ln(MKT vol) is constructed from the same vendor as the dependent variable, while ln(VIXY) is from Twelve Data throughout. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

7 Conclusion

We show that disagreement across prediction-market venues contains incremental, selective information about the volatility of politically exposed equities. A formal selectivity test confirms that the relationship is significantly stronger for a partisan long–short portfolio than for the broad market or a clean non-political benchmark, and the channel operates predominantly through second moments: disagreement predicts

absolute and squared returns far more than signed returns. The incremental effect—a 4% to 11% increase in spread volatility per one-standard-deviation rise in disagreement, after absorbing broad market conditions, with the range spanning our two intraday price vendors—persists for roughly one trading day and survives intraday fixed effects and a lagged dependent variable.

Three caveats bound the claim. First, the evidence is predictive, not causal: temporal ordering, a recursive VAR, and event-day controls point the same way, but we cannot rule out a contemporaneous common shock, and the precise magnitude is sensitive to the price source—a factor of nearly three separates our two vendor estimates, which is why we report a range rather than a point elasticity. Nor can we offer a credible instrument: scheduled shocks such as debates move equity volatility directly and so violate the exclusion restriction, while the cleaner candidate—a platform-specific outage that shifts measured disagreement for reasons unrelated to fundamentals—did not occur during our sample. We flag that design for future data. Second, the event-study evidence (Table 7) is strongest around the two largest shocks—the assassination attempt and President Biden's withdrawal—and weaker around the two debates, which suggests that the relationship may depend on the type of political shock: unscheduled events that abruptly widen the range of electoral scenarios, rather than scheduled ones that markets have already had time to price. Third, the evidence comes from a single, highly salient election. Comparable cross-platform data are not available for earlier U.S. elections—PolyMarket did not yet exist and Betfair was effectively the only liquid exchange in 2016—so cross-platform dispersion could not have been measured then. We therefore read the results as selective predictive content specific to the 2024 episode rather than as a universal pattern, and cross-venue dispersion appears to contain real-time, high-frequency information about election-related uncertainty in this setting. Whether the relationship extends to other high-stakes political events where multiple liquid prediction markets coexist remains an open question for future work.

Data availability

Prediction-market odds were obtained via a paid subscription to The Odds API, with PolyMarket data from the public Polygon blockchain. Baseline intraday equity and VIXY prices are from Twelve Data; daily VIX is from FRED and daily equity data from Compustat. Firm partisan classifications are based on PAC contribution records compiled by OpenSecrets. The non-political benchmark portfolio, the data-vendor robustness check, and the S&P 500 option straddle use institutional-grade intraday data licensed from DataBento. Replication code, the redistributable data, and full instructions for obtaining the licensed sources are deposited at Harvard Dataverse (Aktuğ and Torul, 2026c).

A Appendix

Table A.1: Summary Statistics

Variable	*N*	Mean	SD	Min	Max
Panel A: Hourly					
Disagreement (SD)	1402	0.0254	0.0093	0.0035	0.0477
Disagreement (range)	1402	0.0658	0.0243	0.0080	0.1328
SPREAD vol (hourly)	1402	0.0581	0.0592	0.0103	0.7439
MKT vol (hourly)	1402	0.0634	0.0426	0.0111	0.3690
Mean probability	1402	0.5448	0.0576	0.4717	0.7467
Panel B: Daily					
Disagreement (SD)	201	0.0256	0.0091	0.0102	0.0475
SPREAD vol (daily)	201	0.0642	0.0353	0.0280	0.2999
MKT vol (daily)	201	0.0675	0.0220	0.0332	0.1926

Notes: Hourly realized volatility is the within-hour standard deviation of 5-minute log returns. Disagreement is the cross-sectional standard deviation of probability levels across 7 platforms. Sample: January 22–November 5, 2024 (U.S. trading hours).

References

Addoum, J. M. and Kumar, A. (2016). Political sentiment and predictable returns. *Review of Financial Studies*, 29(12):3471–3518.

Aktuğ, E. and Torul, O. (2026a). Political capital: Real-time electoral probabilities and the value of corporate political alignment. Working Paper, available at https://doi.org/10.2139/ssrn.6624604.

Aktuğ, E. and Torul, O. (2026b). Price discovery across political prediction markets. Working Paper, available at https://doi.org/10.2139/ssrn.6516491.

Aktuğ, E. and Torul, O. (2026c). Replication Data for: When Prediction Markets Disagree: Cross-Platform Dispersion and Equity Volatility in the 2024 U.S. Election. Harvard Dataverse, https://doi.org/10.7910/DVN/O2W2JR.

Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., Levmore, S., Litan, R., Milgrom, P., Nelson, F. D., Neumann, G. R., Ottaviani, M., Schelling, T. C., Shiller, R. J., Smith, V. L., Snowberg, E., Sunstein, C. R., Tetlock, P. C., Tetlock, P. E., Varian, H. R., Wolfers, J., and Zitzewitz, E. (2008). The promise of prediction markets. *Science*, 320(5878):877–878.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636.

Banerjee, S. and Kremer, I. (2010). Disagreement and learning: Dynamic patterns of trade. *Journal of Finance*, 65(4):1269–1302.

Bossaerts, F., Yadav, N., Bossaerts, P., Nash, C., Todd, T., Rudolf, T., Hutchins, R., Ponsonby, A.-L., and Mattingly, K. (2024). Price formation in field prediction markets: The wisdom in the crowd. *Journal of Financial Markets*, 68:100881.

Chakravarty, S., Gulen, H., and Mayhew, S. (2004). Informed trading in stock and option markets. *The Journal of Finance*, 59(3):1235–1257.

Diether, K. B., Malloy, C. J., and Scherbina, A. (2002). Differences of opinion and the cross section of stock returns. *Journal of Finance*, 57(5):2113–2141.

Gonzalo, J. and Granger, C. W. (1995). Estimation of common long-memory components in cointegrated systems. *Journal of Business & Economic Statistics*, 13(1):27–35.

Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *Journal of Finance*, 50(4):1175–1199.

Hong, H. and Stein, J. C. (2007). Disagreement and the stock market. *Journal of Economic Perspectives*, 21(2):109–128.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Kelly, B., Pástor, Ľ., and Veronesi, P. (2016). The price of political uncertainty: Theory and evidence from the option market. *Journal of Finance*, 71(5):2417–2480.

Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *Journal of Finance*, 32(4):1151–1168.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.

Pástor, Ľ. and Veronesi, P. (2012). Uncertainty about government policy and stock prices. *Journal of Finance*, 67(4):1219–1264.

Pástor, Ľ. and Veronesi, P. (2013). Political uncertainty and risk premia. *Journal of Financial Economics*, 110(3):520–545.

Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126.