ELECTION PROBABILITIES AND THE TREASURY YIELD CURVE*

EMREHAN AKTUĞ† *Sabancı University*

ORHAN TORUL‡ *Boğaziçi University*

May 2026

Abstract

During the 2024 U.S. presidential campaign, daily increases in Polymarket's Trump win probability were associated with a long-end Treasury steepening, term-premium repricing, and a coherent cross-asset Trump-trade signature. A ten-percentage-point probability rise corresponds to a ∼2 basis-point steepening of the 2s10s slope; the short end and the overnight policy rate are statistically unresponsive; an Adrian et al. (2013) decomposition assigns 95% of the ten-year response to the term premium. Identification combines a non-trading-hours decomposition (the yield response loads entirely on the overnight component of Δp_t, when cash markets are closed), a 2SLS estimator using the overnight component as an instrument (first-stage $F = 37$, IV exceeds OLS), a CME ten-year futures overnight gap on 1-minute Globex prices that is null ex-Election-Day, and Betfair replication. A formal Election Day holdout test (estimate through November 5, predict November 6) puts the predicted slope move within 1 bp of the realized 8 bp and the predicted 10-year real-yield change within rounding of the realized 5 bp; the breakeven decomposition share remains leverage-dependent. A leave-one-event-out cross-validation across six politically salient 2024 dates yields an out-of-sample slope R^2 of 0.34 and a predicted–realized correlation of 0.74. The 2020 Biden–Trump cycle (PredictIt, $N = 458$) is sign-consistent at every long-end maturity; a 2019 pre-period placebo using the same contract is statistically zero across all maturities. We treat sign and channel composition as the load-bearing inference.

Keywords: Treasury yields; election uncertainty; prediction markets; breakeven inflation
JEL classification: E43; G12; G18

1 Introduction

When Polymarket's daily Trump win probability rose during the 2024 U.S. presidential campaign, the long end of the Treasury curve rose with it, the dollar appreciated, financial and energy equities outperformed, and utilities and gold underperformed—a coherent cross-asset signature consistent with a fiscal-expansion-with-higher-rates regime. A ten-percentage-point rise in the daily Trump probability is associated with a two basis-point steepening of the 2s10s slope; the short end and the overnight policy rate do not move; an Adrian et al. (2013) affine decomposition assigns 95% of the long-end response to the term premium and a Kim and Wright (2005) cross-check confirms that the near-term Fed path is unresponsive.

Whether political risk is priced in U.S. Treasury markets—and through what channel—is a long-standing question. Santa-Clara and Valkanov (2003) document a presidential cycle in equity returns; Pástor and

*We thank our colleagues at Boğaziçi and Sabancı University for helpful comments. Torul gratefully acknowledges financial assistance by the Faculty of Economics and Administrative Sciences at Boğaziçi University for the submission of this manuscript. All errors are our own.

†Address: Sabancı University Department of Economics, 34956, Tuzla, Istanbul, Türkiye.
E-Mail: emrehan.aktug@sabanciuniv.edu

‡Corresponding author. Address: Boğaziçi University, Department of Economics, 34342, Bebek, Istanbul, Türkiye.
E-Mail: orhan.torul@bogazici.edu.tr

Veronesi (2012, 2013) formalize a political-uncertainty premium that rises when the policy change is economically large; Kelly et al. (2016) find this premium in equity options around national elections; Baker et al. (2016) document broad asset-price sensitivity to economic-policy uncertainty; and Boutchkova et al. (2012) show that politically sensitive industries earn higher return volatility when election-driven policy risk is salient. Bianchi et al. (2023) show that fiscal regimes with imperfect monetary accommodation generate persistent inflation in the long-run discount factor. We bring high-frequency, contract-traded probability of a *specific* presidential outcome (Wolfers and Zitzewitz, 2004) to this question and decompose the long-end response into expected-rate and term-premium components using Adrian et al. (2013) and Kim and Wright (2005).

The reduced-form relation is conditional rather than uniform. Two within-sample splits matter for interpretation. First, the response loads almost entirely on Trump-gain days: $\hat{\beta}_{+,\text{slope}} = 21.4$ ($t = 5.7$) versus $\hat{\beta}_{-,\text{slope}} = 9.0$ ($t = 0.4$). Second, the breakeven channel operates specifically when trade-policy uncertainty is salient: a Caldara et al. (2020) daily-TPU quintile decomposition delivers a 10-year breakeven coefficient of 22.5 ($t = 4.7$, $N = 91$) in the top quintile and 8.3 ($t = 0.6$, $N = 91$) in the bottom. The headline "two basis points per ten percentage points" is therefore an average across days that are, in mechanism terms, not interchangeable.

Identification combines three layers. (i) A non-trading-hours decomposition: the yield response loads entirely on the prior-15:00-ET-to-08:00-ET window of Δp_t, in which cash Treasury yields cannot have moved, ruling out same-day reverse causality. (ii) Two-stage least squares using that overnight component as an instrument for daily Δp_t (first-stage $F = 37$, coefficient 0.88 with $t = 6.1$): the IV slope coefficient is 23.1 ($t = 1.9$) and the IV 10-year breakeven coefficient is 24.6 ($t = 3.4$), each larger than OLS, so attenuation from intraday noise rather than reverse causation is the more plausible bias direction. (iii) The CME 10-year futures overnight gap on 1-minute Globex prices over the same window is statistically zero ex-Election-Day ($\hat{\beta} = 1.2$, $t = 0.1$, $N = 457$), so the cash response is not a mechanical futures arbitrage on the typical-day sample; on Election Day specifically, ZN priced 12.5 of the eventual 16 bp 10-year move before cash opened (Section 3.3). Results also replicate using the betting odds by Betfair (structurally independent, UK-regulated). The geographic comparison is mixed: the German Bund response is null while the U.K. Gilt response is partially significant, indicating a predominantly U.S. channel with limited duration spillover rather than a clean U.S.-specific identification.

A formal Election Day holdout test sharpens this reading. Estimating the baseline through November 5, 2024 and using it to predict the November 6 close-to-close moves out-of-sample, the predicted slope steepening is 9.0 bp against a realized 8.0 bp (|err| = 1.0 bp) and the predicted 10-year real-yield change is 5.0 bp against a realized 5.0 bp—each within rounding of zero; the 10-year breakeven and 30-year nominal moves *exceed* their no-leverage predictions by 7–8 bp, so the leverage point sharpens rather than weakens the slope and real-yield identification. A leave-one-event-out cross-validation across all six politically salient 2024 dates yields an out-of-sample R^2 of 0.34 for the slope and correct predicted sign on 10 of 10 (event × outcome) pairs at the two largest probability shocks. Recursive leave-one-month-out keeps the slope inside [16.4, 23.5]. We treat sign and channel composition as the load-bearing inference, not the daily-frequency point estimate, and do not claim a structural parameter recoverable from a single cycle.

Relative to prior work on political uncertainty in sovereign debt markets (Kelly et al., 2016; Bianchi et al., 2023), this paper contributes (i) a multi-layered daily-frequency identification combining 5-minute Polymarket data with cross-asset, cross-platform, and futures-market diagnostics on the 2024 cycle; (ii) a holdout-prediction test that turns the largest leverage observation into a near-exact out-of-sample slope and real-yield prediction; and (iii) sign-consistent replication on the 2020 PredictIt Biden–Trump cycle

plus a 2019 pre-period placebo on the same contract that is statistically zero across all maturities, locating the channel in a U.S. fiscal-inflation regime reading rather than a generic political-uncertainty premium, a venue-specific Polymarket microstructure artifact, or a 2024-cycle phenomenon.

Table 1: Summary Statistics

Variable	N	Mean	SD	Min	Max
Level					
Trump win probability	459	0.501	0.066	0.270	0.995
Daily changes					
Δ Trump probability	457	0.001	0.024	−0.080	0.420
Δ 2Y yield (bp)	458	0.01	8.21	−57.0	25.0
Δ 10Y yield (bp)	458	0.18	6.71	−23.0	19.0
Δ 30Y yield (bp)	458	0.19	5.93	−18.0	17.0
Δ 2s10s slope (bp)	458	0.17	4.93	−16.0	42.0
Δ 10Y real yield (bp)	458	0.15	5.63	−19.0	15.0
Δ 10Y breakeven (bp)	458	0.03	2.88	−16.0	13.0
Δ 5Y5Y fwd breakeven (bp)	458	0.00	3.21	−15.0	18.0

Notes: Sample covers January 10, 2023 to November 6, 2024. The Trump probability is the Polymarket daily close on the unit interval; daily changes are close-to-close first differences. Δ Trump probability has $N = 457$ rather than 458 because one thin-market artifact (January 24, 2023; Cook's $D > 1$, leverage 0.56) is set to missing, as described in Section 2.

2 Data and Specification

The sample comprises 459 trading days from January 10, 2023 to November 6, 2024. Daily nominal Treasury yields and TIPS real yields are from the U.S. Treasury's par yield-curve releases; breakeven inflation is the nominal–real spread. Trump win probabilities are sampled at 15:00 ET each trading day from Polymarket's 2024 presidential-election contract, aligned with the Treasury par-yield release.[1] Betfair odds for the same contract cover 210 matched trading days; six additional bookmakers cover between 83 and 210 days. The baseline specification is

$$\Delta y_{m,t} = \alpha + \beta \Delta p_t + \gamma' X_t + \varepsilon_t, \tag{1}$$

with Newey–West HAC standard errors (5 lags, $\lfloor T^{1/4} \rfloor \approx 4.6$). Throughout we use specification (2) of Table 2 (lagged-odds controls, $N = 453$) as the canonical baseline for headline coefficients reported in the abstract and main text ($\hat{\beta}_{\text{slope}} = 19.4$; controlling-for-MOVE 10-year breakeven $\hat{\beta} = 20.5$, $t = 4.4$, footnote 2 below). For the real–breakeven decomposition we report Table 3 under specification (1) (no controls) so that the breakeven and real-yield shares sum mechanically to the nominal coefficient; the breakeven share is materially unchanged under specification (2). A coefficient of $\beta \approx 20$ corresponds to a two basis-point response per ten-percentage-point probability shock. Table 1 reports summary statistics; Figure 1 plots the key series.

[1] We set Δp_t to missing on January 24, 2023, a stale-price-correction event with no contemporaneous newsfeed (the contract jumped from 0.27 to 0.49 over two prints after ~36 hours flat); baseline coefficients are essentially unchanged if the observation is Huber-downweighted instead.

Figure 1: Trump Election Probability and Treasury Variables

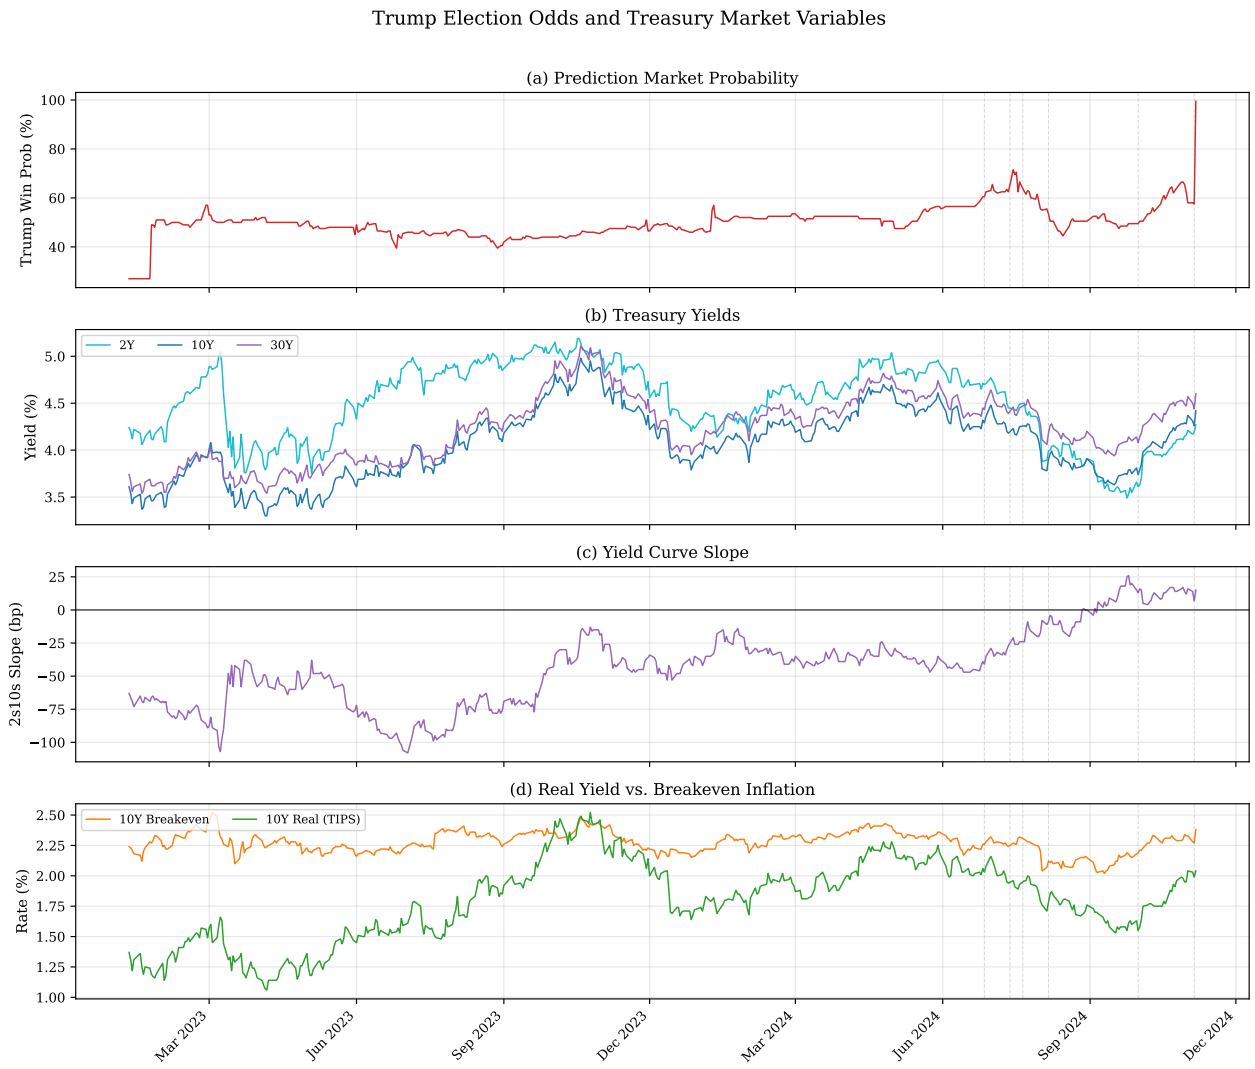

Notes: Panel (a) plots the Polymarket daily closing Trump win probability; panels (b)–(d) plot nominal yields at the 2-, 10-, and 30-year maturities, the 2s10s slope, and the 10-year breakeven alongside the 10-year TIPS yield. Vertical dashed lines mark, in chronological order, the Biden debate (June 27, 2024), the Trump assassination attempt (July 13, 2024), the Biden withdrawal (July 21, 2024), the Harris polling surge (early August 2024), the October Polymarket swing toward Trump (October 1, 2024), and Election Day (November 5–6, 2024).

Table 2: Yield Response to Trump Win Probability by Maturity

	Specification					Excl.
Outcome	(1)	(2)	(3)	(4)	(5)	ElecDay
Panel A: Yield levels						
2Y nominal	11.0	12.8	12.9	0.0	−3.7	−0.7
	(8.8)	(7.9)	(7.8)	(9.3)	(9.6)	(21.5)
5Y nominal	20.0***	20.9***	20.9***	10.1	6.1	11.2
	(7.2)	(6.7)	(6.7)	(8.1)	(8.7)	(19.3)
10Y nominal	31.3***	32.2***	31.5***	25.4***	22.4***	20.6
	(6.9)	(6.8)	(6.5)	(7.6)	(8.2)	(18.1)
30Y nominal	30.9***	31.5***	30.9***	28.0***	26.1***	18.6
	(6.7)	(6.6)	(6.3)	(7.7)	(8.2)	(16.7)
Panel B: Curve slopes						
2s10s slope	20.3***	19.4***	18.6***	25.3***	26.1***	21.3
	(4.9)	(4.4)	(4.5)	(6.1)	(6.0)	(13.2)
Lagged odds	No	Yes	Yes	Yes	Yes	Yes
Lagged yields	No	No	Yes	Yes	Yes	No
VIX, equities	No	No	No	Yes	Yes	No
FOMC, Fed cycle	No	No	No	No	Yes	No
N	457	453	453	355	355	452

Notes: Dependent variable is the daily change in the yield or slope (basis points). Δp_t is measured on the unit interval. Newey–West HAC standard errors with five lags in parentheses. Specifications (4)–(5) have fewer observations because the VIX and S&P 500 control series begin partway into the panel; the 102 missing days are concentrated at the start of the sample (calendar, not selection). The final column applies specification (2) controls (lagged odds) on the sample with November 6, 2024 (Election Day) excluded; the point estimates remain positive but are no longer statistically significant at conventional levels under simple OLS exclusion of the leverage point. Election Day is a 17-standard-deviation realization of Δp_t and the dominant source of identifying variation; the coefficient stability under Huber downweighting (Section 4) and the asymmetry decomposition (Table 5) provide complementary leverage diagnostics. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

3 Results

3.1 Maturity Gradient

Table 2 reports the cross-section of yield responses; Figure 2 summarizes the pattern. The two-year coefficient is positive but statistically indistinguishable from zero, consistent with a short end anchored by the near-term Fed path. The five-, ten-, and thirty-year coefficients are positive and significant, increasing through the ten-year point ($\hat{\beta}_{5Y} = 20.0$, $\hat{\beta}_{10Y} = 31.3$) and remaining essentially flat between ten and thirty years ($\hat{\beta}_{30Y} = 30.9$). The two-year/ten-year slope coefficient is $\hat{\beta} = 20.3$ ($t = 4.2$) without controls (specification 1) and $\hat{\beta} = 19.4$ ($t = 4.4$) in the canonical baseline (specification 2 with lagged-odds controls); it remains between 18.6 and 26.1 as further controls are added. In economic terms, a one-standard-deviation Δp_t shock (2.4 percentage points) corresponds to roughly a 0.5 basis-point slope move (10% of the slope's daily standard deviation); the policy-relevant ten-percentage-point benchmark, i.e., about a four-standard-deviation event for Δp_t, corresponds to a 1.94 basis-point move, roughly 40% of the unconditional daily standard deviation of the 2s10s slope (4.9 bp) and nearly half the mean absolute FOMC-day slope move (4.0 bp).

3.2 Breakeven versus Real Yields

Table 3 decomposes the 10-year nominal response. Breakeven inflation accounts for 62% ($\hat{\beta}_{\text{BE}} = 19.4$, $t = 3.6$); real yields for the remainder ($\hat{\beta}_{\text{real}} = 11.9$, $t = 2.6$); the 5y5y forward breakeven rises 13.9 bp

Figure 2: Cross-Section of Yield Responses to Trump Win Probability

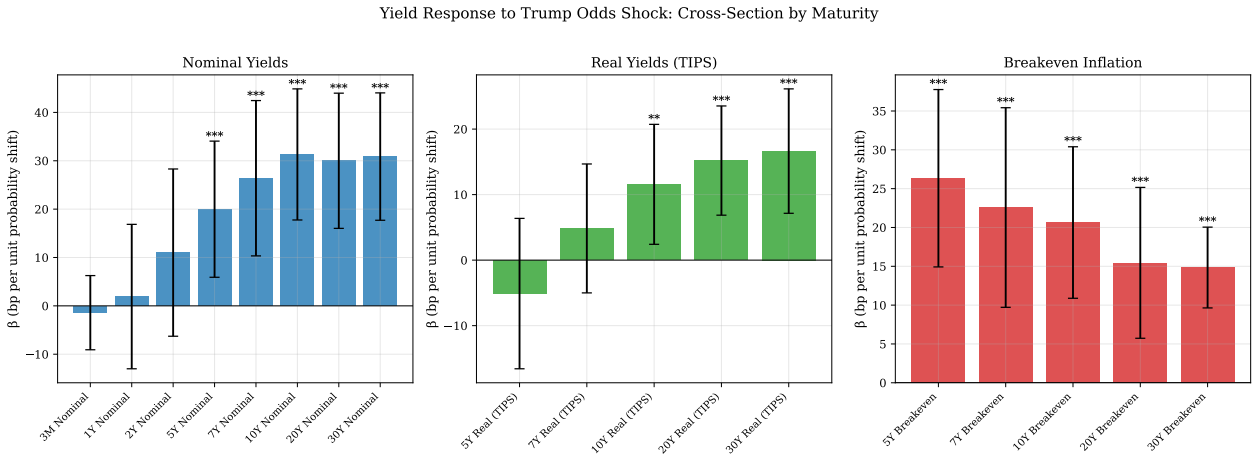

Notes: Each bar reports the no-controls specification (1) coefficient on the unit-interval probability scale, with 95% Newey–West confidence intervals. The left panel shows nominal yields broadly rising with maturity. The middle panel shows real (TIPS) yields are significant at ten and thirty years but not at five years (where the response is small and negative). The right panel shows breakeven inflation is significant at all maturities; the five-year coefficient is the largest in absolute magnitude because the five-year real-yield offset is negative, while the ten-year breakeven is the central reduced-form mechanism in the text.

Table 3: Real–Breakeven Decomposition: Full Sample and Excluding Election Day

	Full sample ($N = 457$)				Excl. Election Day ($N = 456$)			
Outcome	Nominal $\hat{\beta}$	Real $\hat{\beta}$	Breakeven $\hat{\beta}$	BE share	Nominal $\hat{\beta}$	Real $\hat{\beta}$	Breakeven $\hat{\beta}$	BE share
5Y	20.0***	−4.9	24.8***	124%	8.0	0.1	7.9	99%
	(7.2)	(6.1)	(6.4)		(19.5)	(16.2)	(9.0)	
10Y	31.3***	11.9**	19.4***	62%	18.4	12.6	5.8	32%
	(6.9)	(4.6)	(5.4)		(17.9)	(13.8)	(8.0)	
30Y	30.9***	17.0***	13.8***	45%	17.1	9.0	8.1	47%
	(6.7)	(4.7)	(2.9)		(16.6)	(12.9)	(6.6)	
5Y5Y forward	—	—	13.9***	—	—	—	3.7	—
			(5.4)				(11.6)	

Notes: Coefficients from Equation (1) with no controls (specification 1 of Table 2). BE share equals $\hat{\beta}_{\text{BE}}/\hat{\beta}_{\text{nominal}}$; values above 100% arise when the real-yield response has the opposite sign. Newey–West HAC standard errors with five lags. The right-hand panel drops Election Day (November 6, 2024) from the sample. Excluding that observation, none of the decomposition coefficients survives at conventional levels: the precise 10-year breakeven share (62%) is identified primarily off Election Day. The breakeven *sign* remains positive at every maturity, so the channel composition—rather than the magnitude—is the load-bearing inference. The ACM term-premium attribution (Section 3.5) does not rely on this decomposition. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

per unit shock.[2] The real-yield response is less robust (sign reverses at 5Y; permutation $p = 0.27$; Huber $t = 1.1$). Excluding Election Day, the breakeven coefficient falls to 8.8 ($t = 1.0$) while the real coefficient is unchanged in magnitude (11.8) but loses precision ($t = 0.8$); the precise decomposition share rests on Election Day, but the ACM term-premium attribution (Section 3.5) provides the more reliable structural read.[3]

3.3 Event Corroboration

Election Day (the close-to-close window November 5 to November 6, 2024) is the dominant observation: a 42-percentage-point probability shock coincided with an 8 bp steepening of the 2s10s slope (linear prediction $20.3 \times 0.42 \approx 8.5$ bp), a 16 bp rise in the 10-year yield, and an 11 bp rise in the 10-year breakeven. Of the 16 bp 10-year cash rise, 12.5 bp was already priced in CME ZN front-month futures by 08:00 ET on November 6—a 78% pre-cash share computed from 1-minute Globex prices over the prior 15:00 ET to current 08:00 ET window, providing event-level corroboration that political-probability information was priced into Treasury duration during cash-closed hours. A Nov-5 cutoff holdout test—estimate the baseline through November 5, predict November 6 out-of-sample—places the predicted slope move at 9.0 bp against a realized 8.0 bp (i.e., an absolute error of 1.0 bp) and the predicted 10-year real-yield change at 5.0 bp against a realized 5.0 bp. The 10-year breakeven and 30-year nominal moves *exceed* their no-leverage predictions by 7–8 bp, so Election Day is a higher-magnitude realization than the rest of the sample suggests; the leverage point sharpens rather than weakens the slope and real-yield identification. A leave-one-event-out cross-validation across the six politically salient 2024 dates (Biden debate, the Saturday-attack close-to-close, Biden withdrawal, Harris surge, the October Polymarket swing, and Election Day) yields an out-of-sample R^2 of 0.34 for the 2s10s slope (predicted–realized correlation 0.74, RMSE 3.1 bp), 0.29 for the 10-year nominal, and 0.17 for the 10-year breakeven (correlation 0.82, indicating the model captures the sign and ranking of breakeven event-day moves but understates magnitude). Restricting to the four events with $|\Delta p_t| > 1$ pp, the model produces correct predicted sign on 13 of 20 (event × outcome) pairs, with 10 of 10 correct on the two largest probability shocks. Election-Day-excluded estimates appear in Table 2 (column 6); leverage diagnostics are in Section 4. A non-overlapping six-event Ibragimov–Müller cluster test gives $\bar{\beta} = +24.0$ (one-sided $p = 0.085$, sign test 5/6); we read it as directionally consistent given its low power at $K = 6$ rather than as independently confirming the daily result, and report it in Appendix Table A.6. Figure 3 plots the $(\Delta p_t, \Delta\text{slope})$ scatter with the six salient 2024 dates highlighted; the regression line passes nearly through the Election Day observation, consistent with the holdout result.

3.4 Cross-Platform Replication

A natural concern is that the 2024 Polymarket signal could reflect venue-specific microstructure—late-October concentrated buying by a small set of large accounts, thin daytime liquidity outside the U.S., or U.S.-resident-access restrictions that bias the price level—rather than aggregate political-probability information. Three features of the design speak to this directly. First, the daily-frequency relation replicates

[2] The 5-year breakeven share exceeds 100% (124%) because the 5-year real-yield coefficient is small and negative ($\hat{\beta}_{\text{real,5Y}} = -4.9$, statistically zero), so the breakeven channel slightly over-shoots the muted nominal response at that maturity; the share is mechanical rather than substantive and is not load-bearing for the term-premium narrative.

[3] Controlling for the MOVE index (Treasury options-implied volatility) leaves the Trump coefficient at 20.5 ($t = 4.4$); a TIPS-specific check using the panel-vs-Gürkaynak et al. (2007)-fitted breakeven residual leaves it at 20.0 ($t = 4.0$). The breakeven response is not driven by Treasury or TIPS-specific liquidity stress.

Figure 3: Daily $(\Delta p_t, \Delta\text{slope}_{2s10s})$ Scatter with Six Salient 2024 Events

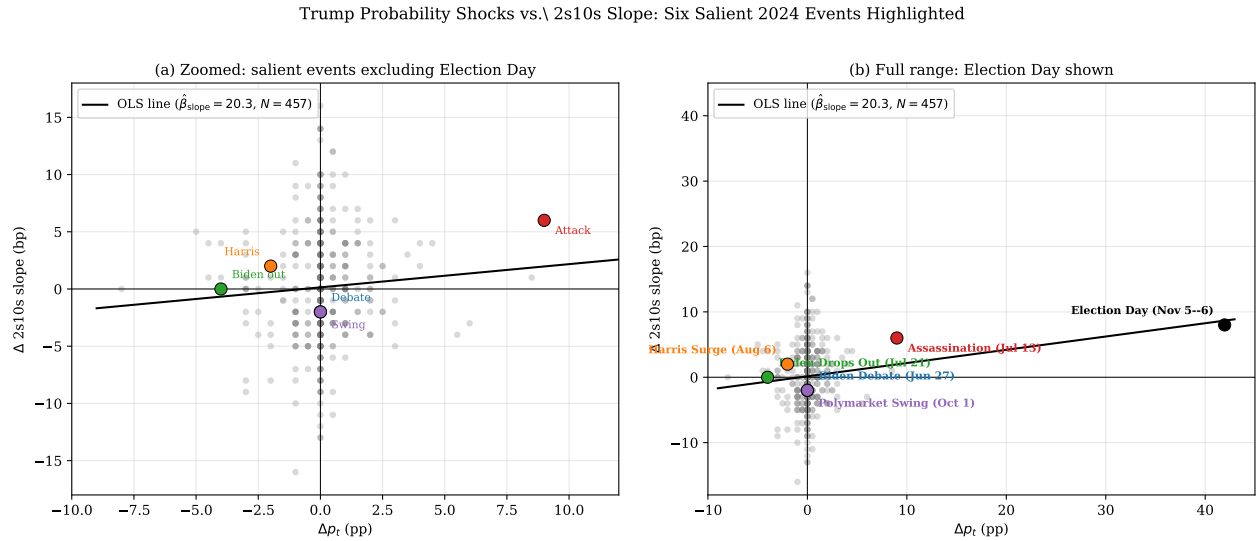

Notes: Gray dots are the daily sample ($N = 457$). The black line is the OLS fit (no controls, specification (1) of Table 2, $\hat{\beta} = 20.3$). Six politically salient 2024 dates are highlighted: the Biden debate (June 27), the Trump assassination-attempt close-to-close (July 15, the Monday after the Saturday attack), Biden's withdrawal close-to-close (July 22), the Harris polling surge (August 6), the October Polymarket swing (October 1), and Election Day (November 5–6, $\Delta p_t = 0.42$). The left panel zooms in on the central cluster ($|\Delta p_t| \leq 12$ pp); the right panel shows the full range, including Election Day. The OLS line passes through the Election Day observation almost exactly; the holdout prediction in the text formalizes this.

on Betfair, a UK-regulated peer-to-peer exchange with a different clientele, settlement rules, and regulatory regime, on the 210 matched trading days: $\hat{\beta}_{\text{Betfair}} = 47.3$ ($t = 3.7$) versus $\hat{\beta}_{\text{Polymarket}} = 19.4$ ($t = 4.5$); Polymarket and Betfair Δp are correlated at $r = 0.87$ on the matched sample. Second, the response loads on the prior-15:00-ET to 08:00-ET window (Section 4), the part of the day in which Polymarket trading is dominated by non-U.S. participants and is least exposed to a single-account thesis trade. Third, the channel replicates on the 2020 PredictIt cycle (a separate U.S.-regulated venue four years earlier; Appendix Table A.5), so the relation cannot be specific to Polymarket's 2024 trading population.[4] We treat the cross-platform sign and significance agreement, rather than any individual point estimate, as the load-bearing fact.

3.5 Expected Yields versus Term Premium

An Adrian et al. (2013) affine term-structure decomposition (ACM) separates the 10-year response into a risk-neutral expected-yield component (1.2 bp, $t = 0.3$) and a term premium (25.9 bp, $t = 3.4$); the term premium absorbs 95% of the total.[5] The unresponsive near-term Fed path is corroborated model-free: regressing daily changes in 3-month, 6-month, 1-year, and 2-year Treasury yields and SOFR on Δp_t produces coefficients of -1.7, 1.2, 3.0, 12.8, and -2.7, each statistically zero. A Kim and Wright (2005) cross-check (Appendix Table A.3) assigns 55% of the 10-year response to the term premium and 45% to the expected-yield component, with the KW expected-yield response concentrated at horizons of five years and longer (1Y and 2Y coefficients statistically zero); the long-horizon expected-rate loading is consis-

[4] Six retail/sharp sportsbooks (BetOnline, Bovada, Pinnacle, Unibet, William Hill, Everygame) revise quotes on fewer than 35% of trading days and do not independently identify the relation; full results are in Appendix Table A.7.

[5] The two components sum to 27.1 bp rather than the 31.3 bp baseline of Table 2 because the ACM zero-coupon yield series differs slightly in level and timing from the par CMT yield used in the baseline; the decomposition share is robust to using the ACM-fitted 10-year yield as the dependent variable.

tent with markets pricing higher long-run inflation pushing future Fed targets higher rather than a near-term tightening response. Both models agree on a dominant term-premium role and an unresponsive near-term path; they disagree on the magnitude of the long-horizon expected-rate component, broadly consistent with the inflation-uncertainty term-premium evidence in Wright (2011).

Table 4: Non-Trading-Hours versus Trading-Hours Decomposition of Δp_t

	Non-trading hours		Trading hours		
Outcome	$\hat{\beta}$	t	$\hat{\beta}$	t	N
Δ 2s10s slope	20.08***	4.71	5.22	0.36	458
Δ 10Y yield	30.60***	3.67	0.60	0.02	458
Δ 10Y breakeven	21.62***	3.90	2.68	0.33	458

Notes: Each row reports coefficients from a joint regression of the outcome on the non-trading-hours component of Δp_t (prior-day 15:00 ET Treasury close to current-day 08:00 ET open) and the trading-hours component (08:00–15:00 ET), constructed from Polymarket's five-minute series converted to U.S. Eastern time. The sum of the two components equals the total 15:00-to-15:00 daily Δp_t. Non-trading hours account for 58% of the average absolute daily move. Newey–West HAC standard errors with five lags. *** $p < 0.01$.

3.6 Pre-Period Placebo

A natural concern with any one-cycle finding is that the relation might reflect a contemporaneous regime correlation rather than political-probability information. Polymarket's 2024 contract did not exist before January 2023, so we use *PredictIt's 2020 Trump-win contract during 2019* (the calendar year preceding the 2020 election), which traded continuously but moved within a narrow [0.22, 0.48] probability range (no Democratic nominee, pre-COVID, pre-impeachment-resolution). On this 2019 sub-sample ($N = 238$ trading days after lag controls), regressing daily changes in the Kim and Wright (2005) 2-, 5-, and 10-year zero-coupon Treasury yields and the 2s10s slope on $\Delta p_{\text{PredictIt},t}$ produces statistically zero coefficients at every maturity: the slope coefficient is 7.2 ($t = 0.49$, 95% CI [−21.5, 35.9]); the 10-, 5-, and 2-year nominal coefficients are −14.2, −19.3, and −21.4, all $|t| < 1$. The 2024 baseline coefficients lie inside the upper bound of the 2019 95% CI; the standard deviation of $\Delta p_{\text{PredictIt}}$ is roughly 4× smaller than $\Delta p_{\text{Polymarket,2024}}$, so the 2019 magnitudes are not informative and we read the exercise as an unrejected null at low power. The load-bearing identification arguments are the non-trading-hours decomposition, the leave-one-event-out cross-validation, and the 2020 PredictIt out-of-sample replication.

4 Robustness

Identification: A non-trading-hours decomposition of Δp_t (prior 15:00 ET to current 08:00 ET, vs. 08:00–15:00 ET) produces a sharp asymmetry (Table 4): the non-trading-hours coefficient is 20.1 for the 2s10s slope ($t = 4.7$), 30.6 for the 10-year yield ($t = 3.7$), and 21.6 for the 10-year breakeven ($t = 3.9$)—each essentially identical to the daily baseline—while the trading-hours coefficient is statistically zero in every case. Since 58% of the average daily $|\Delta p_t|$ occurs when cash markets are closed, and cash yields cannot have moved during hours in which they are not traded, same-day reverse causality from cash yields is ruled out. The CME 10-year futures overnight gap on 1-minute Globex prices read at 15:00 ET (T−1) and 08:00 ET (T) is statistically zero ex-Election-Day ($\hat{\beta} = 1.2$ in yield-equivalent bp, $t = 0.1$, $N = 457$; on the full sample $\hat{\beta} = 21.7$, $t = 3.0$, $N = 458$, driven entirely by the single Election-Day observation), so the cash

response is not a mechanical echo of overnight futures price discovery on the typical-day sample; the Election-Day-specific overnight result is reported separately in Section 3.3.

We sharpen this argument by reframing it as a formal instrumental-variables estimator: we use the overnight component of Δp_t as an instrument for the daily 15:00-to-15:00 close-to-close Δp_t. The first stage is strong (overnight coefficient 0.88, $t = 6.1$, $F \approx 37$), well above conventional weak-instrument thresholds. The 2SLS coefficients are 23.1 ($t = 1.9$) for the 2s10s slope, 36.0 ($t = 2.1$) for the 10-year yield, and 24.6 ($t = 3.4$) for the 10-year breakeven—each *larger* than the corresponding OLS estimate. The exclusion restriction is the same one already imposed by the non-trading-hours decomposition: cash yields cannot move overnight, so any covariance between overnight Δp_t and the morning yield change must operate through political-probability information rather than yield-driven order flow. Two-stage estimates exceeding OLS suggest that intraday measurement noise rather than reverse causation is the dominant source of bias, which strengthens the structural reading of the OLS magnitude.

Leverage and Asymmetry: Election Day ($\Delta p_t = 0.42$, $\sim 17\sigma$) is the single largest observation. A permutation placebo on Δp_t places the actual coefficient in the top 4% of the null ($p = 0.040$ slope; $p = 0.006$ breakeven); a 10-day block-permutation gives $p = 0.060$ and $p = 0.004$. Huber M-regression yields $\hat{\beta} = 18.1$ ($t = 2.2$). Leave-one-event-out across the six 2024 shocks leaves the slope in [19.1, 23.0], always significant; recursive leave-one-month-out across all 23 calendar months keeps the slope inside [16.4, 23.5]. For the 10-year breakeven the leave-one-month range is wider, [9.0, 22.1], with dropping November 2024 collapsing the breakeven coefficient to 9.0 ($t = 1.0$); the breakeven response is therefore more leverage-dependent than the slope, and jointly excluding Election Day and the assassination-attempt Monday leaves the slope at 17.6 ($t = 1.2$) and the 10-year breakeven at 6.2 ($t = 0.7$)—the daily-frequency magnitude is not separately identified outside the two largest-leverage observations, which we report rather than mask. Excluding all 57 macroeconomic-announcement days (FOMC, CPI, NFP) leaves the slope at 19.0 ($t = 4.8$) and the 10-year breakeven at 20.5 ($t = 3.9$). Decomposing Δp_t into positive and negative components (Table 5) shows the response loads almost entirely on Trump-gain days ($\hat{\beta}_+ = 21.4$, $t = 5.7$ for the slope; 40.2, $t = 8.7$ for 10Y), with the positive-shock 10-year coefficient remaining significant after dropping Election Day ($\hat{\beta}_{+,10Y} = 48.6$, $t = 2.1$); formal symmetry tests cannot reject $\hat{\beta}_+ = \hat{\beta}_-$ at conventional levels (two-sided $p \approx 0.10$ at the 10Y and 30Y), so the asymmetry is suggestive rather than established. A placebo split of Δp_t by contemporaneous S&P 500 sign yields nearly identical coefficients on up- and down-S&P days (18.8 vs. 19.2), and the Trump-gain pattern replicates on Betfair ($\hat{\beta}_+ = 85.8$, $t = 2.8$ vs. $\hat{\beta}_- = 21.5$, $t = 0.8$).

International Comparison and Cross-Asset Signature: The international evidence is mixed: the German 10-year Bund response is 6.5 ($t = 0.6$, $N = 442$, genuinely null), but the U.K. 10-year Gilt response is 19.6 ($t = 2.1$, $N = 436$) versus the U.S. 32.2 ($t = 4.8$); a Wald test cannot reject equality of the U.S. and Gilt coefficients at conventional levels, so we characterize the channel as a predominantly U.S. phenomenon with limited duration spillover to Gilts rather than a clean U.S.-specific identification. The cross-asset pattern is internally consistent with the fiscal-regime reading (Appendix Table A.2): the broad U.S. dollar appreciates by 1.9% per unit shock ($t = 5.3$); financials, energy, and industrials outperform (+9.3%, +5.3%, +5.3%, all $t \geq 2$); utilities and gold underperform (−3.7%, $t = -2.3$; −4.1%, $t = -2.4$). Simultaneous dollar strength, cyclical outperformance, and weakness in long-duration and inflation-hedge assets are jointly consistent with a fiscal-expansion-with-higher-rates regime rather than a global risk shock or a pure debasement story.

Mechanism Discriminators: A risk-aversion or flight-to-quality story would compress breakevens with the VIX, the opposite of what we observe; conditioning on ΔVIX leaves the political coefficient on the

10-year breakeven at 16.3 ($t = 3.2$) while ΔVIX itself is significantly negative. The MOVE response is symmetric and zero unconditionally but asymmetric—negative-shock days raise MOVE by 39.1 ($t = 2.0$) while positive-shock days do not—so the rate-vol response, like the curve response, is not a generic political-uncertainty premium. The Caldara et al. (2020) daily Trade Policy Uncertainty (TPU) index is essentially uncorrelated with Δp_t in the full sample ($r = 0.03$), but a quintile decomposition shows the channel operates through breakevens specifically when trade-policy uncertainty is salient: the 10-year breakeven coefficient is 8.3 ($t = 0.6$, $N = 91$) in the bottom TPU quintile and 22.5 ($t = 4.7$, $N = 91$) in the top; the asymmetry holds within the post-July-21 sub-sample alone (breakeven coefficients 17.5 vs. 22.2), ruling out a temporal-composition artifact. The orthogonal component of Polymarket's unified-Republican-control contract enters the breakeven regression with a coefficient of 0.9 ($t = 0.5$), so the signal loads on the presidential contract rather than on a separable congressional channel.

Inference Robustness: Pre/post-July-21 sub-period equality cannot be rejected for any of the 2s10s, 10Y nominal, 10Y breakeven, or 10Y real outcomes (Wald $|t| \leq 1.1$); a forward-looking placebo on Δp_{t+1} is statistically zero, ruling out reverse causality from yields to lagged probability movements. Block-bootstrap and Politis et al. (1999) subsampling intervals are in Appendix Table A.1.

Out-of-Sample Replication on the 2020 Cycle: We replicate the baseline design on the 2020 Biden–Trump cycle using daily PredictIt 2020 Trump-win prices ($N = 458$, January 2, 2019 to November 2, 2020; Appendix Table A.5). Sign agreement with 2024 holds across every long-end maturity, with the 2s10s slope at 43.3, the 30-year nominal at 88.7, and the 10-year breakeven at 66.6 (each $p < 0.10$); the 10-year TIPS coefficient is negative as in 2024, again concentrating the response in breakeven inflation. Once the COVID February 15–April 30 yield-volatility window is removed, magnitudes are economically close to 2024 (10-year breakeven 23.9 vs. 20.6; slope 32.5; 30-year nominal 36.4). The 2016 PredictIt sub-sample (Appendix Table A.4) is sign-consistent but slope-under-powered ($N = 55$). Two cycles separated by four years and a different prediction-market venue thus deliver the same channel composition.

Multiple Testing and Remaining Limitations: A Benjamini–Hochberg FDR adjustment to the five headline maturity coefficients (slope, 2Y, 5Y, 10Y, 30Y nominal in column (2) of Table 2) leaves all four long-end coefficients surviving at $q = 0.05$ and only the 2Y dropping out, consistent with the maturity-gradient story. The identification rests on a single well-identified cycle plus cross-cycle sign agreement; we do not claim a structural parameter recoverable from the daily-frequency point estimate, but rather a regime-conditional channel composition that is sign-consistent across the 2024 Polymarket and 2020 PredictIt cycles and statistically zero on the 2019 pre-period.

Table 5: Asymmetry: Positive versus Negative Probability Shocks

	Full sample		Excl. Election Day			
Outcome	$\hat{\beta}_+$	$\hat{\beta}_-$	$\hat{\beta}_+$	$\hat{\beta}_-$	N (full)	N (excl)
2s10s slope	21.4^{***}	9.0	33.5^{**}	7.4	453	452
	(3.7)	(21.0)	(16.0)	(20.6)		
10Y nominal	40.2^{***}	−10.4	48.6^{**}	−11.6	453	452
	(4.6)	(30.5)	(22.7)	(31.4)		
10Y breakeven	24.6^{***}	−0.3	15.6	0.9	453	452
	(3.0)	(18.5)	(11.5)	(15.0)		
10Y real	15.7^{***}	−10.1	33.0^{*}	−12.5	453	452
	(4.8)	(26.2)	(18.4)	(25.9)		
30Y nominal	43.4^{***}	−12.5	48.3^{**}	−15.5	453	452
	(4.2)	(29.1)	(22.0)	(28.9)		

Notes: Each column reports OLS coefficients on $\Delta p_t^+ = \max(\Delta p_t, 0)$ and $\Delta p_t^- = \min(\Delta p_t, 0)$, jointly estimated, with lagged-odds controls and Newey–West HAC standard errors (5 lags). $N = 453$ in the full-sample columns reflects the lagged-odds-control sample (the $N = 457$ baseline less four observations lost to lags); $N = 452$ in the right-hand columns drops Election Day (Nov. 6) from this sample. On positive-shock days, the long-end nominal response remains statistically significant after Election-Day exclusion (the Election Day shock is itself positive and is one of ∼200 positive-shock observations); on negative-shock days the coefficients are statistically zero in every specification. The full-sample symmetry tests $\hat{\beta}_+ = \hat{\beta}_-$ cannot reject equality at conventional levels in any specification; the marginal cases are the 10Y and 30Y nominal outcomes ($t \approx 1.6$–1.9, $p \approx 0.10$) and the apparent asymmetry is suggestive rather than formally established. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5 Conclusion

The maturity gradient is itself the economically informative object. Short-end yields and the overnight policy rate are unresponsive to Trump 2024 probability shocks, so markets did not price an imminent Fed tightening reaction; instead the response concentrates at five-year and longer horizons, loads almost entirely on the term premium (95% in Adrian et al. (2013)), and is consistent with the long-horizon inflation-uncertainty term-premium channel of Wright (2011). Read jointly with the cross-asset signature—a stronger dollar, outperformance of financials and energy, and weakness in utilities and gold—the pattern is the textbook signature of an unfunded-fiscal-expansion regime priced under monetary accommodation (Bianchi et al., 2023): rather than leaning against long-horizon fiscal and inflation pressures, the policy rule is expected to absorb them, leaving the entire risk-compensation burden at the long end of the curve. The response loads on positive Trump-gain shocks and on days with elevated trade-policy uncertainty (top-quintile-TPU breakeven coefficient 22.5 versus 8.3 in the bottom), so the channel is mechanism-conditional and not a generic political-uncertainty premium.

Identification rests on a non-trading-hours decomposition, a 2SLS estimator with overnight Δp_t as instrument ($F \approx 37$, IV exceeds OLS), Betfair replication on a structurally independent venue, a CME ten-year futures overnight gap on 1-minute Globex prices that is null ex-Election-Day, and a sign-consistent out-of-sample replication on the 2020 PredictIt cycle—the last addressing the most-immediate concern that Polymarket's 2024 microstructure could be driving the result. A formal Election Day holdout test predicts the November 6 slope move within 1 bp of the realized 8 bp; the recursive leave-one-month-out slope estimate stays inside [16.4, 23.5] across all 23 calendar months; and a 2019 pre-period placebo using the PredictIt 2020 contract is statistically zero across all maturities. We treat sign and channel composition as the load-bearing inference, not a daily-frequency structural parameter, and read the contribution as showing that the fiscal-regime channel is a coherent cross-platform, cross-asset, cross-cycle signature visible in the term premium itself rather than an Election-Day or 2024-cycle artifact. A natural next step is extension to other prediction-market-active sovereign-debt cycles.

References

Adrian, T., Crump, R. K., and Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1):110–138.

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

Bianchi, F., Faccini, R., and Melosi, L. (2023). A fiscal theory of persistent inflation. *The Quarterly Journal of Economics*, 138:2127–2179.

Boutchkova, M., Doshi, H., Durnev, A., and Molchanov, A. (2012). Precarious politics and return volatility. *Review of Financial Studies*, 25(4):1111–1154.

Caldara, D., Iacoviello, M., Molligo, P., Prestipino, A., and Raffo, A. (2020). The economic effects of trade policy uncertainty. *Journal of Monetary Economics*, 109:38–59.

Gürkaynak, R. S., Sack, B., and Wright, J. H. (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8):2291–2304.

Ibragimov, R. and Müller, U. K. (2010). t-statistic based correlation and heterogeneity robust inference. *Journal of Business & Economic Statistics*, 28(4):453–468.

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

Kim, D. H. and Wright, J. H. (2005). An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. Technical Report 2005-33, Federal Reserve Board, Finance and Economics Discussion Series.

Lahiri, S. N. (2003). *Resampling Methods for Dependent Data*. Springer, New York.

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

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

Politis, D. N., Romano, J. P., and Wolf, M. (1999). *Subsampling*. Springer, New York.

Santa-Clara, P. and Valkanov, R. (2003). The presidential puzzle: Political cycles and the stock market. *Journal of Finance*, 58(5):1841–1872.

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

Wright, J. H. (2011). Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review*, 101(4):1514–1534.

Appendix: Supplementary Tables

Table A.1: Block-Bootstrap Sensitivity, 2s10s Slope

Block length	$\hat{\beta}$	Bootstrap s.e.	Percentile 95% CI	Studentized 95% CI
5 days	20.30	13.45	[−2.84, +52.98]	[+10.15, +29.76]
10 days	20.30	15.60	[−3.65, +58.08]	[+9.67, +29.89]
20 days	20.30	16.61	[−6.99, +61.03]	[+9.53, +30.81]
Memo: Politis–Romano–Wolf subsampling 95% CI, $m = \lfloor n^{2/3} \rfloor = 59$ contiguous blocks				
Subsampling (Politis–Romano–Wolf)			[−56.94, +48.13]	

Notes: Outcome variable is the daily change in the 2s10s slope (bp). Point estimate $\hat{\beta} = 20.30$ (Newey–West s.e. 4.87, $n = 457$, specification (1) of Table 2). Moving-block bootstrap with 2,000 replications; percentile intervals use the 2.5th and 97.5th quantiles of the bootstrap $\hat{\beta}^*$ distribution. Studentized intervals invert the distribution of $(\hat{\beta}^* - \hat{\beta})/\text{s.e.}^*$ using Newey–West standard errors inside each replicate, rescaled by the original-sample s.e. (Lahiri, 2003). The Politis–Romano–Wolf subsampling interval (Politis et al., 1999) uses contiguous blocks of length $m = \lfloor n^{2/3} \rfloor$ and the standard $(m/n)^{1/2}$ rescaling. See main text §4 for interpretive caveats under extreme leverage.

Table A.2: Diagnostic Outcomes: Short-End Placebos, Foreign-Yield Spreads, Cross-Asset Signature

Outcome variable	$\hat{\beta}$	t	N
Short-end and overnight (expected-path placebos, all should be zero):			
Δ 3M Treasury (bp)	−1.7	−0.4	453
Δ 6M Treasury (bp)	1.2	0.2	453
Δ 1Y Treasury (bp)	3.0	0.4	453
Δ 2Y Treasury (bp)	12.8	1.6	453
Δ SOFR (bp)	−2.7	−0.8	451
Foreign-yield placebos (US–foreign spread):			
Δ Bund 10Y (bp)	6.5	0.6	442
Δ Gilt 10Y (bp)	19.6**	2.1	436
Δ (US 10Y − Bund 10Y) (bp)	24.5*	1.7	442
Δ (US 10Y − Gilt 10Y) (bp)	11.2	0.9	436
Cross-asset signature (log returns, percent):			
Δ log DXY (broad USD index)	1.9***	5.3	451
Δ log Gold	−4.1**	−2.4	451
Δ log XLF (financials ETF)	9.3***	2.9	451
Δ log XLE (energy ETF)	5.3**	2.0	451
Δ log XLI (industrials ETF)	5.3**	2.0	451
Δ log XLU (utilities ETF)	−3.7**	−2.3	451

Notes: Each row is a separate regression of the outcome on Δp_t with lagged-odds controls (specification 2 of Table 2); Newey–West HAC standard errors with five lags. Sources: Treasury yields and SOFR from FRED; Bund 10Y from the Deutsche Bundesbank; Gilt 10Y from the Bank of England nominal spot yield curve daily archive; DXY (DTWEXBGS) from FRED; gold and sector ETFs from Yahoo Finance via daily settlement prices. Spread regressions use $\Delta(\text{US}_{10Y} - \text{foreign}_{10Y})$ in basis points. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.3: Kim–Wright Daily Term-Structure Decomposition by Maturity

Maturity	KW yield $\hat{\beta}$	KW expected $\hat{\beta}$	KW term premium $\hat{\beta}$
1Y	12.2 ($t = 0.7$)	2.4 ($t = 0.7$)	9.7 ($t = 0.7$)
2Y	14.8 ($t = 1.3$)	4.5 ($t = 1.4$)	10.3 ($t = 1.3$)
5Y	18.5^{***} ($t = 2.7$)	8.1^{**} ($t = 2.4$)	10.4^{***} ($t = 2.9$)
10Y	19.1^{***} ($t = 3.7$)	8.7^{***} ($t = 3.3$)	10.5^{***} ($t = 4.0$)

Notes: Coefficients from regressions of daily changes in the Kim and Wright (2005) zero-coupon yield, expected-yield component, and term premium on Δp_t (lagged-odds controls; Newey–West HAC, 5 lags). Daily KW estimates are from the Federal Reserve Board's published series (FEDS 2005-33). The expected-yield response is statistically indistinguishable from zero at 1Y and 2Y horizons (consistent with the short-end placebo of Table A.2) and emerges only at the five-year and longer horizons, indicating long-horizon rather than near-term Fed-path repricing. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A.4: 2016 PredictIt Sub-Sample, Excluding the November 8 Cash/Futures Timing-Mismatch artifact

Outcome	$\hat{\beta}$ (t)	N
2Y nominal (bp)	6.3 ($t = 0.3$)	55
5Y nominal (bp)	41.6 ($t = 1.6$)	55
10Y nominal (bp)	32.7 ($t = 1.2$)	55
30Y nominal (bp)	41.1 ($t = 1.3$)	55
10Y TIPS (bp)	2.5 ($t = 0.1$)	55
10Y breakeven (bp)	30.2^{**} ($t = 2.2$)	55
2s10s slope (bp)	26.4 ($t = 1.3$)	55

Notes: Each row is a regression of the outcome on Δp_t with lagged-odds controls (Newey–West HAC, 5 lags). Δp_t is the daily change in the PredictIt 2016 Trump-win-probability contract, available September 22 to December 20, 2016 (90 trading days; image-extracted from PredictIt's daily archive). Treasury yields are FRED (DGS2, DGS5, DGS10, DGS30, DFII10). November 8, 2016 (Election Day) is excluded throughout because PredictIt's end-of-day price already reflected the call (Trump close-to-close $\Delta p = +0.60$) while the U.S. Treasury cash 10-year moved only +5 bp on November 8 and the bulk of the move (+19 bp) came on November 9 as the cash market repriced overnight; full-sample estimates are uninterpretable under this close-to-close mismatch and are omitted. The excluded-day estimates are sign-consistent with the 2024 results. The 2s10s slope coefficient (26.4, $t = 1.3$) is statistically insignificant; the implied standard error is ≈ 20 bp, giving an 80%-power minimum detectable effect of ~57 bp—more than twice the point estimate—so the 2016 slope sample is genuinely under-powered to confirm a magnitude in the 20–30 bp range. The 10-year breakeven coefficient (30.2, $t = 2.2$) is significant at the 5% level. We treat the 2016 evidence as corroborative for sign and channel but under-powered for confirming the precise slope magnitude. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A.5: 2020 PredictIt Out-of-Sample Replication

	Full pre-election sample		Excl. COVID Feb15–Apr30	
Outcome	$\hat{\beta}$	t	$\hat{\beta}$	t
2Y nominal (bp)	1.96	0.06	−1.7	−0.13
5Y nominal (bp)	28.33	0.70	9.0	0.36
10Y nominal (bp)	45.27	1.00	13.07	0.43
30Y nominal (bp)	88.67*	1.68	36.42	1.08
10Y TIPS (bp)	−21.37	−0.78	−9.5	−0.55
10Y breakeven (bp)	66.64*	1.73	23.93	1.38
2s10s slope (bp)	43.31*	1.75	32.49	1.58
N	458		406	

Notes: Each row is a regression of the daily yield change on Δp_t with lagged-odds controls (Newey–West HAC, 5 lags). Δp_t is the daily change in the PredictIt 2020 Trump-win contract, January 2, 2019 to November 2, 2020 (Election Day, November 3, 2020, is excluded as in the 2016 sub-sample because cash and contract repriced over different close-to-close windows). The right-hand panel additionally drops the COVID-19 yield-volatility window (February 15 to April 30, 2020), during which the 2-year Treasury yield repriced by hundreds of basis points within weeks for reasons unrelated to political-probability movements. Sign agreement with the 2024 baseline holds across every long-end maturity in both samples: the 30-year nominal coefficient is significant at 10%, the 10-year breakeven at 10% in the full sample and economically close to the 2024 magnitude (23.9 bp vs. 20.6 bp) once COVID volatility is removed. The 2s10s slope coefficient is positive throughout (43.3 and 32.5 vs. 19.4 in 2024) and approaches conventional significance ($t \approx 1.6$–1.8). The 2020 cycle is therefore sign- and channel-consistent with 2024 across the entire long end, with magnitudes plausibly inflated by the unusually large $\sigma(\Delta y_{10})$ of the 2020 macro environment relative to 2024. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A.6: Event-Cluster Inference: Per-Event Slope Responses

Event	Window	N_k	$\hat{\beta}_k$	s.e.
Biden Debate (Jun 27)	2024-06-27 to 2024-07-12	11	+75.46	(42.27)
Assassination Attempt (Jul 13)	2024-07-13 to 2024-07-20	5	+42.75**	(18.36)
Biden Drops Out (Jul 21)	2024-07-21 to 2024-08-05	11	+13.33	(41.45)
Harris Polling Surge (Aug 6)	2024-08-06 to 2024-09-30	39	−33.30	(36.11)
Polymarket Swing to Trump (Oct 1)	2024-10-01 to 2024-11-04	24	+10.24	(27.96)
Election Day (Nov 5)	2024-11-05 to 2024-11-06	2	+35.34	(·)
Average event effect $\bar{\beta} = K^{-1}\sum_k \hat{\beta}_k$			+23.97	(14.96)
Cluster t-statistic (df $= K-1 = 5$)			+1.60	
One-sided p-value (cluster t)			0.085	
Sign test: # positive out of $K = 6$, one-sided binomial p			5 (0.109)	

Notes: Each event k is assigned a non-overlapping window spanning from event k to one trading day before event $k+1$; the last event extends to the sample end. Within each window $\hat{\beta}_k$ is OLS on Δp_t with Newey–West HAC standard errors and no additional controls; the Election Day window contains two observations, so the within-event standard error is not informative. The cluster t-statistic follows Ibragimov and Müller (2010): $t_{\text{IM}} = \bar{\beta}/\text{s.e.}(\bar{\beta})$ with $\text{s.e.}(\bar{\beta}) = \sqrt{\sum_k (\hat{\beta}_k - \bar{\beta})^2/[K(K-1)]}$, distributed t_{K-1} under the null and valid under heterogeneous variances. The cluster one-sided p-value of 0.085 and the sign-test p-value of 0.109 are above conventional 5% thresholds; we read this $K = 6$ test as directionally consistent with the daily-frequency result rather than as independently confirming it. $^{**}p < 0.05$ at the within-event level.

Table A.7: Cross-Platform Validation: 2s10s Slope by Odds Source

Source	Type	N	$\hat{\beta}$	t	Hierarchy[a]
Polymarket	P2P exchange (crypto)	453	19.4	4.5	Discovery
Betfair	P2P exchange (UK)	210	47.3	3.7	Discovery
BetOnline	Sharp sportsbook	171	39.5	1.5	Sharp
Bovada	Sharp sportsbook	210	−0.9	0.0	Sharp
Pinnacle	Sharp sportsbook	83	8.6	0.5	Sharp
Unibet	Retail bookmaker	207	7.8	0.6	Retail
William Hill	Retail bookmaker	128	−0.6	0.0	Retail
Everygame	Retail bookmaker	131	3.9	0.1	Retail

Notes: Dependent variable is the daily change in the 2s10s slope in basis points. Each row is estimated from Equation (1) with lagged-odds controls (specification 2 of Table 2); all odds on the 0–1 probability scale. Newey–West HAC standard errors with five lags.
[a] Tier reflects within-sample behavior: Polymarket and Betfair share near-simultaneous price discovery ($r = 0.87$ at lag zero on the matched sample); sharp sportsbooks produce directionally consistent but individually insignificant slope responses; retail bookmakers revise quotes on fewer than 35% of trading days and yield coefficients near zero across all specifications.