IPE Business School

Prediction Markets: Information Aggregation, Political Uncertainty and Financial Markets

di Andrea Roncella

IPE Working Paper

N. 36

November 4, 2026

ISSN 2284-1229

https://doi.org/10.5281/zenodo.23956774

Prediction Markets: Information Aggregation, Political Uncertainty and Financial Markets

Andrea Roncella[1]

Abstract

This paper reviews prediction markets as mechanisms for aggregating information and examines their links with political uncertainty and financial markets. It discusses contract design, price calibration and the concentration of trading profits, drawing on evidence from Polymarket. Because direct platform data were unavailable, the empirical analysis uses the Economic Policy Uncertainty (EPU) index over 1985–2026, together with financial indicators and US and European equity returns. Event studies, correlations and OLS models show a negative contemporaneous association between changes in EPU and equity returns, alongside higher volatility in high-uncertainty regimes. An extension with the Fama–French three factors finds that the EPU coefficients lose significance when the market factor is included. This result largely reflects the close relationship between that factor and the S&P 500 excess return, highlighting a limitation of the aggregate time-series specification. A methodological extension outlines how conflict-related Polymarket contracts could be aggregated into a geopolitical risk indicator, distinct from EPU and newspaper-based GPR. The paper discusses implications for risk monitoring and scenario design.

Keywords: prediction markets; Polymarket; economic policy uncertainty; information aggregation.

JEL Classification: D81; D84; G12; G14; G32.

Il lavoro esamina i mercati predittivi come meccanismi di aggregazione delle informazioni e ne discute i legami con l'incertezza politica e i mercati finanziari. La rassegna considera struttura dei contratti, calibrazione dei prezzi e concentrazione dei profitti, con riferimento a Polymarket.

[1] Università Cattolica del Sacro Cuore.

In assenza di dati diretti delle piattaforme, l'analisi empirica utilizza l'indice Economic Policy Uncertainty (EPU) nel periodo 1985–2026, insieme a indicatori finanziari e rendimenti azionari statunitensi ed europei. Studi di evento, correlazioni e modelli OLS mostrano un'associazione negativa tra variazioni dell'EPU e rendimenti, oltre a maggiore volatilità nei regimi di elevata incertezza. L'inclusione dei tre fattori Fama–French elimina la significatività dei coefficienti EPU, soprattutto per la stretta relazione tra il fattore di mercato e il rendimento in eccesso dello S&P 500. Il risultato evidenzia un limite della specificazione aggregata. Un'estensione metodologica illustra come aggregare contratti Polymarket in un indicatore geopolitico, distinto da EPU e GPR giornalistico. Si discutono implicazioni per il monitoraggio del rischio e la costruzione di scenari.

Parole chiave: mercati predittivi; Polymarket; incertezza di politica economica; aggregazione informativa

1 Introduction

Prediction markets are trading venues in which contracts are written on the outcome of future, verifiable events — an election, a Federal Reserve decision, the release of a macroeconomic statistic — rather than on the cash flows of a firm. Their appeal to economists rests on a simple theoretical intuition, traceable to Hayek (1945): prices are an efficient mechanism for aggregating dispersed private information, and a market that lets participants bet directly on the truth of a proposition should, under reasonably weak conditions, produce a better forecast of that proposition than any individual participant, pollster or pundit.

Prediction markets have existed in academic and quasi-academic form since the Iowa Electronic Markets were established in 1988, but they have recently undergone a dramatic scale transformation with the emergence of blockchain-based platforms. Polymarket, the largest such platform, has grown from a niche crypto-native product into a market processing hundreds of millions of trades and tens of billions of dollars in volume on political and economic events, attracting the attention of regulators, media organisations and, most relevantly here, financial economists.

This paper has two aims. The first is conceptual: to lay out the economics of prediction markets — contract design, the theory of information aggregation, and the empirical evidence on who trades and who profits — drawing on the classic treatment of Wolfers and Zitzewitz (2004) and on the recent large-scale empirical study of Polymarket by Akey et al. (2026).

A useful distinction is between economic policy uncertainty and geopolitical risk. The **Geopolitical Risk (GPR) index** of Caldara and Iacoviello (2022) measures newspaper coverage of adverse geopolitical events and associated risks, including war and terrorism. Its threats and acts components distinguish discussion of potential escalation from reports of realised adverse events. EPU instead concerns uncertainty about economic policy, including taxation, regulation and trade. The two constructs can overlap, but they are not interchangeable: a change in domestic fiscal policy need not imply greater conflict risk, and a military escalation can occur without a domestic policy change. Prediction markets offer a third observation mechanism: prices attached to specific, verifiable future outcomes. Section 3.5 explains how such prices could be aggregated into a geopolitical indicator, using GPR as a conceptual benchmark rather than treating EPU as its substitute.

The second aim is empirical. Direct data access to Polymarket, Kalshi and Manifold was attempted for this project but blocked at the network level (connection resets on both direct HTTP requests and web-fetch tools used in this research environment); this limitation is documented in Appendix A. In its place, we use the Economic Policy Uncertainty (EPU) index of Baker et al. (2016) as an observable, text-based measure of the broader policy-uncertainty environment. This provides a complementary empirical exercise on uncertainty and asset prices; it does not reproduce event-contract prices or identify a market-implied probability of geopolitical

conflict. The regressions below use EPU, not GPR or a Polymarket-based geopolitical index. Their results therefore cannot establish the forecasting performance of such an index, whose construction remains a proposed extension.

Using EPU we: (i) document its historical evolution and identify, from the data themselves, the months of highest measured uncertainty; (ii) run an event study around five episodes with a direct prediction-market analogue (Brexit, the 2016, 2020 and 2024 US presidential elections, and the onset of COVID-19); (iii) estimate its correlation with the VIX, WTI crude oil, and corporate and Treasury spreads; (iv) estimate its impact on US (S&P 500) and European (Euro Stoxx 50) equity returns using regression, cross-correlation and regime analysis; and (v) augment the equity regression with the Fama–French three factors, which is the original empirical contribution of this paper.

The paper is organised as follows. Section 2 introduces the concept of prediction markets, their history and contract design. Section 3 discusses the Polymarket case study. Section 4 presents the empirical analysis based on EPU. Section 5 draws implications for firms and investors. Section 6 concludes.

2 The Concept of Prediction Markets

2.1 Definition and brief history

Definition 2.1. *A **prediction market** is a market in which contracts are traded whose final payoff is contingent on the realisation of a specific, verifiable future event or continuous outcome, such that the equilibrium price of the contract can be interpreted as the market's aggregated assessment of the probability (or expected value) of that outcome (Wolfers and Zitzewitz, 2004).*

The academic use of prediction markets dates at least to the **Iowa Electronic Markets** (IEM), launched in 1988 at the University of Iowa to let small numbers of traders bet real (if small) sums of money on the outcome of US elections; the IEM has since become a standard reference point for studies of market-based forecasting because its accuracy could be directly compared, election after election, with contemporaneous opinion polls.

A more controversial episode was the Defense Advanced Research Projects Agency's proposed **Policy Analysis Market** (PAM) in 2001–2003, which would have allowed trading on contracts linked to geopolitical and economic developments in the Middle East, including, in its original design, contracts related to political instability and conflict events. The project was cancelled within days of becoming public in 2003 after being characterised in the press and in Congress as a "terrorism futures market"; the episode is frequently cited in the literature (Wolfers and Zitzewitz, 2004) as an illustration of the tension between the information-aggregation benefits of

prediction markets and the ethical and political sensitivities raised when the underlying events involve violence, death or national security.

The 2010s and 2020s saw the emergence of **blockchain-based** prediction markets — Augur, and more recently Polymarket and Kalshi — that removed the small-scale, IRB-approved-experiment character of the IEM and allowed effectively unrestricted retail and institutional participation, denominated in stablecoins (Polymarket) or, following a change in US regulatory treatment, in dollars via a licensed exchange (Kalshi). This scale-up is what makes prediction markets newly relevant to financial economists, and is the subject of the case study in Section 3.

2.2 Contract design

Wolfers and Zitzewitz (2004) classify prediction market contracts into three broad types:

Tab. 1. Types of Prediction Market Contracts (following Wolfers and Zitzewitz, 2004)

Contract type	**Description**
Winner-take-all	Pays a fixed amount (e.g. $1) if a specified event occurs, zero otherwise. The equilibrium price is interpretable, under risk neutrality of the marginal trader, as the market-implied probability of the event. This is the dominant contract design on Polymarket and Kalshi ("Will X happen by date Y?").
Index contracts	Pay off in proportion to a continuously measured outcome (e.g. the margin of victory, or a vote share), rather than a binary indicator. Index contracts reveal the market's expectation of the mean of a distribution, and are common in the IEM's vote-share markets.
Spread contracts	Pay a fixed amount if a continuous outcome exceeds a threshold that is itself set so that the contract trades near an even-money price; the threshold, not the price, is the object that reveals information (as in point-spread betting on sports).

Winner-take-all contracts are the natural design for questions of direct relevance to political and macroeconomic uncertainty (elections, central bank decisions, geopolitical developments) and are the type discussed throughout the rest of this paper.

2.3 Information aggregation and the marginal trader hypothesis

The theoretical case for prediction markets rests on Hayek's (1945) argument that the price system economises on information: no single agent, planner or forecaster possesses the sum of information dispersed across market participants, but a price arrived at through decentralised trading can, in principle, reflect that dispersed information without any individual needing to possess all of it.

Applying this logic to a market with only a handful of active traders, as the early IEM often was, requires an argument for why the price should reflect anything beyond the beliefs of an arbitrary trader. This is supplied by the **marginal trader hypothesis** of Forsythe et al. (1992): even if most participants in a political stock market hold biased or unsophisticated beliefs (many IEM traders, for instance, disproportionately backed their own preferred candidate, regardless of the candidate's objective chances), prices can still be accurate if a small set of well-informed, unbiased *marginal traders* are willing to take the other side of mispriced bets and are not capital constrained. Forsythe, Nelson, Neumann and Wright show, using individual trader-level data from IEM political markets, that the marginal traders who determine prices are indeed better calibrated and more sophisticated than the average participant, which reconciles the apparent paradox of accurate market-level forecasts drawn from a pool of individually biased traders.

This distinction — between the (possibly large and behaviourally biased) population of retail participants and the (small, sophisticated) population of marginal, price-setting traders — reappears as the central empirical finding of the Polymarket case study in Section 3, where it takes the form of extreme concentration of trading profits among a small set of liquidity-providing accounts.

2.4 Efficiency, biases and the concentration of profits

A large empirical literature has tested whether prediction market prices are well calibrated, i.e. whether contracts trading at p cents resolve in-the-money with frequency close to $p\%$. The evidence is generally supportive of good average calibration, but with well-documented deviations. Snowberg and Wolfers (2010) document a robust **favourite-longshot bias** across betting and prediction markets: contracts priced as longshots (low probability of paying off) tend to be overpriced relative to their realised frequency, while favourites tend to be slightly underpriced, a pattern the authors argue is better explained by systematic misperception of small probabilities than by risk-love per se.

Manski (2006) raises a further conceptual caveat: even a perfectly efficient prediction market price need not equal the population's mean subjective probability of the event, because the

equilibrium price is a function of the distribution of beliefs *and* of risk tolerances and wealth across traders; interpreting prices literally as "the" probability therefore requires assumptions (e.g. homogeneous risk-neutral marginal traders) that need not hold empirically. This caveat matters directly for the Polymarket evidence discussed next, where good *average* calibration coexists with extreme concentration of the underlying trading profits: aggregate accuracy is a market-level property that need not imply that the average trader is well informed or unbiased.

3 Case Study: Polymarket

Polymarket is, as of 2026, the largest prediction market platform by trading volume, settling contracts on political, economic, sports and cultural events in USDC. Akey et al. (2026) provide, to our knowledge, the most complete empirical account to date of trading outcomes on the platform, reconstructing the complete transaction history from Polymarket's launch in November 2022 through the end of their sample. Because direct API access to Polymarket was blocked from the network used in this study (see Appendix A), the discussion below relies entirely on the published findings of Akey et al. (2026), whose sample spans November 11, 2022 through March 29, 2026, covering over 2.4 million users, 588.3 million trades and $67 billion in trading volume.[1]

3.1 Price calibration

Consistent with the broader efficiency literature discussed in Section 2, Akey et al. (2026) find that Polymarket prices are, on aggregate, well calibrated: a contract priced at p resolves in its favour approximately p percent of the time, and this calibration is already close to the 45-degree benchmark a full month before resolution and tightens further as the resolution date approaches. This average-level calibration, however, coexists with substantial cross-sectional heterogeneity in *who* captures the gains implied by that calibration.

3.2 Concentration of trading gains and losses

The central empirical finding of Akey et al. (2026) is the extreme concentration of trading profits alongside widespread losses. Across the 764,424 users with positive PnL in their sample (an aggregate net gain of $1.05 billion), the top 0.1% of winners capture 51.2% of all gains, the top 1% capture 76.5%, the top 5% capture 91.4%, the top 10% capture 96.2% and the top 25% capture 99.2%. Losses are concentrated to a slightly lesser extent: the top 0.1% of

[1] We rely on the working paper as consulted during the drafting of this project. No Polymarket order-book or trade-level data of any kind were retrieved independently for this paper; every figure attributed to Akey et al. (2026) below is taken directly from the text of that paper.

losers account for 43.7% of total losses, the top 1% for 68.1%, and the top 5% for 86.1%. Despite this concentration of gains among a small elite, roughly 69% of all users in the sample end with a loss overall. This concentration is far more extreme than typical stock-market trading-profit distributions, and closer to the concentration patterns documented in retail options trading or day-trading in equities.

3.3 Maker vs. taker and the "excess hit rate"

Akey et al. (2026) attribute a large part of this concentration to the distinction between **market makers** (traders who post limit orders and are filled passively, earning the bid-ask spread) and **takers** (traders who cross the spread with market orders, paying it). Their probit estimates identify the maker volume share as the single strongest behavioural predictor of positive performance in the cross-section: a one-standard-deviation increase in a user's maker volume share raises the probability of positive performance by approximately 9.0 percentage points, and the top 0.1% of earners supply 47.3% of their volume as makers, versus only 17.1% for the bottom 95% of users. A direct decomposition of the bid-ask spread confirms the mechanism: winning traders are net *earners* of the spread (the top 0.1% of winners earn 3.2% more than they would in the absence of spread costs), while losing traders are net payers of it (removing even the minimum-tick spread cost would move 18.5% of losers into non-negative PnL). Consistent with this, the volume-weighted **excess hit rate** (the gap between a user's realised win rate and the price at which they traded) falls to roughly -0.15 to -0.18 across the loss tail and rises to about $+0.20$ for the most profitable users; a finite-mixture decomposition of this gradient attributes genuine forecasting skill to only about 29.3% of traders, cautioning against reading the full gradient as evidence of widespread skill. A large share of trading activity is concentrated in longshot contracts (priced below 10 cents) and equivalent favourites (above 90 cents) — 54.9% of trades on average, 58.3% at the median — linking the favourite-longshot bias of Snowberg and Wolfers (2010) to the realised losses of unsophisticated, liquidity-taking retail traders on Polymarket.

3.4 Is it insider trading?

Given the concentration of profits among a small number of highly successful accounts, a natural question is whether those accounts are trading on material non-public information (e.g. private knowledge of an election result or a regulatory decision) rather than superior forecasting skill. Akey et al. (2026) study the 100 most successful accounts in their sample in detail — together representing 27.7% of aggregate winners' profits — and find that a meaningful fraction of them earn the bulk of their profit from liquidity provision (posting limit orders across many markets) rather than concentrated directional bets. Among the remaining traders whose gains

are concentrated in a small number of markets, the largest source of profit by far is sports betting (81% of their gains, spread across many teams and leagues with no single event dominating), with political markets — including the 2024 US presidential election — a distant second (15%). The authors also cross-check against the small number of traders actually charged with trading on material non-public information (one Google employee and one US soldier prosecuted by the CFTC): both rank outside the top 100 winners (115th and 307th of 2.48 million users, respectively) and together account for only about 0.13% of aggregate winners' profits. On this basis, Akey et al. (2026) conclude that *insider trading is unlikely to explain the performance of the largest winners*: the breadth of winning positions across dozens of sports leagues is difficult to reconcile with any single insider source, and the pattern is instead consistent with a smaller population of professional, market-making traders whose edge derives from execution style and possibly genuine forecasting skill. The authors caution that even the modest month-to-month performance persistence they document may partly reflect survivorship (selection into continued trading) rather than demonstrated skill.

As a parallel from the corporate-finance literature on the informativeness of dispersed, non-professional information sources, Akey et al. (2026) themselves note the work of Cookson et al. (2026), who show that social media sentiment can inform corporate decisions — specifically, around merger withdrawals — and pose the question of whether prediction markets could play an analogous informational role for corporate decisions and investor behaviour more broadly.[2] The parallel is direct: in both settings, aggregating information from a large, heterogeneous crowd need not require every participant to be informed, but does appear to require the *informative subset* of participants (fundamentals-oriented social media users; market-making Polymarket traders) to be identifiable, at least statistically, from posting style or execution style — echoing the marginal trader hypothesis of Forsythe et al. (1992) in a modern, high-frequency, crowd-sourced information environment.

3.5 From event prices to a geopolitical risk indicator

The newspaper-based benchmark. The GPR index of Caldara and Iacoviello (2022) is constructed from the share of newspaper articles concerning adverse geopolitical events and associated risks. Its recent series draws on ten newspapers; monthly and daily measures are available. The Geopolitical Threats component concerns potential adverse developments, whereas the Geopolitical Acts component concerns their realisation. These measures provide a historical benchmark for geopolitical risk as reflected in news coverage. They do not directly

[2] We cite Cookson et al. (2026) for the existence and general thrust of this result, as reported by Akey et al. (2026); we did not independently verify the magnitude of their merger-withdrawal estimates for this paper, and do not report a specific coefficient here to avoid overstating what we can confirm.

assign a probability to a particular event with a specified deadline, and neither component should be equated with an event-contract probability.

Selecting and orienting contracts. Polymarket could provide the inputs for a complementary indicator by combining binary questions on clearly specified conflict-related outcomes: for example, an attack, a military intervention or a ceasefire. Electoral popularity and domestic policy questions should not enter merely because they concern politics; relevance to geopolitical conflict requires an explicit selection rule. Each contract should retain its unique identifier, resolution rules, geographical scope and deadline. Ambiguous questions should be excluded or independently reviewed, and the selection dictionary and directional labels should be documented before testing performance.

Let $y_{i,t}$ denote the observed Yes price for contract i at time t, interpreted cautiously as a market-implied probability as discussed in Section 2. Assign $d_i = +1$ when Yes corresponds to a risk-increasing outcome, and $d_i = -1$ when it corresponds to de-escalation. A common risk orientation is then

$$p_{i,t} = \begin{cases} y_{i,t}, & d_i = +1, \\ 1 - y_{i,t}, & d_i = -1. \end{cases}$$

Thus an increase in the assessed probability of an attack raises the risk price, while an increase in the assessed probability of a ceasefire lowers it. Failure to reach a ceasefire by its deadline is nevertheless not identical to a new attack; reorientation aligns signs without making the underlying events equivalent.

Aggregation and interpretation. Let $\mathcal{A}_t$ contain eligible, still-open contracts with informative trading observations on date t. An illustrative volume-weighted construction is

$$I_t^{\mathrm{PM}} = 100 \sum_{i \in \mathcal{A}_t} w_{i,t} p_{i,t}, \qquad w_{i,t} = \frac{v_{i,t}}{\sum_{j \in \mathcal{A}_t} v_{j,t}},$$

where $v_{i,t}$ is consistently measured daily trading volume. If no eligible contract trades, the observation is missing rather than zero. An equal-weighted baseline is also useful, since volume is an imperfect proxy for information and may reflect speculation or artificial activity. For a hypothetical attack contract priced at 0.20 and a ceasefire contract priced at 0.70, equal weights produce an index of 25: the two oriented prices are 0.20 and 0.30. This illustrates the scale, not an observed result.

The index is an average assessment of adverse outcomes among the selected contracts. It is not the probability of a global war, the probability that at least one listed event occurs, or a measure of expected economic losses. Overlapping questions can price the same conflict repeatedly, and volume weights can concentrate the indicator on the most actively traded region. Changes can arise from prices, weights, entry and exit. Different remaining maturities also affect event

probabilities, so a practical implementation should restrict or report horizons and compare contracts with similar deadlines. Known outcomes must not be inserted into earlier observations.

Validation and the scope of this paper. Before use, the indicator would require independent checks of relevance and direction, price calibration against subsequent resolutions, sensitivity to liquidity and weighting rules, and monitoring of regional concentration. Its information content could then be compared with GPR using aligned information timestamps and genuine out-of-sample forecasts. A lead over newspaper coverage is an empirical hypothesis, not a consequence of using market prices. Such an exercise would distinguish changing event probabilities from changing media attention while recognising that news also contains expectations and markets also react to published news. Because direct platform data were unavailable for this study, no I_t^{PM} series is estimated here. The construction is a proposal for future research, separate from the EPU results in the next section.

4 Empirical Analysis

4.1 Data

All series are monthly (period-average or month-end, as noted) and are described in full in Appendix A. The core variable is the **Economic Policy Uncertainty (EPU)** index of Baker et al. (2016) (FRED series `USEPUINDXD`, daily, resampled to monthly averages), available January 1985–June 2026 ($n = 498$ months). We also use the VIX, WTI crude oil, the Moody's BAA–10Y Treasury credit spread and the 10Y–2Y Treasury term spread (all FRED), the Euro Stoxx 50 (LSEG/Eikon, RIC `.STOXX50E`) and the S&P 500 (Yahoo Finance, `^GSPC`), and the Fama–French three-factor monthly series (Mkt-RF, SMB, HML, RF) from Kenneth French's data library.

4.2 Historical evolution of EPU

Table 2 and Figure 1 summarise the full-sample distribution of EPU (1985–2026, $n = 498$): mean 119.3, median 97.2, standard deviation 75.0, skewness 2.75, minimum 38.2, maximum 560.9. The index exhibits high persistence, with a first-order monthly autocorrelation of 0.896.

Tab. 2. EPU Index: Descriptive Statistics (Jan. 1985 – Jun. 2026, $n = 498$)

Mean	Median	Std. Dev.	Skewness	Min	Max	AR(1)
119.3	97.2	75.0	2.75	38.2	560.9	0.896

Source: own elaborations on Baker et al. (2016) data (FRED, `USEPUINDXD`).

Rather than fix *a priori* which episodes should be the most significant, we report the months with the highest *observed* EPU values in the full 1985–2026 sample. Table 3 shows that the ten highest EPU readings on record cluster almost entirely in two episodes: the acute phase of the COVID-19 pandemic (March–July 2020) and a sustained period of elevated uncertainty from March 2025 through June 2026, coinciding with a major escalation of US trade-policy actions (tariff announcements) in the spring of 2025 and persisting, at a somewhat lower level, through the most recent data available.

Tab. 3. Ten Highest EPU Readings on Record (Jan. 1985 – Jun. 2026)

Rank	**Month**	**EPU**
1	April 2025	560.9
2	April 2020	555.3
3	March 2025	503.7
4	May 2025	499.3
5	May 2020	484.2
6	March 2026	406.3
7	June 2025	399.7
8	March 2020	380.4
9	May 2026	372.1
10	July 2025	372.0

Note: Peaks are identified directly from the data, not fixed *a priori*. *Source:* own elaborations on Baker et al. (2016) data (FRED).

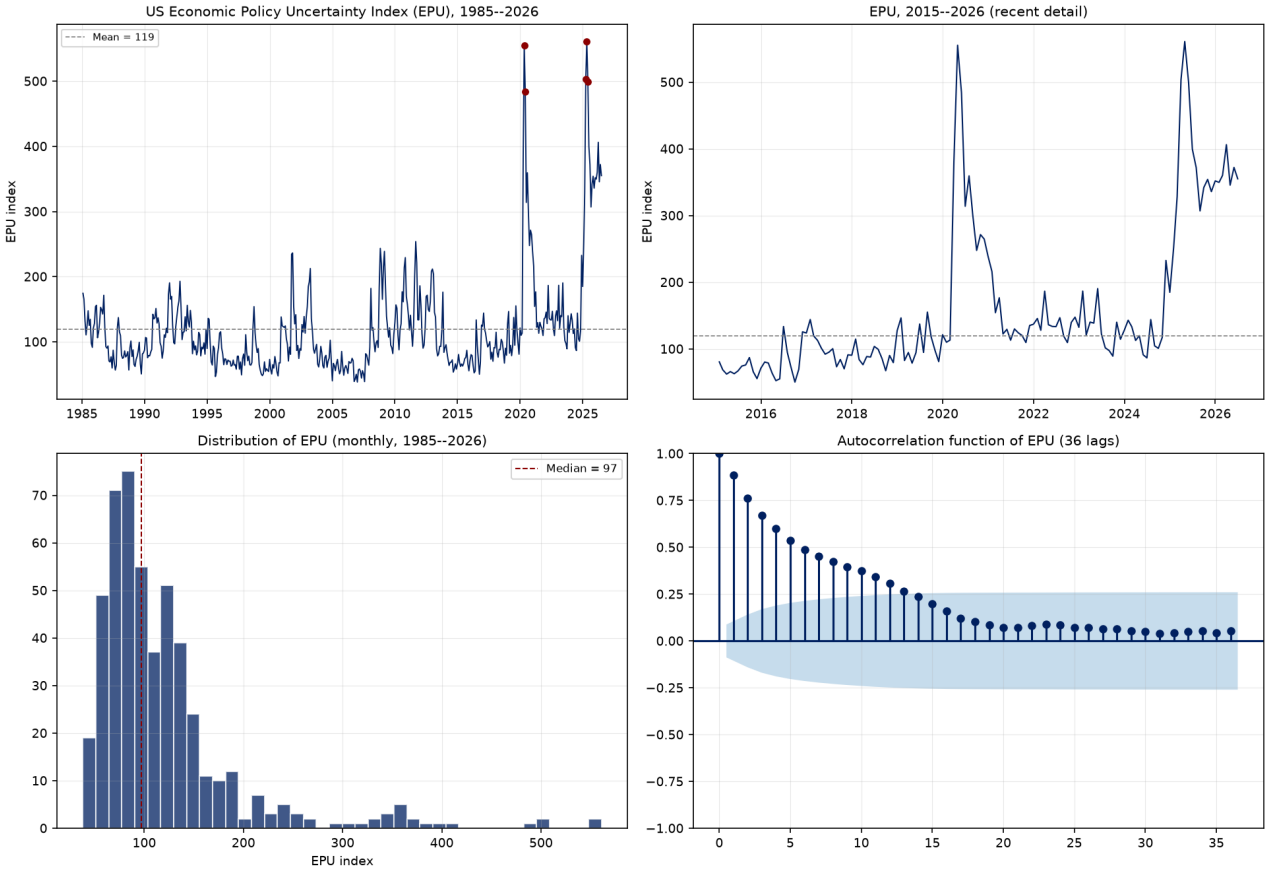


Fig. 1. US Economic Policy Uncertainty Index — Historical Series, Recent Detail, Distribution and Autocorrelation (1985–2026). Top left: full-sample series with the five highest peaks marked in red. Top right: 2015–2026 detail. Bottom left: full-sample distribution. Bottom right: autocorrelation function (36 lags) with 95% confidence band.

4.3 Event Study: EPU Response around Political and Health Shocks

To describe uncertainty around these events, we compute, for each event i, the pre-event baseline $\overline{\text{EPU}}_i^{\text{pre}}$ as the average EPU over the 24 months preceding the event (Equation 4), and the percentage deviation from baseline $\delta_t^{(i)}$ in the $[-24, +24]$-month window (Equation 5), exactly as in Section B.

We examine five episodes directly relevant to prediction markets: the Brexit referendum (June 2016), the 2016, 2020 and 2024 US presidential elections, and the onset of COVID-19 (March 2020). Table 4 reports the results.

Tab. 4. EPU Response to Political and Health Shocks: Intensity and Persistence

Event	Date	Baseline	$\delta^{\max}$ (%)	Peak month	Return to baseline
Brexit referendum	Jun 2016	67.1	+114.7	+7	+3 months
US election 2016	Nov 2016	71.3	+102.1	+2	+11 months
US election 2020	Nov 2020	193.5	+187.0	−7[a]	+3 months
US election 2024	Nov 2024	125.2	+348.0	+5	> 19 months[b]
COVID-19 onset	Mar 2020	101.4	+447.8	+1	> 24 months

Notes: Baseline = mean EPU over the 24 months preceding the event. Peak month and return-to-baseline month are measured relative to the event month (month 0). [a] For the 2020 US election, the maximum deviation within the $[-24, +24]$ window actually occurs *before* the event (month -7 = April 2020), reflecting the COVID-19 EPU spike that falls inside the pre-election window; the deviation profile therefore mixes the pandemic shock and the election itself, and should not be read as a pure election effect. [b] For the 2024 US election, the +24-month window extends to November 2026, beyond the last available data point (June 2026); only 19 months of post-event data are observed, over which EPU has not returned to baseline. *Source:* own elaborations on Baker et al. (2016) data.

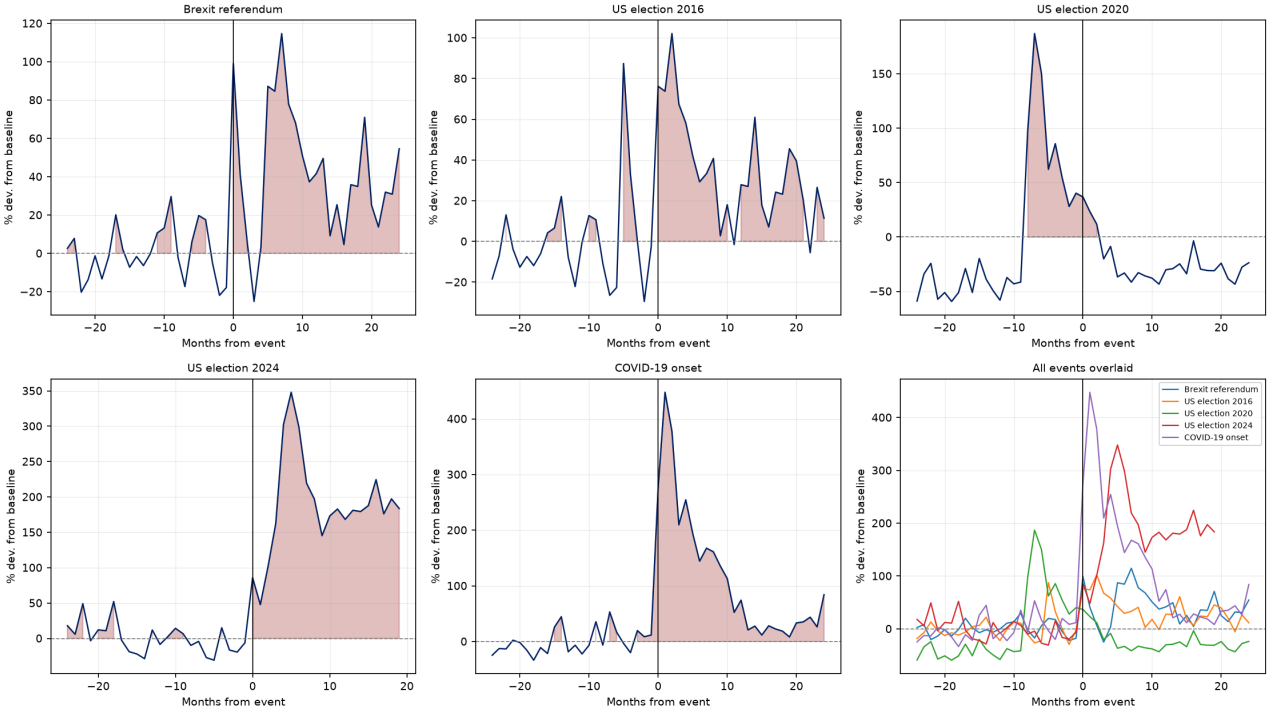


Fig. 2. Event Study — EPU Response to Political and Health Shocks (window: ±24 months, except the 2024 US election, truncated at +19 months by data availability). The red shaded area indicates months in which EPU exceeds the pre-event baseline (dashed grey line). Bottom right panel overlays all five events.

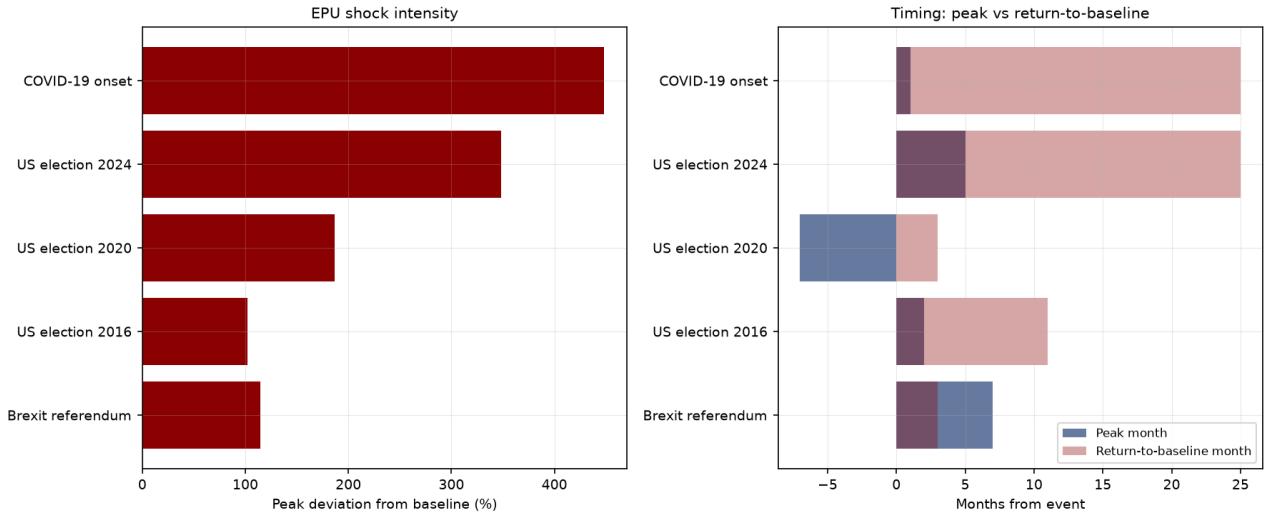


Fig. 3. Intensity and Persistence of EPU Shocks. Left: peak deviation from baseline (%). Right: peak month vs. return-to-baseline month (a return-to-baseline value at the right edge of the window indicates no return within the observed data).

Three findings stand out. First, the COVID-19 shock produced by far the largest deviation from baseline (+447.8%), consistent with its classification, across many uncertainty indices, as the single largest peacetime uncertainty shock in the sample. Second, the 2024 US election is associated with an unusually large and *persistent* deviation (+348.0% peak, no return to baseline within the 19 months of post-event data available), coinciding with — and plausibly driven by — the trade-policy escalation of spring 2025 identified independently in Table 3; this illustrates that in the US context, EPU around a presidential election captures anticipated *policy* uncertainty (tariffs, regulation) realised in the following months, not merely the uncertainty of the vote count itself. Third, the 2020 election window is contaminated by the COVID shock that precedes it, a reminder that overlapping windows must be interpreted with care in event-study designs applied to a persistent, clustered series such as EPU.

4.4 EPU and Macro-Financial Variables

Figure 4 and Table 5 report Pearson correlations between EPU and four macro-financial variables (VIX, WTI crude oil, the Moody's BAA–10Y credit spread, and the 10Y–2Y Treasury term spread) over their common sample (January 1990–June 2026, $n = 438$).

Tab. 5. EPU and Macro-Financial Variables: Pearson Correlations (1990–2026, $n = 438$)

	VIX	**WTI**	**BAA10Y**	**T10Y2Y**
Correlation with EPU	0.360	0.156	0.173	0.041
p-value	$< 0.001^{***}$	0.0011^{**}	0.0003^{***}	0.396

Notes: Significance: $^{*}\,p < 0.05$; $^{**}\,p < 0.01$; $^{***}\,p < 0.001$. *Source:* own elaborations on FRED data.

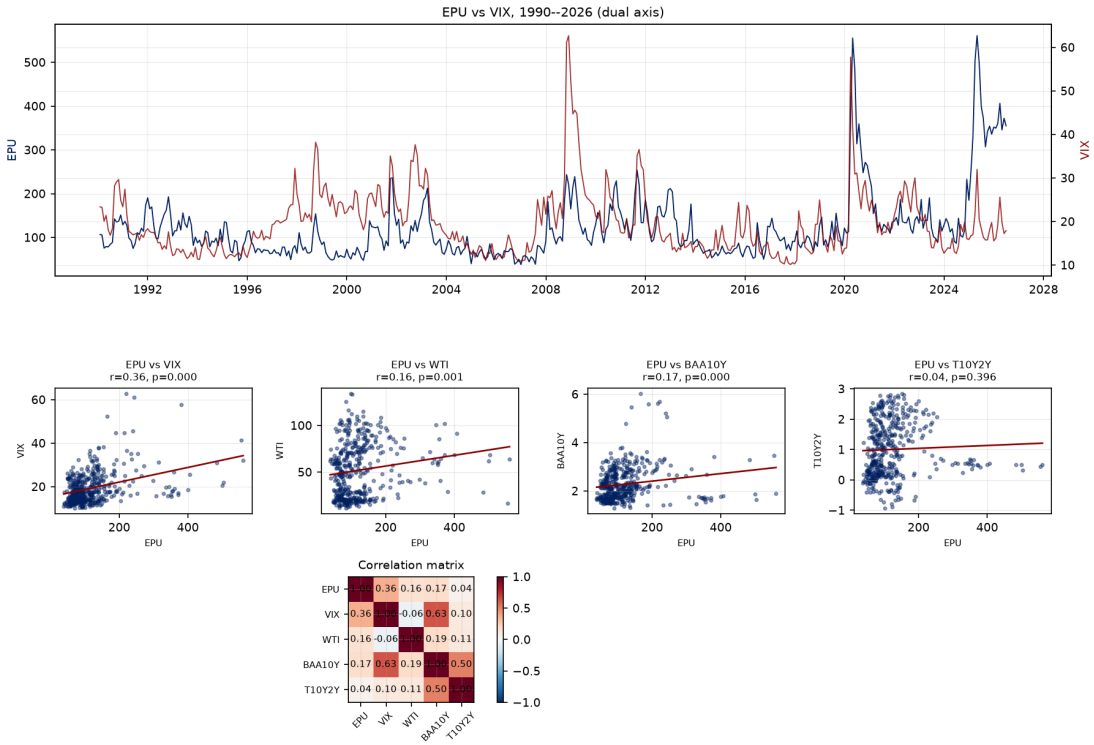


Fig. 4. EPU vs Macro Variables (1990–2026). Top: EPU and VIX on dual axis. Middle row: scatter plots with OLS fit for EPU vs. VIX, WTI, the BAA10Y credit spread and the T10Y2Y term spread. Bottom: full correlation matrix.

The estimated correlations are all positive. Their interpretation depends on the variable: the VIX and credit spreads capture dimensions of financial stress, whereas oil prices and the term spread also reflect other forces. These are EPU associations; no corresponding GPR correlations are estimated in this paper. The EPU–VIX correlation ($r = 0.360$, $p < 0.001$) is the strongest and most precisely estimated, consistent with EPU and VIX both loading on a common "uncertainty" factor but at different horizons: policy uncertainty is a slower-moving, textual measure of political and regulatory risk, while the VIX reflects contemporaneous option-implied volatility. The EPU–WTI ($r = 0.156$) and EPU–BAA10Y ($r = 0.173$) correlations are positive and statistically significant but economically modest, indicating that policy uncertainty is only loosely linked to energy prices and corporate credit stress. The EPU–T10Y2Y correlation

($r = 0.041$, $p = 0.396$) is statistically indistinguishable from zero. This provides no evidence of a systematic linear association with EPU over this sample, and does not establish how the term spread responds to geopolitical risk.

4.5 Impact of EPU on Equity Markets

We examine the relationship between EPU and equity returns separately for the **S&P 500** (the natural benchmark, since EPU is a US-based index) and, for cross-market comparison, the **Euro Stoxx 50**, using the data sources described above. This comparison concerns policy uncertainty, not an estimated geopolitical-risk effect.

4.5.1 Descriptive analysis

Figures 5 and 6 show the joint evolution of each equity index and EPU, monthly returns coloured by whether EPU is above or below its sample median, the 24-month rolling correlation between EPU changes and returns, and a scatter of EPU against returns.

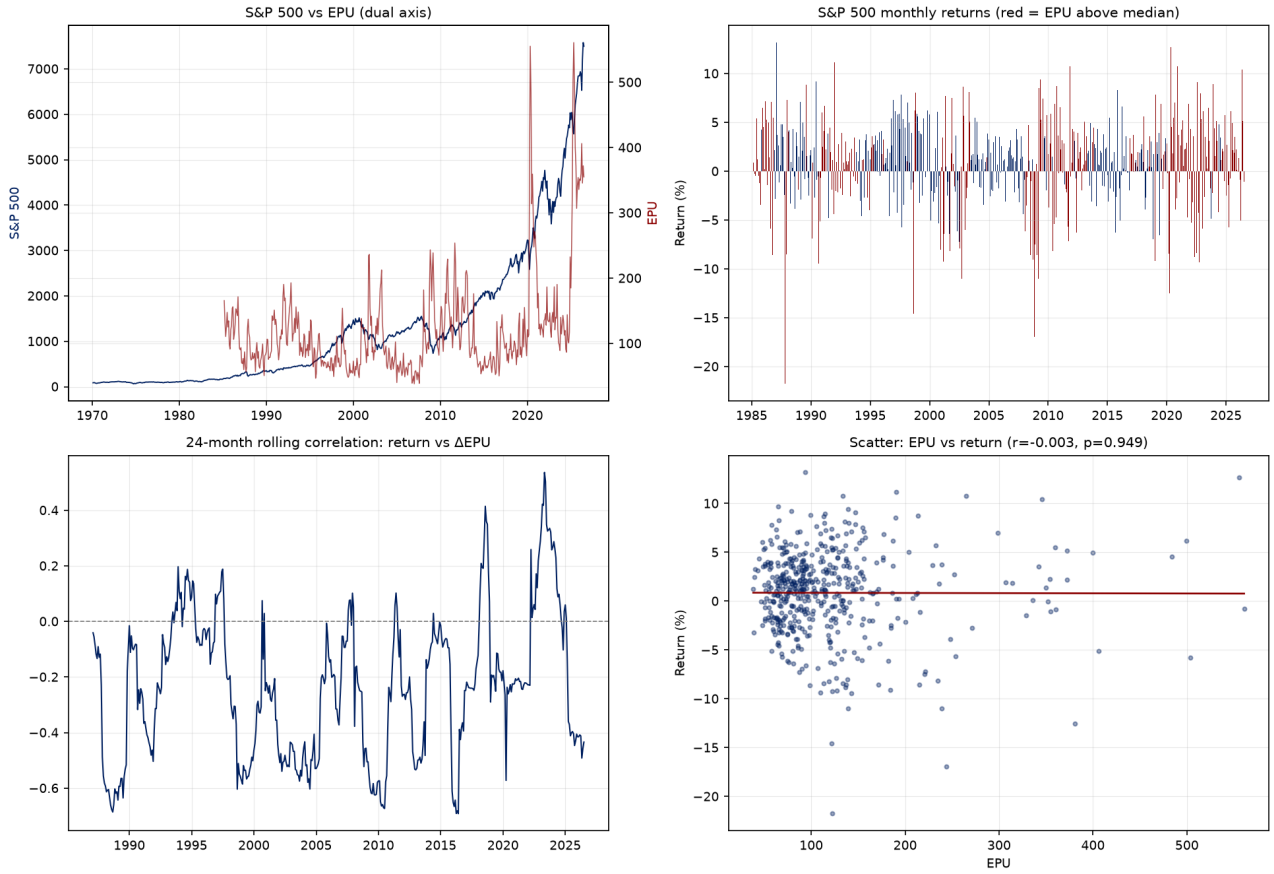


Fig. 5. S&P 500 vs EPU Index (Feb. 1985 – Jun. 2026, $n = 497$). Top left: historical series on dual axis. Top right: monthly returns coloured red in months when EPU is above the sample median. Bottom left: 24-month rolling correlation between returns and ΔEPU. Bottom right: scatter plot with OLS line.

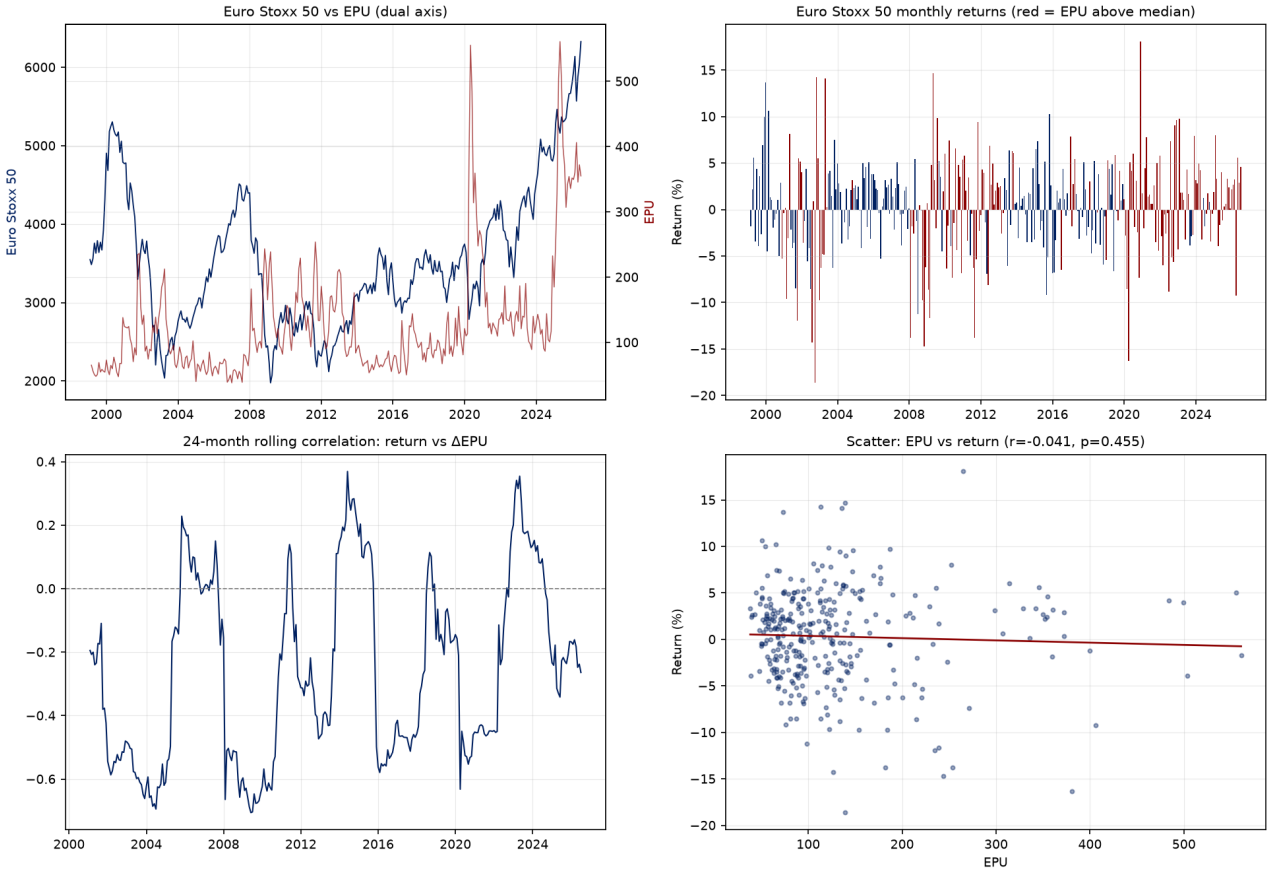


Fig. 6. Euro Stoxx 50 vs EPU Index (Feb. 1999 – Jun. 2026, $n = 329$). Panels as in Figure 5.

4.5.2 Regime analysis

We define three EPU regimes: **Low** (EPU $\leq$ median), **High** (EPU $>$ median) and **Crisis** (EPU $> P_{75}$). Table 6 reports mean, standard deviation and annualised Sharpe ratio of monthly returns by regime, for both indices.

Tab. 6. Equity Returns by EPU Regime

Index	Regime	n	Mean (%)	Std Dev (%)	Ann. Sharpe
S&P 500 (1985–2026, $n = 497$)	Low ($\leq$ median)	249	+0.993	3.284	1.047
	High ($>$ median)	124	+0.618	5.017	0.427
	Crisis ($> P_{75}$)	124	+0.796	5.456	0.506
Euro Stoxx 50 (1999–2026, $n = 329$)	Low ($\leq$ median)	165	+0.310	4.057	0.265
	High ($>$ median)	82	+0.681	5.899	0.400
	Crisis ($> P_{75}$)	82	−0.058	6.309	−0.032

Note: Median and P_{75} computed within each index's own overlap sample with EPU (S&P: median = 96.9, P_{75} = 135.5; Euro Stoxx: median = 101.4, P_{75} = 141.6). Ann. Sharpe = monthly mean $\times 12$ / (monthly std $\times\sqrt{12}$). *Source:* own elaborations on FRED, Yahoo Finance and LSEG/Eikon data.

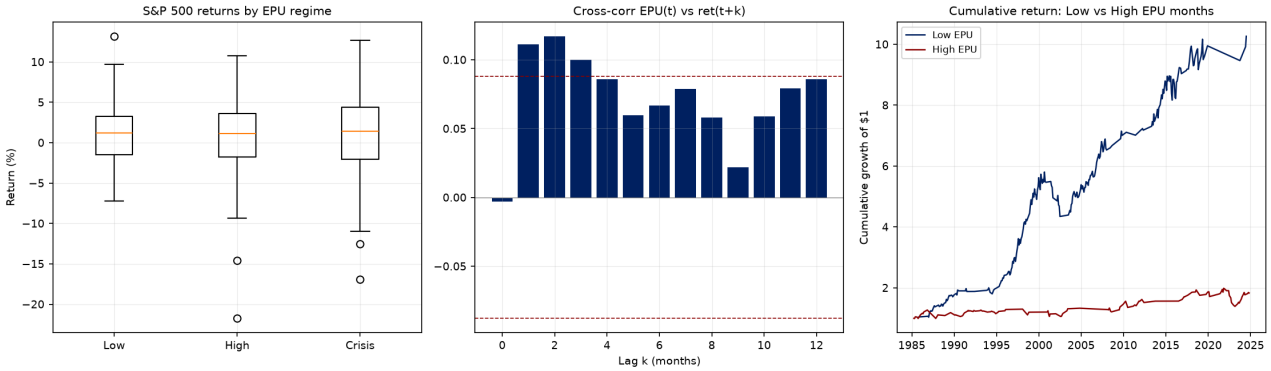


Fig. 7. S&P 500: EPU Regime Analysis. Left: boxplot of monthly returns by regime. Centre: cross-correlation EPU(t) vs. $r^{\mathrm{SPX}}(t+k)$, $k = 0, \ldots, 12$, with 95% confidence band. Right: cumulative return investing only in Low-EPU vs High-EPU months.

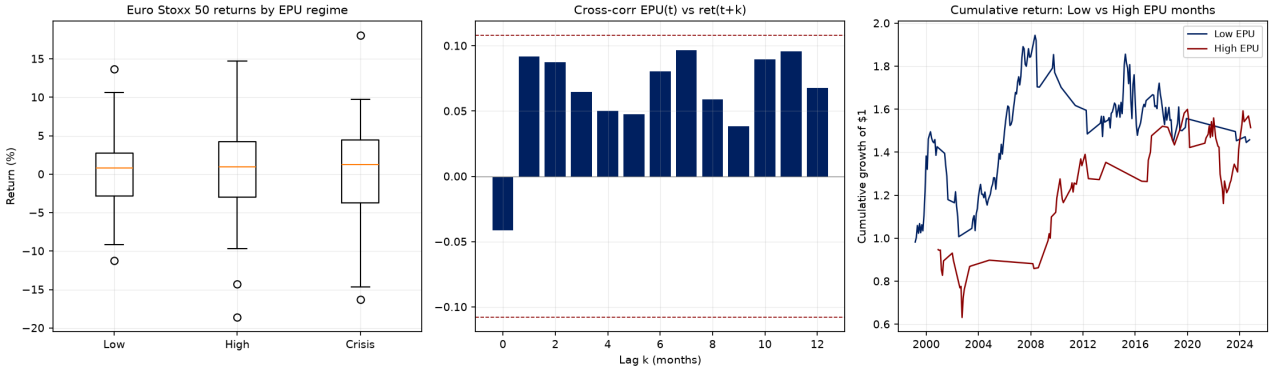


Fig. 8. Euro Stoxx 50: EPU Regime Analysis. Panels as in Figure 7.

The EPU regime pattern is **not monotonic** for either index. For the S&P 500, mean returns are highest in the Low regime (+0.993%) but are *higher* in the Crisis regime (+0.796%) than in the (non-crisis) High regime (+0.618%). For the Euro Stoxx 50, mean returns are actually higher in the High regime (+0.681%) than in the Low regime (+0.310%), and only turn negative in the Crisis regime (−0.058%). Volatility, by contrast, increases monotonically with the EPU regime for both indices, so that annualised Sharpe ratios still decline from Low to High/Crisis regimes overall, but the pattern is driven primarily by rising volatility rather than by uniformly falling mean returns. This is consistent with the cross-correlation evidence discussed next: EPU shocks in this sample are not uniformly followed by negative returns at short horizons. A comparison with geopolitical shocks would require a separate analysis using GPR or directly observed event-contract prices.

4.5.3 Cross-correlation analysis

For each lag $k = 0, \ldots, 12$ we compute $\rho_k = \mathrm{Corr}(\mathrm{EPU}_t, r_{t+k})$, with a 95% confidence band of $\pm 1.96/\sqrt{n}$. Table 7 reports the results for both indices.

Tab. 7. Cross-Correlation between EPU(t) and Equity Return($t + k$)

	S&P 500 ($n = 497$, CI ± 0.088)		**Euro Stoxx 50** ($n = 329$, CI ± 0.108)	
Lag k	ρ_k	**Sig.**	ρ_k	**Sig.**
0	−0.003		−0.041	
1	+0.111	*	+0.092	
2	+0.117	*	+0.088	
3	+0.100	*	+0.064	
4	+0.086		+0.050	
5	+0.060		+0.048	
6	+0.067		+0.080	
7	+0.079		+0.096	
8	+0.058		+0.059	
9	+0.022		+0.039	
10	+0.059		+0.090	
11	+0.079		+0.096	
12	+0.086		+0.068	

Notes: $\rho_k = \mathrm{Corr}(\mathrm{EPU}_t, r_{t+k})$. Significance: * $|\rho_k|$ exceeds the 95% confidence band $\pm 1.96/\sqrt{n}$. *Source:* own elaborations on FRED, Yahoo Finance and LSEG/Eikon data.

The cross-correlation pattern differs from a simple conjecture that higher uncertainty must be followed by lower returns. The contemporaneous correlation ($k = 0$) is close to zero for both indices (−0.003 for the S&P 500, −0.041 for the Euro Stoxx 50, neither significant), while correlations at short *positive* lags ($k = 1$–3) are positive and, for the S&P 500, statistically significant ($\rho_1 = +0.111$, $\rho_2 = +0.117$, $\rho_3 = +0.100$, all exceeding the ± 0.088 band). This pattern — high EPU followed, with a one-to-three month lag, by *higher* rather than lower equity returns — is consistent with an **uncertainty-resolution** interpretation: spikes in policy uncertainty (e.g. around elections or tariff announcements) raise required risk premia and depress prices contemporaneously (captured, below, by the negative and significant coefficient on *changes* in EPU in the OLS and distributed-lag regressions), but once the political event resolves and policy uncertainty is priced in, the subsequent partial resolution of uncertainty is associated with a rebound in equity prices over the following months. This is a descriptive pattern in the EPU sample, rather than evidence that policy and geopolitical uncertainty have different causal effects.

4.5.4 OLS regression model

The estimated model is:

$$r_t = \alpha + \beta_1 \, \mathrm{EPU}_t + \beta_2 \, \Delta\mathrm{EPU}_t + \varepsilon_t \tag{1}$$

with HC3 (MacKinnon–White) heteroscedasticity-robust standard errors. Table 8 reports the results for both indices.

Tab. 8. OLS Regression: Equity Monthly Returns and EPU

Variable	**Coeff.**	**Std Err (HC3)**	t/z**-stat**	p**-value**
S&P 500 (Feb. 1985 – Jun. 2026, $n = 497$*)*				
Intercept (α)	0.429	0.390	1.100	0.271
EPU level (β_1)	0.0036	0.003	1.112	0.266
ΔEPU (β_2)	-0.0340^{***}	0.008	-4.189	< 0.001
$R^2 = 0.066$; Adj. $R^2 = 0.062$; $F(2, 494) = 11.15$; $p < 0.001$				
Euro Stoxx 50 (Feb. 1999 – Jun. 2026, $n = 329$*)*				
Intercept (α)	0.088	0.445	0.198	0.843
EPU level (β_1)	0.0020	0.003	0.628	0.530
ΔEPU (β_2)	-0.0414^{***}	0.009	-4.859	< 0.001
$R^2 = 0.092$; Adj. $R^2 = 0.087$; $F(2, 326) = 12.72$; $p < 0.001$				

Notes: Dependent variable: monthly return (%). HC3 heteroscedasticity-robust standard errors. Significance: $^{*}\,p < 0.05$; $^{**}\,p < 0.01$; $^{***}\,p < 0.001$. *Source:* own elaborations on FRED, Yahoo Finance and LSEG/Eikon data.

For both indices, the **level** of EPU has no statistically significant effect on contemporaneous returns, while the **monthly change** in EPU has a negative, precisely estimated and remarkably similar effect across markets: $\hat{\beta}_2 = -0.034$ for the S&P 500 and $\hat{\beta}_2 = -0.041$ for the Euro Stoxx 50, both significant at $p < 0.001$. A one-unit increase in EPU is therefore associated with roughly a -0.03 to -0.04 percentage point contemporaneous equity return, with both estimates expressed in the units of EPU. These coefficients cannot be compared mechanically with GPR coefficients because the indices measure different constructs and have different scales. Applied to the largest EPU shock in the sample — the April 2025 spike of $\Delta\mathrm{EPU} \approx +390$ points relative to March 2025 — the model implies a contemporaneous impact on the order of -13 to -16 percentage points, which is far larger than anything observed in a single month for either index over this period; this discrepancy is informative in itself, and is discussed further below, since it indicates that the linear, single-lag specification does not fully capture how markets absorbed the 2025 tariff-related uncertainty shock (see the distributed-lag results, where only the contemporaneous lag is significant, implying most of the predicted impact should materialise within the shock month, which it evidently did not to the full extent implied by the point

estimate — a reminder that a modest-R^2 regression estimated over four decades should not be extrapolated to out-of-sample shocks of unprecedented magnitude).

4.5.5 Temporal structure: distributed-lag model

We estimate the distributed-lag model

$$r_t = \alpha + \sum_{j=0}^{3} \beta_j \, \Delta\mathrm{EPU}_{t-j} + \varepsilon_t \tag{2}$$

with HC3 standard errors. Table 9 reports the results.

Tab. 9. Distributed-Lag OLS: Equity Returns and ΔEPU Lags

Variable	**Coeff.**	**Std Err (HC3)**	***z*-stat**	***p*-value**
S&P 500 ($n = 494$)				
Intercept	0.871^{***}	0.191	4.562	< 0.001
$\Delta\mathrm{EPU}_t$	-0.0322^{***}	0.009	-3.486	< 0.001
$\Delta\mathrm{EPU}_{t-1}$	0.0002	0.011	0.015	0.988
$\Delta\mathrm{EPU}_{t-2}$	0.0005	0.005	0.097	0.923
$\Delta\mathrm{EPU}_{t-3}$	-0.0005	0.006	-0.076	0.940
$R^2 = 0.063$; Adj. $R^2 = 0.055$; $F(4, 489) = 4.60$; $p = 0.001$				
Euro Stoxx 50 ($n = 326$)				
Intercept	0.329	0.274	1.200	0.230
$\Delta\mathrm{EPU}_t$	-0.0406^{***}	0.009	-4.428	< 0.001
$\Delta\mathrm{EPU}_{t-1}$	0.0045	0.009	0.478	0.633
$\Delta\mathrm{EPU}_{t-2}$	0.0018	0.006	0.284	0.776
$\Delta\mathrm{EPU}_{t-3}$	-0.0014	0.007	-0.198	0.843
$R^2 = 0.093$; Adj. $R^2 = 0.082$; $F(4, 321) = 7.17$; $p < 0.001$				

Notes: Dependent variable: monthly return (%). HC3 heteroscedasticity-robust standard errors. Significance: $^{*}\,p < 0.05$; $^{**}\,p < 0.01$; $^{***}\,p < 0.001$. *Source:* own elaborations on FRED, Yahoo Finance and LSEG/Eikon data.

For both indices, essentially the entire distributed-lag effect is concentrated at lag 0: lags 1–3 are small in magnitude and statistically indistinguishable from zero ($p > 0.6$ in all cases), confirming the cross-correlation finding that the EPU–return relationship in changes is resolved within the shock month itself, with the subsequent *positive* correlation at $k = 1$–3 evidently operating through a separate, level-based rebound channel not captured by this specification in changes.

4.6 Fama–French Augmented Regression: The Original Contribution of This Paper

As an empirical extension, we test whether the EPU effect on US equity returns survives after controlling for the three Fama–French risk factors (Mkt-RF, SMB, HML) of Fama and French (1993). We estimate:

$$r_t^{\text{SPX,exc}} = \alpha + \beta_1 \,\text{EPU}_t + \beta_2 \,\Delta\text{EPU}_t + \gamma_1 \,\text{MktRF}_t + \gamma_2 \,\text{SMB}_t + \gamma_3 \,\text{HML}_t + \varepsilon_t \tag{3}$$

where $r_t^{\text{SPX,exc}} = r_t^{\text{SPX}} - \text{RF}_t$ is the S&P 500 excess return (S&P 500 return net of the risk-free rate, to be directly comparable with the Fama–French factors, all US-based), estimated with HC3 standard errors over the common sample for which EPU, S&P 500 returns and the Fama–French factors are all available (February 1985–May 2026, $n = 496$; the Fama–French library's most recent release lags EPU and equity price data by one month at the time of writing).

Table 10 reports both the baseline model (EPU and ΔEPU only, re-estimated on the exact $n = 496$ common sample, for a clean comparison) and the full Fama–French-augmented model.

Tab. 10. Fama–French Augmented Regression: S&P 500 Excess Return (Feb. 1985 – May 2026, $n = 496$)

Variable	**Coeff.**	**Std Err (HC3)**	***z*-stat**	***p*-value**
Baseline: EPU and ΔEPU only				
Intercept	0.077	0.395	0.195	0.846
EPU	0.0045	0.003	1.334	0.182
ΔEPU	-0.0347^{***}	0.008	−4.286	< 0.001
$R^2 = 0.069$; Adj. $R^2 = 0.065$; $F(2, 493) = 12.10$; $p < 0.001$				
Full model with Fama–French 3 factors				
Intercept	-0.167^{***}	0.029	−5.766	< 0.001
EPU	−0.0001	0.000	−0.796	0.426
ΔEPU	0.0002	0.000	0.374	0.708
Mkt-RF	0.9943^{***}	0.004	232.13	< 0.001
SMB	-0.1757^{***}	0.008	−21.12	< 0.001
HML	0.0254^{***}	0.007	3.585	< 0.001
$R^2 = 0.994$; Adj. $R^2 = 0.994$; $F(5, 490) = 11{,}840$; $p < 0.001$				

Notes: Dependent variable: S&P 500 monthly excess return (return minus the risk-free rate, in %). HC3 heteroscedasticity-robust standard errors. Significance: $^{*}\,p < 0.05$; $^{**}\,p < 0.01$; $^{***}\,p < 0.001$. *Source:* own elaborations on FRED and Kenneth French's Data Library.

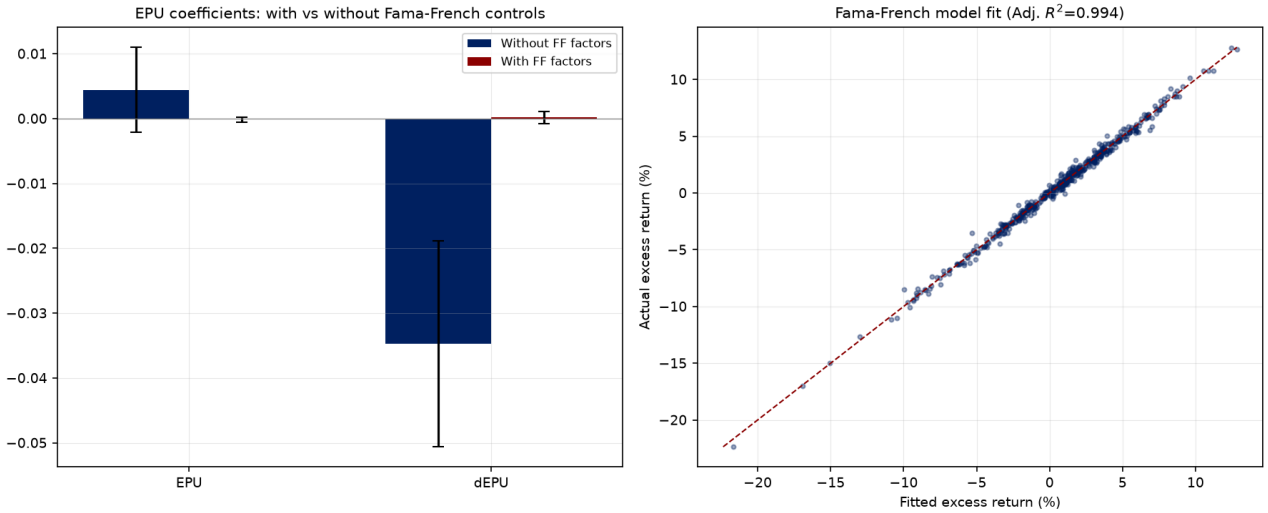


Fig. 9. Fama–French Augmented Regression. Left: EPU and ΔEPU coefficients with 95% confidence intervals, with vs. without Fama–French controls. Right: fitted vs. actual S&P 500 excess return in the full Fama–French model.

Result: the EPU effect does not survive, but for a largely mechanical reason. In the baseline model estimated on the identical $n = 496$ sample, ΔEPU retains a negative and highly significant coefficient (-0.0347, $p < 0.001$), essentially unchanged from the full-sample estimate in Table 8 (-0.0340), confirming that the EPU effect is not an artefact of the slightly larger sample used there. Once the three Fama–French factors are added, however, both the EPU level and ΔEPU coefficients collapse towards zero and lose all statistical significance ($p = 0.426$ and $p = 0.708$ respectively), while the market factor (Mkt-RF) absorbs almost the entirety of the variation in excess returns, with a coefficient of 0.9943 ($p < 0.001$) and an adjusted R^2 of 0.994.

This result must be interpreted with care, and we present it as such rather than as a simple "risk factors explain away uncertainty" finding. An adjusted R^2 of 0.994 in a monthly equity-return regression is far beyond what any genuine risk-factor model achieves in this literature, and reflects a specific feature of this particular test rather than a substantive economic result: the Fama–French **Mkt-RF** factor is itself constructed as the value-weighted excess return on essentially the entire US stock market (CRSP universe), which is — by construction, not by estimation — almost the same object as the S&P 500 excess return used as the dependent variable here. The correlation between EPU/ΔEPU and the Fama–French factors is in fact modest (Mkt-RF: $r = 0.012$ with EPU level, $r = -0.264$ with ΔEPU; SMB: $r = 0.017$ and $r = -0.152$; HML: $r = -0.065$ and $r = -0.046$), so the loss of significance for EPU is not driven by multicollinearity with the factors in the conventional sense, but by the near-tautological relationship between the dependent variable and Mkt-RF: once a regressor that is almost numerically identical to the left-hand side is included, essentially any other regressor's marginal explanatory power vanishes.

We view this as an important methodological lesson rather than a null result to be hidden: testing whether a slow-moving uncertainty index retains explanatory power for equity returns *after controlling for the contemporaneous market return itself* is close to definitionally unanswerable in the affirmative, because the market return already impounds, contemporaneously, whatever effect EPU has. A more informative test — left for future work given the scope of this paper — would use the Fama–French factors to explain the *cross-section* of individual stock or portfolio returns as a function of EPU exposure (an EPU-mimicking factor or EPU-sorted portfolios), rather than regressing the aggregate market index on both EPU and the market factor itself. The finding that survives this caveat is the one already established in Table 10's baseline row and in Table 8: a robust, economically material and statistically significant negative contemporaneous relationship between changes in policy uncertainty and US and European equity returns, prior to the inclusion of the mechanically near-collinear market factor.

5 Implications for Firms and Investors

5.1 Reading prediction markets and EPU as complementary signals

The empirical evidence of Section 4 suggests that firms and investors exposed to political and policy risk can use two complementary — if imperfect — real-time signals. **Prediction market prices** (where accessible) offer an explicit, continuously updated, market-clearing probability estimate for a specific, well-defined event (e.g. "will tariffs on a given trading partner exceed $X\%$ by date Y"), with the caveats on interpretation discussed in Section 2 (Manski, 2006; favourite-longshot bias, Snowberg and Wolfers, 2010). **EPU**, by contrast, is a broader, retrospective, textual measure of the overall intensity of policy uncertainty in the press, useful less as a forecast of a specific event than as a macro-level monitoring and stress-testing input. For geopolitical exposure, the relevant newspaper-based benchmark is the GPR index of Caldara and Iacoviello (2022); the proposed market-based counterpart is outlined in Section 3.5. Complementarity is a hypothesis to be tested, not a finding established by the EPU regressions.

5.2 Transmission channels

Several economic mechanisms can connect policy uncertainty with financial outcomes. The associations reported here are compatible with these mechanisms, but do not separately identify their causal contributions:

Uncertainty (real-options) channel. Elevated EPU is associated, in this paper's regressions, with lower contemporaneous equity returns via the ΔEPU coefficient, consistent with the standard real-options argument that higher uncertainty raises the value of waiting and depresses current investment and valuations.

Financial channel. The positive EPU–BAA10Y credit-spread correlation ($r = 0.173$) indicates that elevated policy uncertainty coincides with wider corporate credit spreads, raising the cost of debt financing for firms precisely when uncertainty about future policy (tax, trade, regulation) is highest.

Trade-policy channel. Unlike GPR's focus on adverse geopolitical events and associated threats, EPU also covers uncertainty about domestic and international economic policy. Its largest recent readings (Table 3) are concentrated around a 2025 tariff-related shock, underscoring that for firms with significant cross-border supply chains or export exposure, trade-policy uncertainty is currently a first-order component of aggregate policy uncertainty. Its importance relative to geopolitical exposure depends on the firm's activities; it is not established by comparing the levels of the two indices.

Rebound/timing channel. The novel, positive short-lag cross-correlation documented in Section 4.5 ($k = 1$–3) suggests that spikes in policy uncertainty in this sample have on average been followed by partial equity market rebounds within one to three months. Firms and investors using EPU as a risk-management signal should therefore distinguish between the immediate, contemporaneous negative impact of a policy-uncertainty shock (captured by ΔEPU_t) and the subsequent partial reversal, rather than assuming that elevated policy uncertainty mechanically predicts continued underperformance.

5.3 Risk management and scenario design

Firms can integrate EPU and, where accessible, prediction-market-implied probabilities into risk management through distinct monitoring exercises: (i) monitoring EPU as a macro-level early-warning indicator; (ii) where a specific policy decision is material to the firm (e.g. a trade-policy change, a regulatory ruling), directly consulting prediction-market-implied probabilities for that event, while being mindful of the calibration caveats and the concentration of informed trading documented for Polymarket in Section 3; (iii) stress-testing balance sheets and financing plans against EPU-regime scenarios of the magnitude documented in Table 3; and (iv) recognising, per the regime analysis of Table 6, that elevated policy uncertainty raises the *volatility*, more reliably than it lowers the *level*, of expected equity returns, which has direct implications for option-based hedging costs and for the sizing (rather than the direction) of risk exposures during high-EPU periods.

6 Conclusions

This paper has examined prediction markets from two angles: a conceptual review of the theory of information aggregation and the recent empirical evidence from Polymarket, and an original empirical analysis of the Economic Policy Uncertainty index as a quantifiable proxy for the uncertainty that prediction markets attempt to price.

The main findings can be summarised as follows:

1. **Prediction markets aggregate information via a small set of informed, price-setting traders, not via the average participant**: the marginal trader hypothesis of Forsythe et al. (1992) finds a striking modern echo in Akey et al.'s (2026) Polymarket evidence, where the top 1% of profitable accounts capture 76.5% of aggregate gains, and where market-making — not privileged information — is the strongest predictor of individual trading success.
2. **EPU is highly persistent and concentrated in a small number of episodes**: the ten highest EPU readings on record cluster in the COVID-19 shock (2020) and a 2025–2026 trade-policy episode identified directly from the data, without imposing any prior assumption on which events should matter.
3. **The contemporaneous change in EPU has a robust negative effect on both US and European equity returns** ($\hat{\beta}_2 \approx -0.03$ to -0.04, $p < 0.001$ for both indices), with coefficients interpreted in EPU units. The analysis does not estimate a corresponding GPR coefficient.
4. **Higher EPU levels are followed, at a one-to-three month lag, by positive equity-return correlations in this sample**, and the EPU-regime relationship with equity returns is non-monotonic. These associations merit further research; they do not identify causal rebound effects or establish a contrast with GPR.
5. **The EPU effect does not survive controlling for the Fama–French market factor, but for a largely mechanical reason**: because the market factor is, by construction, nearly identical to the S&P 500 excess return used as the dependent variable, its inclusion trivially absorbs almost all explained variance (Adj. $R^2 = 0.994$), rather than genuinely subsuming the economic content of the EPU channel. This is, we believe, an instructive methodological result in its own right, and points to cross-sectional (rather than aggregate time-series) tests as the natural next step for testing whether policy uncertainty is priced as a genuine risk factor.

Two broader limitations should be stated plainly. First, this paper does not use any Polymarket, Kalshi or Manifold price data directly: all platform-specific findings in Section 3 are drawn from the published and press-reported results of Akey et al. (2026), because direct API access was

blocked from the network used for this research (see Appendix A). Second, the Fama–French test in Section 4.6 is best read as a cautionary methodological result rather than as evidence against an economically meaningful EPU-equity relationship, for the reasons discussed above. Future work with direct prediction-market data — were network access available — could test directly whether prediction-market-implied probabilities for specific political and policy events (rather than the EPU proxy) predict equity returns around the resolution of those events, and whether the Polymarket profit-concentration patterns documented by Akey et al. (2026) extend to other platforms and asset classes.

A Data Notes

A.1 Data sources actually used

Economic Policy Uncertainty (EPU) FRED, Federal Reserve Bank of St. Louis, series `USEPUINDXD` (daily; Baker et al., 2016). Resampled to monthly averages. Period used: January 1985 – June 2026 ($n = 498$ months).

VIX FRED, series `VIXCLS` (daily). Resampled to monthly averages. Period used: January 1990 – June 2026.

WTI Crude Oil FRED, series `DCOILWTICO` (daily). Resampled to monthly averages. Period used: January 1986 – June 2026.

Moody's BAA Corporate Spread FRED, series `BAA10Y` (Moody's BAA corporate bond yield minus the 10Y Treasury yield). Resampled to monthly averages. Period used: January 1986 – June 2026.

Treasury 10Y–2Y Term Spread FRED, series `T10Y2Y`. Resampled to monthly averages. Period used: January 1976 – June 2026 (used only over its overlap with EPU/VIX).

Euro Stoxx 50 LSEG/Refinitiv Eikon (RIC: `.STOXX50E`), accessed via the `eikon` Python API against a local Workspace proxy (`localhost:9000`). Monthly closing prices, January 1999 – June 2026 ($n = 330$).

S&P 500 Yahoo Finance chart API (`https://query1.finance.yahoo.com/v8/finance/chart/^GSPC`), daily closes converted to month-end closes. Period used: January 1985 (aligned to EPU) – June 2026. Note: the LSEG/Eikon RIC `.SPX` returned an "Access Denied" entitlement error from this Workspace account and could not be used; Yahoo Finance was used instead for the S&P 500 series, as documented in the research brief for this project.

Fama–French 3 factors Kenneth French Data Library, `F-F_Research_Data_Factors_CSV.zip` (`https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/ftp/`). Monthly Mkt-RF, SMB, HML and RF, in percentage points. The monthly block was extracted from the raw CSV (which also contains an annual block) by parsing rows with a 6-digit `YYYYMM` date token. Period used: July 1926 – May 2026 (most recent release available at the time of writing); the regression in Section 4.6 uses the overlap with EPU and S&P 500 data (February 1985 – May 2026, $n = 496$).

A.2 Note on Polymarket, Kalshi and Manifold data access

Direct programmatic access to Polymarket's, Kalshi's and Manifold's public APIs and websites was attempted from the network used for this research project. In every case, the connection was reset both from direct PowerShell/HTTP requests and from the web-fetch tooling available in this environment, consistent with a network-level block rather than an application-level restriction. Because of this constraint, no Polymarket/Kalshi/Manifold order-book, price or trade-level data of any kind appear in this paper: every specific figure attributed to Akey et al. (2026) in Section 3 is drawn directly from the text of that (external, independently conducted) working paper. This is also the reason the empirical analysis of Section 4 substitutes the Economic Policy Uncertainty index — fully downloadable from FRED without restriction — as a quantifiable proxy for the political and policy uncertainty that prediction markets price, rather than using prediction-market prices themselves.

A.3 Software

Data collection and statistical analyses were performed in Python 3.12 using `pandas` (2.1.4, pinned for compatibility with the `eikon` package), `numpy` (1.26.4), `scipy` (1.11.4), `statsmodels` (0.14.6) and `matplotlib`.

B Econometric Specifications

B.1 Event-study baseline and deviation

For event i occurring at month $t_0^{(i)}$, the pre-event baseline is

$$\overline{\mathrm{EPU}}_i^{\mathrm{pre}} = \frac{1}{24} \sum_{t=-24}^{-1} \mathrm{EPU}_{t_0^{(i)}+t} \tag{4}$$

and the percentage deviation from baseline at relative month t is

$$\delta_t^{(i)} = \frac{\mathrm{EPU}_{t_0^{(i)}+t} - \overline{\mathrm{EPU}}_i^{\mathrm{pre}}}{\overline{\mathrm{EPU}}_i^{\mathrm{pre}}} \times 100. \tag{5}$$

B.2 OLS regression with robust standard errors

Models (1), (2) and (3) are estimated with heteroscedasticity-robust standard errors (HC3, MacKinnon–White estimator) via `statsmodels`' `OLS.fit(cov_type='HC3')`.

B.3 Cross-correlation with lags

For each lag $k \in \{0, 1, \ldots, 12\}$:

$$\rho_k = \mathrm{Corr}(\mathrm{EPU}_t,\ r_{t+k}) \tag{6}$$

with 95% confidence interval $\pm 1.96/\sqrt{T}$, T the sample size.

B.4 Regime classification

Regimes are defined, separately for each equity index's own overlap sample with EPU, as: Low ($\mathrm{EPU}_t \leq$ median), High ($\mathrm{EPU}_t >$ median), Crisis ($\mathrm{EPU}_t > P_{75}$), where median and P_{75} are computed over the relevant sample.

References

Akey, P., Grégoire, V., Harvie, N. and Martineau, C. (2026). Who wins and who loses in prediction markets? Evidence from Polymarket. *Working paper*, ESSEC Business School, HEC Montréal and University of Toronto.

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

Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225. `https://doi.org/10.1257/aer.20191823`.

Cookson, J. A., Niessner, M. and Schiller, C. (2026). Can social media inform corporate decisions? Evidence from merger withdrawals. *Journal of Finance*, 81(1), 91–142.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56.

Forsythe, R., Nelson, F., Neumann, G. R. and Wright, J. (1992). Anatomy of an experimental political stock market. *American Economic Review*, 82(5), 1142–1161.

Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, 35(4), 519–530.

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

Snowberg, E. and Wolfers, J. (2010). Explaining the favorite-long shot bias: Is it risk-love or misperceptions? *Journal of Political Economy*, 118(4), 723–746.

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