How do beliefs about future events affect asset prices?

Evidence from Prediction Markets

Simeng Li Jiarui Wang

Imperial College London

17 June 2026

Abstract

How beliefs about future events—and the disagreement among them—shape asset prices is a central question in finance, but existing empirical measures of investor beliefs are indirect, low-frequency, or aggregated. We address this using the universe of transaction-level data from Polymarket, a money-backed prediction market, over June 2024 through March 2026. From the same transaction tape, in the common metric of probabilities, we construct two high-frequency primitives: a price-entropy measure of consensus *uncertainty* and a dollar-volume-weighted measure of investor *disagreement*. We establish three results. First, the prediction market and the U.S. equity market are informationally integrated in both directions: stocks lead Polymarket within the trading day, while overnight Polymarket innovations predict the next morning's equity open. Second, disagreement positively predicts next-day stock returns—opposite in sign to analyst-forecast dispersion—consistent with compensation for unspanned macro state risk rather than mispricing. Third, consensus uncertainty predicts *lower* near-term equity volume and realised volatility, with the volatility released as uncertainty resolves, so it governs the *timing* of repricing rather than its level. A parsimonious heterogeneous-beliefs model rationalises the full pattern: disagreement is priced, while uncertainty is traded.

Keywords: Investor beliefs; Disagreement; Prediction markets; Polymarket; Stock returns.

1 Introduction

How investor beliefs and the disagreement among them shape asset prices is a central question in finance. Theory has long maintained that both matter — under short-sale constraints, the average belief is reflected in prices and disagreement tilts those prices toward optimists (Miller, 1977); under speculative resale motives, disagreement generates bubbles (Harrison and Kreps, 1978; Scheinkman and Xiong, 2003); under heterogeneous priors, it drives trading volume and momentum (Hong and Stein, 1999, 2003) and shapes equilibrium risk premia in general (Atmaz and Basak, 2018). Empirical work, however, has been constrained by a fundamental measurement problem. Neither the level of consensus belief nor the cross-sectional dispersion is directly observable in conventional financial data. Existing proxies — analyst forecasts, options, surveys, text sentiment — are *indirect*, *low-frequency*, or *aggregated*, and each faces interpretive objections that have produced decades of contested empirical findings.

This paper uses the universe of on-chain transactions from Polymarket, a money-backed prediction market on the Polygon blockchain, to recover these objects directly. Polymarket lists contracts that pay one dollar in specified states of the world; by no-arbitrage, the contract price is the marginal trader's risk-neutral probability of that state (Wolfers and Zitzewitz, 2004; Manski, 2004; Gjerstad, 2004). Each signed trade reveals one wallet's posterior about the same event, with real money at stake. We construct three primitives from this tape: (i) the contract price itself, a direct measure of the consensus *belief*; (ii) the entropy of the price-implied outcome distribution across the mutually exclusive states within an event, an *uncertainty* index Ψ that measures how unresolved the state distribution is; and (iii) the dollar-volume-weighted dispersion of signed wallet trades within a contract, a direct measure of the cross-sectional *disagreement* σ. All three are built from transaction-level data timestamped to the second, across a broad menu of political, economic, financial, and geopolitical events in the window 2024-06 through 2026-03. Critically, the primitives are extracted from the same data, in the same metric (probabilities in $[0,1]$), eliminating the apples-to-oranges problem that has dogged the existing literature.

Our main empirical results are threefold. We first establish that the belief object is informationally connected to the U.S. equity market, and that the connection runs in *both* directions. Within the trading day, stocks lead Polymarket: equity returns are rapidly impounded into the implied probabilities of the corresponding prediction markets, as one would expect if prediction-market traders watch the equity tape. In the overnight window, when U.S. equity markets are closed, the lead-lag reverses, and innovations in Polymarket prices *predict* the next morning's equity open on the underlying stocks. The two markets therefore draw on a common information flow. This two-way linkage is the headline of the section, and it validates the belief object: the Polymarket price moves with public information yet is not redundant with public prices. It is also the nexus on which the rest of the analysis rests. Only

because the two markets are informationally integrated can disagreement and uncertainty measured on Polymarket speak to equity returns and volatility.

We turn next to the stock return prediction. A one-standard-deviation rise in the dollar-volume-weighted disagreement index σ predicts a 17 basis point higher next-day return on the underlying stocks, with a positive and statistically significant Clark–West out-of-sample test statistic in our headline specification, and the political sub-index does most of the work. What makes the result compelling is that the predictability survives a demanding battery of controls. It is not absorbed by the Fama–French five factors, and not subsumed by the standard gauges of uncertainty and disagreement themselves, namely analyst forecast dispersion, the VIX, economic policy uncertainty (Baker et al., 2016a). It turns out that σ captures a distinct, priced source of information that existing predictors do not span. The sign is the opposite of the familiar *negative* relation between analyst forecast dispersion and returns (Diether et al., 2002; Miller, 1977). But because Polymarket disagreement concerns aggregate macro states rather than firm-specific cash flows, higher dispersion here is priced as compensation for aggregate state risk (Atmaz and Basak, 2018) rather than signalling the over-optimistic pricing that later mean-reverts.

We finally turn to the prediction of stock volatility, where the active object is consensus uncertainty Ψ. A high level of Ψ predicts *lower* next-day realised volatility on the underlying stocks. This is not merely a by-product of thin trading, since the volatility effect survives after controlling for contemporaneous volume. Thus Ψ carries independent information about *when* equity prices move. The sharper test is on the change in uncertainty: a fall in Ψ, the natural profile as an event nears resolution, predicts a subsequent *rise* in realised volatility. Volatility is released as the state distribution sharpens and the resolved state is finally incorporated into equity prices. We read Ψ as a measure of how much aggregate state risk is still unresolved. Because this risk is spread across many independent events, it is diversifiable and carries no return premium. What it governs instead is the *timing* of repricing: while the state stays unresolved, equity investors hold back, and the volatility is realised only once the uncertainty clears. This pattern is in the spirit of the literature on ambiguity aversion: when the relevant state is poorly resolved, a situation of Knightian uncertainty rather than mere risk (Ellsberg, 1961; Gilboa and Schmeidler, 1989; Fox and Tversky, 1995; Machina, 2014), ambiguity-averse investors optimally hold back and display portfolio inertia, so equity prices adjust sluggishly while the state is opaque and the deferred adjustment is realised as a burst of volatility once the ambiguity resolves (Epstein and Schneider, 2008; Illeditsch, 2011). Viewed through this lens, Ψ furnishes a direct, market-based proxy for the unresolved ambiguity that this literature has typically had to infer indirectly. The pattern is asymmetric across the two indices and holds out of sample: disagreement σ does not predict volatility, just as uncertainty Ψ does not predict returns.

To interpret these patterns we develop a parsimonious heterogeneous-beliefs framework, and three insights from it are worth previewing. First, risk aversion is the engine of the dis-

agreement premium: the compensation that risk-averse investors demand for disagreeing about unspanned macro states loads positively on returns, and it vanishes under log utility, so the effect is a genuine risk premium rather than a mispricing. Second, the object of disagreement must be *unbounded*, an aggregate growth rate rather than a single bounded event, for that premium to be unconditional and non-reversing, which is what separates our positive sign from the negative analyst-dispersion premium. Third, the two moments dissociate because uncertainty is spread across many independent events: consensus uncertainty is therefore diversifiable and earns no premium, yet it still drives each market's volatility and hence the timing of repricing. The upshot is that one standard risk-premium mechanism, with no frictions or mispricing, reproduces the entire pattern of estimates.

Our paper contributes to three strands of literature. First, it speaks to the broad literature on *beliefs and asset pricing*. Asset prices are conditional expectations taken under the marginal investor's beliefs (Lucas, 1978), so the level, dispersion, and resolution of those beliefs are first-order for prices; a recurring difficulty, recognised at least since the joint-hypothesis problem in Euler-equation tests (Hansen and Singleton, 1982), is that beliefs are not directly observed, which leaves preference-based and belief-based explanations of asset-pricing puzzles hard to separate. Survey evidence shows that investor expectations depart systematically from the rational benchmark (Greenwood and Shleifer, 2014), and a large theoretical literature shows that belief *heterogeneity* alone moves prices: under short-sale constraints disagreement tilts prices toward optimists and predicts reversal (Miller, 1977), speculative resale motives generate bubbles (Harrison and Kreps, 1978; Scheinkman and Xiong, 2003), and differences of opinion drive volume, momentum, and crashes (Hong and Stein, 1999, 2003; Banerjee and Kremer, 2010a); in equilibrium, disagreement is priced as a risk premium even without frictions (Atmaz and Basak, 2018). A parallel strand concerns aggregate *uncertainty*, from the Knightian distinction between measurable risk and unmeasurable uncertainty (Knight, 1921) to the real and financial effects of uncertainty shocks (Bloom, 2009). Whether dispersion is informative depends, in turn, on *who* disagrees: informed traders are compensated for processing costly information (Grossman and Stiglitz, 1980a; Kyle, 1985), whereas noise traders push prices away from fundamentals (Black, 1986; De Long et al., 1990). We contribute a single, direct, money-backed measurement of the consensus belief, its cross-sectional dispersion, and its residual uncertainty for the same events, and show that each maps to a distinct moment of equity returns.

Second, the paper contributes to the empirical literature that measures investor beliefs from observable data, every branch of which confronts a known limitation that our money-backed, event-level measure is designed to sidestep. *Surveys* elicit expectations directly but are low-frequency, rest on small panels, and put no capital at stake (Greenwood and Shleifer, 2014; Coibion and Gorodnichenko, 2015). *Trading volume* proxies for the disagreement that generates trade (Karpoff, 1987; Harris and Raviv, 1993a; Kandel and Pearson, 1995; Lo and Wang, 2000), yet conflates belief heterogeneity with liquidity, hedging, and rebalancing and

leaves the sign of disagreement unidentified. *Analyst forecast dispersion* is the workhorse cross-sectional proxy (Diether et al., 2002), but it is contaminated by idiosyncratic risk (Johnson, 2004), illiquidity (Sadka and Scherbina, 2007), and credit quality (Avramov et al., 2009), and reflects only a small, incentive-constrained slice of agents. *Option-implied* measures recover risk-neutral moments and disagreement (Whaley, 2000; Bakshi et al., 2003; Carr and Wu, 2009; Golez and Goyenko, 2022; Salome et al., 2024), but require modelling the pricing kernel and exist only where options trade. Three more recent families are closest to ours because they too exploit non-traditional, high-frequency data, and we discuss them in more detail. *News and text* indices score the tone and content of articles and official documents: media pessimism forecasts returns (Tetlock, 2007), text-based indices quantify economic-policy and geopolitical uncertainty (Baker et al., 2016b; Caldara and Iacoviello, 2022), and news-implied measures capture disaster and narrative risk (Manela and Moreira, 2017); but these gauge the attention of journalists rather than the beliefs of investors and are not anchored to capital at risk. *Social-media* measures mine message boards and platforms such as Twitter, StockTwits, and Seeking Alpha for sentiment and revealed positioning (Antweiler and Frank, 2004; Bollen et al., 2011; Chen et al., 2014), and have been used to build direct investor-disagreement indices (Cookson and Niessner, 2020; Cookson et al., 2026); yet posters are a self-selected, mostly retail subset, posts are cheap talk, and in episodes such as meme-stock runs the measure can drive flows rather than reveal beliefs. *Google search* volume proxies for retail attention (Da et al., 2011), including around earnings announcements (Drake et al., 2012) and as a predictor of volatility (Vlastakis and Markellos, 2012), and can separate retail from institutional attention (Ben-Rephael et al., 2017); but it captures attention rather than beliefs and is silent on direction. Against all of these, Polymarket trades are direct, money-backed, event-specific, and resolved to the individual wallet, which lets us recover the consensus belief, its dispersion, and its residual uncertainty in a common probability metric.

Third, the paper contributes to the literature on prediction markets. Classical work emphasises the forecasting power, information-aggregation properties, and interpretation of these markets (Wolfers and Zitzewitz, 2004, 2006; Snowberg et al., 2012; Manski, 2004), and more recent work uses them as event-study benchmarks (Snowberg et al., 2011). On the macroeconomic side, Diercks et al. (2026) study the rapid rise of Kalshi's regulated macro contracts and show that their prices deliver an accurate, high-frequency, real-time gauge of macroeconomic outcomes such as inflation prints and monetary-policy decisions, frequently rivalling or improving on survey and market-based forecasts; their object, however, is the prediction-market price as a *macroeconomic expectation*, a forecast of a single macro release. Turning to the microstructure of these markets, Gómez-Cram et al. (2026b) ask why prediction-market prices are so accurate and show that the accuracy is driven by a small, persistently skilled minority, roughly 3% of accounts, whose order flow predicts subsequent prices and final outcomes, while the crowd supplies most of the volume but little

of the information. This evidence is complementary to ours: by establishing that wallet-level heterogeneity is economically meaningful rather than incidental, it supports our use of transaction-level wallet data to measure disagreement and sharpens the question we go on to pursue, namely whose beliefs actually matter for asset prices.

Closest to our setting, Gómez-Cram et al. (2026a) use Polymarket's new firm-level earnings contracts to build a high-frequency, stake-backed measure of earnings expectations that is more accurate than the analyst consensus, approximately unbiased, incrementally informative for earnings-announcement returns, and leads analyst forecast revisions. Our paper differs in that we use Polymarket to recover investor beliefs, and especially their *uncertainty* and *disagreement*, as asset-pricing objects, and study how these objects are transmitted to external financial markets. Whereas their focus is the informativeness of Polymarket prices *within* prediction markets, we exploit the information embedded in Polymarket prices and wallet-level trading behaviour to study how revealed beliefs and disagreement reach broader asset markets.

The remainder of the paper proceeds as follows. Section 2 describes Polymarket and the data. Section 3 establishes Polymarket's relevance as a macro-financial laboratory. Section 4 constructs the uncertainty index Ψ and the disagreement index σ, then validates each against internal Polymarket price formation. Section 5 establishes the two-way informational connection between Polymarket and the U.S. equity market. This connection is the nexus on which the rest of the analysis rests. Section 6 is the headline empirical chapter. It shows how disagreement maps into next-day stock returns and how uncertainty governs the timing of equity volatility and volume. Section 7 develops a heterogeneous-beliefs framework that rationalises the full pattern of estimates, organised around a single thesis: disagreement is priced, while uncertainty is traded. Section 8 concludes.

2 Polymarket

This section describes the institutional features of Polymarket that matter for the measurement strategy that follows. The reader who is already familiar with prediction-market mechanics can skim Sections 2.1 and 2.2 and proceed directly to the description of the trading data in Section 2.4.

There are three features of Polymarket that distinguish it from the prediction markets studied in earlier work, such as the Iowa Electronic Markets or Intrade. *First,* Polymarket is a blockchain-native market: every order match and every fill is recorded on Polygon in a public ledger that is immutable and downloadable without intermediaries. *Second,* traders are identified by persistent wallet addresses, which lets us follow individual trading behaviour over time without the platform's cooperation. *Third,* Polymarket runs on a fully-collateralised binary-contract design — every YES share is matched by a NO share so that the unit always pays one dollar at resolution — which gives the contract price an unambiguous

probability interpretation that does not depend on counterparty risk or platform mark-ups.

2.1 A simple contract

It is easiest to begin with a concrete contract. Suppose the market asks: *Will the Federal Reserve cut rates at the January meeting?* A *Yes* share pays \$1 if the event occurs and \$0 otherwise; a *No* share pays \$1 if the event does not occur and \$0 otherwise. The two shares are complementary by construction: whichever outcome realises, the holder of the winning share receives \$1 and the holder of the losing share receives nothing. Together they cover the unit.

This structure has a direct implication that will be central to everything that follows. If the *Yes* share trades at a price $p \in [0,1]$, then by no-arbitrage the *No* share must trade at $1-p$: any other configuration leaves a riskless profit on the table. A trader who is willing to buy *Yes* at p is therefore revealing, by capital commitment, that she values the future *Yes* payoff at least at p. Equivalently, she is revealing that she believes the probability of *Yes* is at least p. The trade price is, up to a small liquidity wedge, the marginal trader's risk-neutral probability of the event. This is the sense in which a Polymarket price is a *direct* measure of a belief: the interpretation does not require a structural model of risk premia, a calibration of a stochastic discount factor, or any auxiliary assumption beyond the absence of arbitrage between the two complementary shares.

The same logic implies that every trade reveals disagreement in the most literal sense. The buyer of *Yes* at p believes the probability is at least p; the buyer of *No* at $1-p$ believes the probability is at most p. The two are on opposite sides of the same future state, with real money at stake on each side. Section 4.2 builds on this observation to construct a wallet-level measure of cross-sectional disagreement; here, the point is simply that the data already contain, at the level of every single fill, the raw material from which both belief and disagreement measurements can be recovered.

2.2 Binary markets and grouped event markets

Polymarket supports two kinds of contract structures, and the distinction matters for how we measure disagreement. A *binary market* is a single *Yes*/*No* contract on a well-defined future state — for example, *Will Donald Trump be inaugurated as President on January 20, 2025?* The two prices sum to approximately one, the object of disagreement is a single state, and the disagreement structure is captured entirely by the cross-section of who is buying which side at what price. We use binary markets as the building block for the dispersion measure $\sigma_{m,t}$ in Section 4.2.

Many other interesting events, however, are represented by *grouped markets* that partition the event into several mutually exclusive outcomes. The Federal Reserve example illustrates the structure. A single economic event — the January FOMC decision — can generate

a family of related contracts: *25+ bps increase*, *no change*, *25 bps decrease*, *50 bps decrease*, and *75+ bps decrease*. Each contract within the group is itself a binary *Yes*/*No*, but the contracts share the same underlying event and their *Yes* prices must sum to approximately one by no-arbitrage.

This grouped structure introduces a second dimension of disagreement. Within a given contract — say, the *no change* sub-market — traders disagree about whether that specific outcome will realise. Across contracts within the same event, the cross-section of prices itself reveals how diffuse the market's beliefs are: a probability vector that is concentrated on one outcome (e.g., 0.90 on *no change*, 0.025 on each alternative) reflects a tight consensus; a probability vector that is spread evenly (e.g., 0.20 on each of five outcomes) reflects a wide-open disagreement about which state will realise. In Section 4.1 we use exactly this idea to construct Ψ_t, the entropy of the price-implied outcome distribution within each event (a grouped event's contracts, or the two outcomes of a binary market). The two measures — Ψ from prices across contracts within a grouped event, σ from wallet-level positions within a single contract — are complementary: each captures a piece of disagreement that the other misses.

A practical note: because grouped events are sometimes split into many sub-markets (with new sub-markets occasionally added as additional candidate outcomes become relevant), we treat the event, not the individual sub-market, as the unit of analysis when computing Ψ. The construction is described in detail in Section 4.1; here, we simply flag that the existence of grouped markets is one of the institutional features that lets us measure disagreement at all.

2.3 Who creates markets and how are they resolved?

Markets are created by Polymarket's internal markets team, and users can suggest new markets through community channels but cannot list arbitrary contracts themselves. Each listed market comes with pre-specified resolution rules that identify the source used to determine the outcome, the date on which the market becomes eligible for resolution, and the treatment of edge cases. These rules are not cosmetic: the market title is only shorthand for the contract, and the formal resolution criteria determine how the contract will ultimately settle. The practical consequence is that traders are not betting on vague narratives. They are trading on contracts tied to explicit, verifiable resolution criteria, so that diffuse beliefs about the future are translated into prices on sharply defined state-contingent claims — exactly the object that the disagreement theories we set out in the introduction require.

Resolution itself is handled not by Polymarket but by an external verification mechanism, UMA's Optimistic Oracle, which operates as a proposer–disputer system. When the underlying event becomes known, a participant proposes an answer to the market's resolution question by consulting the source named in the resolution rules — an official Federal

Reserve statement, an election certification, a court docket — and posts a collateral with the proposal. The system is "optimistic" in the sense that the proposed answer is treated as correct unless it is challenged during a specified dispute window. If no challenge arrives, the proposal stands, the market resolves on that basis, and the proposer recovers the collateral. If the proposal is disputed, the case escalates and UMA tokenholders vote on the correct outcome; the losing side forfeits its collateral, which is transferred to the winning side. The economic incentives on both sides — to propose carefully, to challenge incorrect proposals — are what makes the final resolution credible.

This resolution structure matters for our paper because it removes a specific kind of platform opacity. Contracts are not settled by discretionary judgment from Polymarket or any other single party after the fact; each market is tied to ex ante resolution rules, and disputed outcomes are handled through a formal challenge process in which participants have economic incentives to propose and defend the correct result. It also matters because every wallet observes the same resolution outcome at the same time: there is a single, unambiguous, publicly observable realisation $y \in \{0,1\}$ for each resolved market, and that realisation is the same for every wallet that ever traded the contract. This means the disagreement we measure prior to resolution is anchored to a well-defined common state, not to subjective interpretations of an ambiguous outcome.

2.4 Price formation and trading

The market microstructure that produces our prices is a hybrid central limit order book, with off-chain order matching and on-chain settlement. A *maker* supplies liquidity by posting a resting limit order to the book. A *taker* removes liquidity by submitting an order that executes against the best available quote. Suppose, for example, that a maker has posted a resting bid to buy *Yes* at 0.60. The order can be completed either by a counterparty willing to sell an existing *Yes* share at 0.60 or, equivalently, by a counterparty willing to take the complementary *No* side at 0.40 — the two are economically the same transaction.

The off-chain matching engine is what makes the market practical for active trading. Quote revisions and order placements do not have to be written to the Polygon blockchain in real time, which avoids the gas fees and block-confirmation latency that would otherwise prevent market makers from posting tight spreads. Quotes update at sub-second frequency, prices respond to new information rapidly, and the order book is deep enough on the most actively traded markets to support short-horizon price discovery. Once a trade is matched, however, settlement is finalised on-chain and the completed transaction is recorded on Polygon. This hybrid design preserves the transparency and auditability that motivate the use of blockchain data in the first place while keeping trading fast enough to support an active order-driven market. It also means that trading is non-custodial: Polymarket organises trading and matches orders but is never the counterparty, and so cannot unilaterally execute

unauthorised transactions or appropriate customer assets. Trading is also effectively fee-free over our sample period, which further supports active participation and frequent price updating.

For our purposes the key consequence is that every executed trade leaves a permanent, machine-readable record on the Polygon ledger. For each fill we observe the block number and timestamp, the wallet addresses of the maker and the taker, the asset identifiers of the *Yes* and *No* tokens involved (which we link to market-level metadata to recover the question text, event category, contract structure, and resolution outcome), the number of shares exchanged, the per-share transaction price, and the implied dollar value of the transfer. The wallet addresses are persistent across markets and over time, which allows us to follow individual trading behaviour without the platform's cooperation. This is the feature that makes the disagreement measure σ in Section 4.2 possible at all.

Our analysis sample is constructed from this transaction tape. We collect the `OrderMatched` stream from Polygon blocks ~59M through ~84M, corresponding to the window 2024-06-10 through 2026-03-10, and merge in market-level metadata from the Polymarket API. Across the entire on-chain ledger in this window we observe 656,848,877 executed trades involving 2,257,528 distinct wallet addresses and roughly $24.6 billion in cumulative dollar volume. After restricting to the political, economic, geopolitical, and financial sub-sample that drives most of our analysis (Section 3), we are left with 127,950,212 trades, 1,639,425 wallets, and $15.26 billion in cumulative dollar volume. A more detailed breakdown is in Section 3.1; the headline point is that the data are big enough, broad enough, and granular enough to support the measurement strategy that follows.

3 Polymarket as a Macro-Financial Laboratory

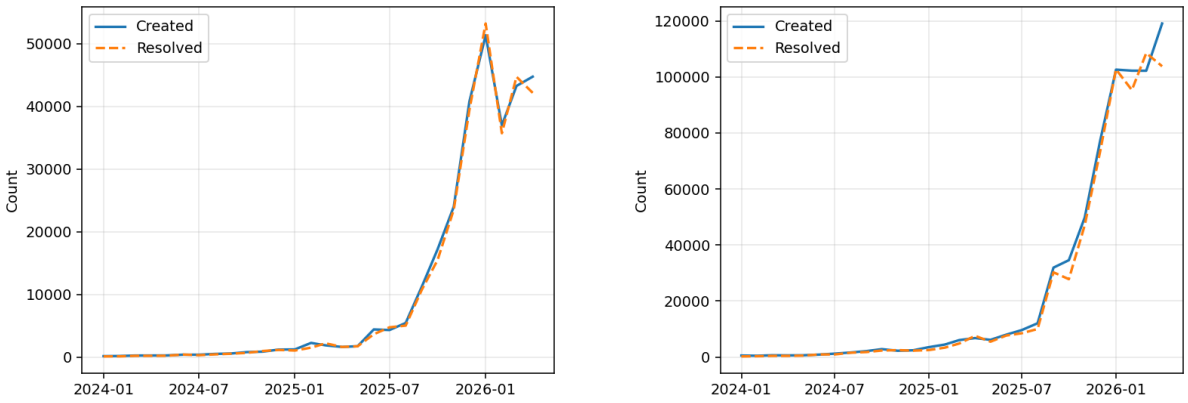
Whether Polymarket is a serious enough venue to be useful for asset-pricing measurement turns on three empirical questions: is it big, is it broad, and is it accurate? We address each in turn. The answers are: yes, increasingly so; yes, across all major macro and political event categories; and yes, both relative to alternative forecasting benchmarks and relative to the eventual realisations.

3.1 Scope and economic relevance

The first question a finance reader will ask about prediction-market data is whether the markets are about anything that matters. If Polymarket were dominated by sports betting and entertainment contracts, it would be a poor venue from which to draw inferences about investor beliefs over macro-financially relevant states. We address this concern directly.

Contracts on Polymarket are written on sharply defined events with binary or discrete outcomes, which allows prices to be interpreted as risk-neutral probabilities for the reasons

set out in Section 2.1. The platform hosts a broad range of events, spanning politics, macroeconomics, cryptocurrencies, sports and weather among other domains. Figure 1 documents that market creation and resolution proceed at a steady pace throughout our sample: in a typical month, several hundred new events and several hundred new markets are introduced, with a comparable number resolving as their underlying events realise. The sample period as a whole spans roughly 748,000 markets across 323,000 distinct events. This flow structure indicates a continuously active marketplace rather than a static collection of stale contracts.



(a) Events

(b) Markets

Figure 1: Monthly counts of newly created and newly resolved events (Panel a) and individual binary markets (Panel b) on Polymarket over June 2024–March 2026. In each panel the two series report, respectively, the number of contracts whose trading opens and the number whose underlying event realises and settles within the month. Creation and resolution proceed at a steady pace throughout the sample—several hundred new events and markets per month, with a comparable number resolving—indicating a continuously active marketplace rather than a static collection of stale contracts. The sample as a whole spans roughly 748,000 markets across 323,000 distinct events.

For the remainder of the paper we restrict attention to what we will call the *macro-financial* sample, defined as the union of Polymarket's primary-category labels POLITICS, ECONOMY, FINANCE, and GEOPOLITICS. The political category captures electoral outcomes, institutional transitions, court rulings, and policy decisions; the economic category captures macroeconomic data releases, monetary-policy decisions, and central-bank communications; the finance category captures markets directly written on stock prices, equity-index levels, and corporate events; and the geopolitics category captures contracts on cross-border conflicts and sanctions whose outcomes routinely move global asset prices. The restriction excludes cryptocurrency, sports, weather, esports, and entertainment — domains in which disagreement is interesting in its own right but whose connection to the asset-pricing literature is indirect. After this filter we are left with roughly 130,000 markets across the four categories, comprising 108,090 political, 2,855 economic, 15,179 finance, and 3,642 geopolitical markets. Cumulative reported notional volume on the macro-financial sample is approximately

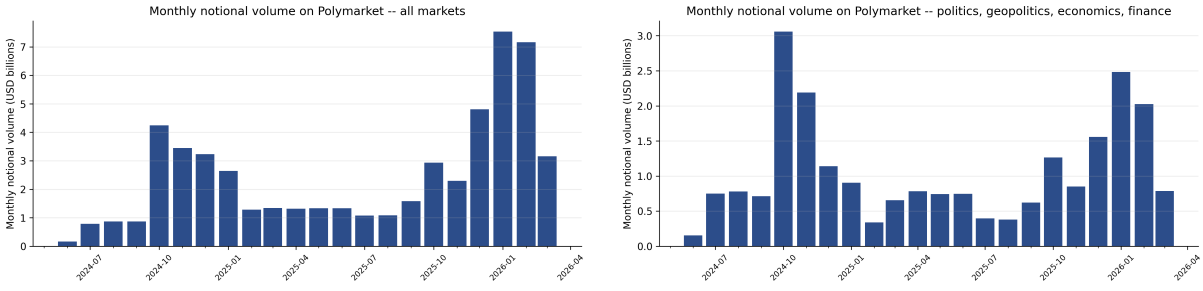
$23.9 billion across the sample window, with politics and economy together accounting for the large majority and geopolitics and finance making up the remainder.

Even on this richer sample, the dataset is large and economically deep. The sample window covers periods of substantial macro-financial importance, including the 2024 U.S. presidential election cycle and its aftermath, an unusually active sequence of Federal Open Market Committee meetings, and several episodes of heightened trade-policy uncertainty driven by tariff announcements. These features ensure that the disagreement we measure is about events whose realisations carry real consequences for asset prices — and so that the predictive content we will document in Section 6 is built on a foundation of economically meaningful underlying states.

3.2 Market participation and economic significance

The second question is whether participation in these markets is deep enough to take their prices seriously. A thin or unrepresentative market would give us prices, but those prices would not reflect any serious aggregation of beliefs and would not be informative about the broader investor population. The concern is non-trivial: many earlier prediction-market venues, including the Iowa Electronic Markets, were subject to strict position limits that capped participation at a few hundred dollars per trader.

The data on Polymarket strongly reject the thin-participation hypothesis. In our matched-order tape, cumulative notional volume across all contract categories totals approximately $58.8 billion over the sample window. Restricting to politics, geopolitics, economics, and finance markets, the sample contains 1,675,827 distinct wallets and roughly $23.9 billion in notional volume. This equals about $1.04 billion per month on average, with a peak monthly volume of about $3.06 billion in October 2024, immediately before a major political catalyst. On a typical day, the platform supports tens of thousands of unique wallets interacting across many thousands of active markets. Figure 2 plots the two monthly notional-volume series side by side: the all-markets series on the left and the politics–geopolitics–economics–finance sub-sample on the right. The two series share the same broad shape but differ in scale and composition: the all-markets series also includes high-volume categories, such as sports, that we exclude from the empirical chapters. The sub-sample in panel (b), while smaller in total, is precisely the universe whose beliefs are likely to be relevant for the asset-pricing questions we ask. These magnitudes point to sustained, large-scale participation with meaningful capital at risk, and to a market structure that supports high-frequency information aggregation rather than sporadic speculative betting.



(a) All markets.

(b) Macro finance markets subsample.

Figure 2: Monthly notional trading volume on Polymarket, June 2024 through March 2026. Notional volume is measured in face-value dollars of binary contracts traded (size summed over fills) — the same measure that Polymarket's API and third-party trackers such as The Block report. Panel (a) covers all contract categories; panel (b) restricts to politics, geopolitics, economics, and finance markets. Same time axis on both panels; note the difference in vertical scale.

Beyond the aggregate scale, the cross-section of wallet activity is itself informative. Table 1 reports, for wallets active in the macro-financial sample, the cross-sectional distribution of four quantities: total notional volume traded, average notional size per trade, the number of distinct tokens the wallet trades, and the wallet's realised dollar profit-and-loss on markets that have resolved by the end of our sample. The distribution is heavily right-skewed — the median wallet trades roughly $856 of notional, while the 99th-percentile wallet trades about $164,756. The average trade size is similarly dispersed: the median wallet places trades of about $85, while the top one percent of wallets place trades averaging about $22,003. The realised-payoff distribution is centered close to zero for the median wallet, with a median payoff of $−0.80, but the cross-section is wide: the 10th percentile loses about $103.62, whereas the 99th percentile earns about $1,682.53. The average wallet loses a small amount overall (mean realised payoff $−34.85), consistent with the bid–ask spread and trading costs that a marginal participant pays. This wide heterogeneity across wallets in trade size, breadth, and realised payoff is the raw material on which the disagreement measure of Section 4.2 operates. Figure 3 shows the three distributions graphically.

	N	Mean	p10	p25	p50	p75	p90	p99
Total notional ($)	1,675,827	14,247	18.32	120.11	856.38	4,398	15,423	164,756
Avg trade size ($)	1,675,827	1,100	5.16	19.70	85.41	380.25	1,374.50	22,003
# trades	1,675,827	34.59	1.00	3.00	7.00	20.00	45.00	279.00
# distinct tokens	1,675,827	19.66	1.00	2.00	6.00	15.00	33.00	175.00
Realised P&L ($)	1,588,986	-34.85	-103.62	-13.51	-0.80	1.55	30.04	1,682.53

Table 1: Wallet-level summary statistics on the macro-financial sub-sample. Reports the cross-sectional distribution across all wallets that traded at least once in the macro-financial sample of (a) total notional volume traded, (b) average notional size per trade, (c) the number of trades, (d) the number of distinct tokens the wallet traded, and (e) the wallet's realised dollar payoff on markets that have resolved by the end of the sample window. The macro-financial sample is the union of the POLITICS, ECONOMY, FINANCE, and GEOPOLITICS categories. The payoff row excludes wallets with no resolved-market activity. Notional is measured in face-value dollars of binary contracts traded.

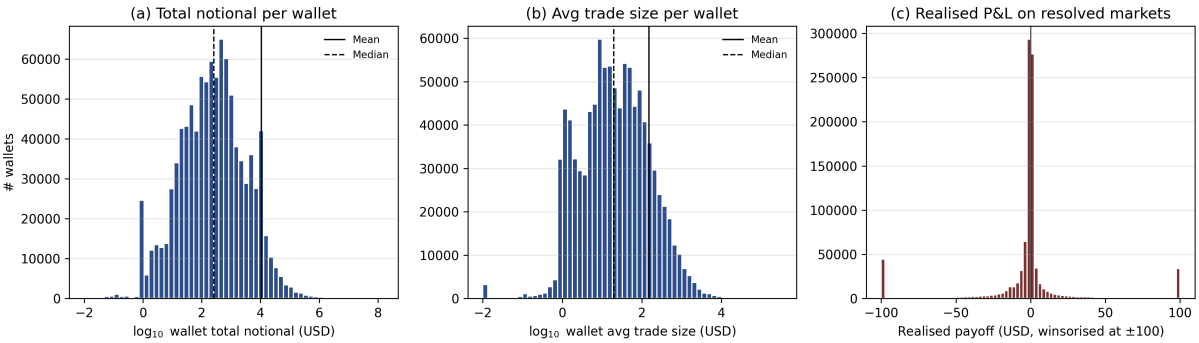


Figure 3: Wallet-level distributions. Panel (a): total dollar volume traded per wallet, on a base-10 log scale. Panel (b): average dollar size per trade per wallet, also on a base-10 log scale. Panel (c): realised dollar payoff per wallet on markets that have resolved by the end of the sample, winsorised at ±$100 for visibility. The payoff distribution has long tails in both directions, underscoring the substantial cross-sectional heterogeneity in trader sophistication on the platform.

The growing relevance of these markets is also reflected in their adoption by mainstream financial-data providers. In late 2025, Google Finance began incorporating data from Polymarket and Kalshi, allowing users to see real-time probability-based forecasts derived from market prices alongside conventional asset quotes; Bloomberg Terminal added a similar feed shortly thereafter. We do not mean to oversell what these integrations imply: a Bloomberg listing is not a substitute for asset-pricing evidence that the prices contain information. But it does indicate that the institutions whose prices we will be predicting in Section 6 regard the same prediction-market signals as worth displaying to their professional clients, which is at least consistent with the view that the signals are taken seriously by some fraction of the marginal investor base.

3.3 Predictive accuracy and informational efficiency

The third and final question we address in this section is whether Polymarket prices actually contain meaningful information about future outcomes, or whether they instead reflect noise and speculative behaviour. For our disagreement measures to be useful, the underlying prices must aggregate information well enough to provide reliable forecasts. We document this in two complementary ways: against external professional benchmarks, and against the eventual realisation of the underlying events themselves.

We begin with the comparison to professional forecasters. Recent work shows that regulated, macro-focused prediction markets can match or beat the Bloomberg Consensus on key macroeconomic releases; Diercks et al. (2026) document this for Kalshi. We perform the analogous exercise for Polymarket, comparing the market-implied mean to the Bloomberg Consensus and to the Kalshi market-implied mean for the two macro releases for which Polymarket runs a structured prediction market: the headline year-over-year CPI release and the unemployment-rate release. Our evaluation window spans November 2024 through July 2025, covering nine releases of each indicator on the immediate eve of the release. Table 2 reports the mean absolute error (MAE) for each forecast source.

	Polymarket (2024–2025)		Kalshi benchmark		Bloomberg
Macroeconomic indicator	Mean	Mode	Mean	Mode	Consensus
Headline CPI (YoY)	0.063	0.056	0.069	0.063	0.081
Unemployment rate	0.085	0.078	0.117	0.107	0.109

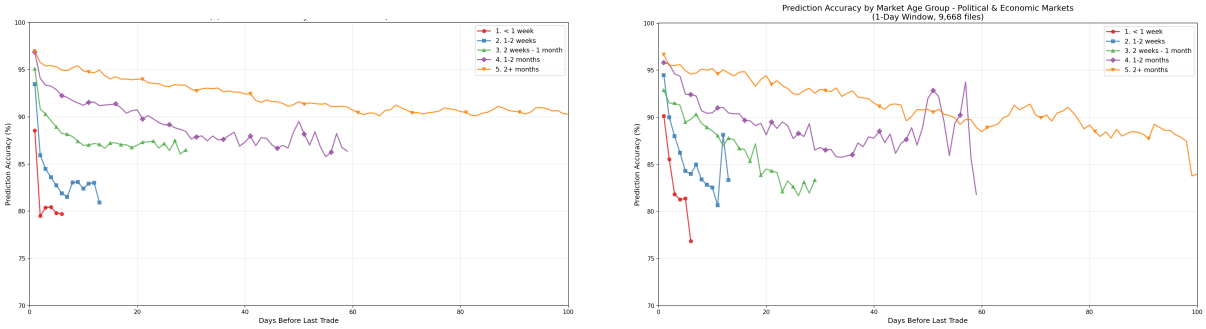
Table 2: Forecast accuracy on U.S. macroeconomic releases: mean absolute error (MAE) of the Polymarket-implied forecast against professional benchmarks. For headline CPI (year-over-year) and the unemployment rate, the table reports the MAE of the Polymarket-implied mean and mode, the corresponding Kalshi mean and mode, and the Bloomberg Consensus, each measured on the eve of the release. The Polymarket-implied mean is at least as accurate as the consensus of professional macro forecasters—roughly a 22% MAE improvement over Bloomberg on CPI, with a wider gap on the unemployment rate.

Notes: MAE on the eve of the release. Polymarket errors computed on nine releases each between November 2024 and July 2025. Kalshi and Bloomberg Consensus benchmarks are from Diercks et al. (2026).

The numbers in Table 2 are striking. On headline CPI, the Polymarket-implied mean has an MAE of 0.063, compared with 0.069 for Kalshi and 0.081 for the Bloomberg Consensus — a roughly 22% improvement over the consensus benchmark. On the unemployment rate the gap is wider: Polymarket achieves an MAE of 0.085 against 0.117 for Kalshi and 0.109 for Bloomberg. The sample is small, and the differences across forecast sources reflect, in part, slightly different sample windows. But the qualitative pattern is unambiguous: on the macroeconomic releases for which a direct comparison is feasible, Polymarket prices are at

least as accurate as the consensus of professional macro forecasters.

The second exercise looks at the full set of resolved markets and asks how accurately prices have predicted the eventual outcome at different points in each market's lifespan. We group markets by their total lifespan and, for each day before resolution, code a market as "correct" if the outcome that ultimately realises is also the one priced above 0.5 on that day. Within each age bin, prediction accuracy is the fraction of markets satisfying this condition. Figure 4 reports the result.



(a) Full sample

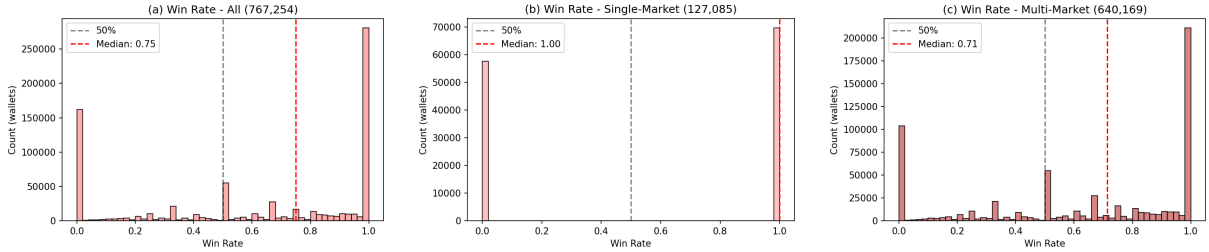
(b) Political and economic markets

Figure 4: Prediction accuracy as a function of time to resolution. Markets are grouped by total lifespan; for each day before resolution a market is coded "correct" if the outcome that ultimately realises is the one priced above 0.5 on that day, and accuracy within an age bin is the share of markets meeting this condition. Panel (a) uses all resolved markets; Panel (b) restricts to the political and economic sub-sample. Accuracy is high at every horizon—around 80% even at the earliest and closer to 90% for most cohorts—and rises monotonically toward resolution, reaching nearly 95% on the final day in the political and economic sample, consistent with prices aggregating incremental information over a market's life.

Two patterns stand out. First, accuracy is high throughout the lifespan of a market: even at the earliest horizons it sits around 80%, and for most market-age cohorts it is closer to 90%. Second, accuracy improves monotonically as markets approach resolution — exactly what we would expect if prices were aggregating incremental information over time. The political and economic sub-sample (Panel b) shows the pattern even more cleanly: accuracy rises from ~ 86% several days before resolution to nearly 95% on the final day. These patterns are difficult to reconcile with a noise-driven view of prediction-market trading.

The macro-forecast benchmark and the market-level accuracy series speak to the aggregate behaviour of prices. A natural next question is whether trader-level behaviour is also consistent with information aggregation. We can answer this directly using the persistent wallet identifiers described in Section 2.4. Figure 5 shows the cross-sectional distribution of wallet win rates, defined for each wallet as the fraction of resolved markets in which the wallet's final position matches the realised outcome. Two features stand out. First, the median wallet has a win rate well above 0.5 — traders on average outperform a coin-flip benchmark. Second, the distribution is highly dispersed and skewed toward higher values, with a substantial mass at 1. Even after restricting to multi-market wallets (Panel c) — which

removes the mechanical bimodality of single-trade wallets — the win-rate distribution remains skewed toward accuracy. This documents substantial heterogeneity in wallet-level forecasting skill, consistent with the view that Polymarket aggregates beliefs from a meaningfully informed cross-section of traders rather than a homogenous pool of noise participants.



(a) All wallets (b) Single-trade wallets (c) Multi-trade wallets

Figure 5: Cross-sectional distribution of wallet-level win rates, where a wallet's win rate is the fraction of resolved markets in which its final position matches the realised outcome. Panel (a) shows all wallets, Panel (b) single-trade wallets, and Panel (c) multi-trade wallets (which removes the mechanical bimodality induced by wallets active in only one market). The median wallet wins well above the 0.5 coin-flip benchmark and the distribution is dispersed and right-skewed, with substantial mass at 1; the skew toward accuracy survives among multi-market wallets. The pattern documents substantial heterogeneity in forecasting skill, consistent with Polymarket aggregating beliefs from a meaningfully informed cross-section of traders rather than a homogeneous pool of noise traders.

Taken together, the three pieces of evidence in this subsection — matching professional forecasters, monotonically improving accuracy through each market's lifespan, and a wallet-level win-rate distribution that is heterogeneous but informed on average — support the view that Polymarket prices aggregate dispersed information effectively. The platform is a serious price-discovery mechanism, not a noise-driven betting venue, and the disagreement measures we construct from it in the next section have economic content rather than reflecting random trading.

4 Measuring Uncertainty and Disagreement

In this section we construct the two measures of belief distribution, both built from the same source: the transaction-level Polymarket fills described in Section 2.4. Section 4.1 and Section 4.2 describe the constructions. The first, $\Psi_{m,t}$, operates on prices and is meant to capture the aggregate uncertainty about which future state will realise, that is, how unresolved the world is as priced by the market. The second, $\sigma_{m,t}$, operates on signed wallet trades and is meant to capture the cross-sectional disagreement across traders within a single market. Our guiding principle is to keep both as simple and intuitive as possible: each takes a narrow primitive that is directly observable in the data, a contract price or a signed trade, and

applies a small number of transparent operations to produce a scalar summary at the event-hour level, with no structural assumptions about risk premia, utility functions, or noise distributions.

This intuitive construction comes at a price: $\Psi_{m,t}$ and $\sigma_{m,t}$ are proxies, and it is not obvious a priori that they capture what their names claim. We then test whether each behaves as a genuine measure of its object should, subjecting Ψ to a battery of validity tests in Section 4.3 and σ in Section 4.4. Finally, Section 4.5 asks whether the two indeed capture distinct information rather than relabelling the same variation.

4.1 Uncertainty: the entropy index Ψ

The first measure operates on prices. Start with the simplest case, a single binary market on a *Yes*/*No* event. Its price is, directly, the market's implied probability that the event occurs, so it already reveals how uncertain the market is: a contract trading near 0.5 describes a coin flip, the state of maximal uncertainty, while one trading near 0 or 1 describes a question the market regards as all but settled. A *Yes* price of 0.90, for instance, leaves little doubt about the outcome, whereas a price of 0.50 leaves the most.

The same logic extends to an event represented by a family of mutually exclusive sub-markets, as in the FOMC example of Section 2.2. There the cross-section of *Yes* prices across the sub-markets is itself a vector of implied probabilities, and the dispersion of that vector measures how unresolved the event is as priced by the market. Suppose the FOMC market at time t shows *No Change* at 0.90, *25 bp cut* at 0.07, and three other outcomes summing to 0.03: the mass is concentrated on a single state and the event is, in pricing terms, essentially resolved. Compare a market showing the same five outcomes at $(0.20, 0.20, 0.20, 0.20, 0.20)$: every state has equal weight, and the uncertainty over which will realise is at its maximum.

We formalise this with Shannon entropy. For event e with sub-markets $\{m_1, \ldots, m_{n_e}\}$ and observed sub-market prices $\widehat{p}_{m_i,t}$, define the normalised price-implied probability of outcome i:

$$\widetilde{p}_{e,i,t} = \frac{\widehat{p}_{m_i,t}}{\sum_{j=1}^{n_e} \widehat{p}_{m_j,t}}. \tag{1}$$

The normalisation absorbs any small price-vector slack from illiquidity or bid–ask spread and makes $\widetilde{p}_{e,i,t}$ a proper probability vector. Event-level entropy is then

$$H_{e,t} = -\frac{1}{\log n_e} \sum_{i=1}^{n_e} \widetilde{p}_{e,i,t} \log \widetilde{p}_{e,i,t}, \tag{2}$$

normalised by $\log n_e$ so that $H_{e,t} \in [0,1]$ regardless of how many sub-markets the event contains. A binary market is simply the $n_e = 2$ case of the same definition. It has a single traded contract with observed price $\widehat{p}_{m,t}$ (the YES price), and its two mutually exclusive outcomes carry implied probabilities $\widetilde{p}_{e,1,t} = \widehat{p}_{m,t}$ and $\widetilde{p}_{e,2,t} = 1 - \widehat{p}_{m,t}$, which already sum to one so

the cross-market normalisation above is a no-op. Event-level entropy is then the standard binary entropy

$$H_{e,t} \;=\; -\frac{1}{\log 2}\big[\,\widehat{p}_{m,t}\log\widehat{p}_{m,t} + (1-\widehat{p}_{m,t})\log(1-\widehat{p}_{m,t})\,\big] \;\in\; [0,1], \tag{3}$$

maximised at $\widehat{p}_{m,t} = 0.5$, where uncertainty about which state will realise is greatest, and zero at $\widehat{p}_{m,t} \in \{0,1\}$, where the event is resolved. Binary markets therefore enter Ψ_t exactly as grouped events do, through the entropy of their price-implied outcome distribution.

To aggregate event-level entropies into a single time-series, we weight each event by its cumulative trading volume. Let $V_{m,s}$ denote dollar trading volume in sub-market m at time s; the cumulative event volume is $V^{\text{cum}}_{e,t} = \sum_{m\in e}\sum_{s\le t} V_{m,s}$. The uncertainty index is the volume-weighted average of event entropies:

$$\Psi_t \;=\; \frac{\sum_{e\in\mathscr{E}_t} H_{e,t}\, V^{\text{cum}}_{e,t}}{\sum_{e\in\mathscr{E}_t} V^{\text{cum}}_{e,t}}. \tag{4}$$

The cumulative dollar volume is our proxy for attention: it measures how much capital traders have actually committed to an event, and hence how much the market cares about the question, so the index loads on the states that are economically important to investors rather than on obscure or stale contracts. It also prevents a thinly traded new event from dominating the index simply because it has many near-equally priced sub-markets in its first few hours of life.

Ψ_t has three properties that make it useful as a measure of the unresolvedness of the world embedded in prices. It rises when the market spreads probability mass across the alternatives within an event, which exactly the configuration in which the eventual outcome is most uncertain. It falls when prices concentrate on a likely outcome, either because public information has resolved the question or because trading has converged on a consensus. And because it is constructed from prices alone, Ψ_t can be computed at any frequency at which prices update: in our applications we use daily end-of-day values, but a higher-frequency version is mechanical. The cost of relying only on prices is that Ψ_t cannot speak to cross-sectional disagreement among traders within a single market; the wallet-level measure $\sigma_{m,t}$ constructed next addresses this.

4.2 Investor disagreement: the dispersion index σ

The second measure operates on the cross-section of signed wallet trades within a single market. The starting observation again is simple: a transaction at price p reveals a bound on the wallet's risk-neutral belief about the event probability. A wallet that *buys YES at p* reveals that its belief about the YES probability is at least p; a wallet that *sells YES at p* reveals that its belief is at most p. By symmetry, a wallet that *buys NO at p* reveals that its YES probability

is at most $1-p$; a wallet that *sells NO at p* reveals that it is at least $1-p$. Every signed trade therefore generates an interval on the YES-probability scale within which the wallet's belief must lie.

Since the wallet's exact belief is unobserved, we summarise the interval by its midpoint. For trade f executed at price p_f, the revealed midpoint on the YES scale is

$$\bar{q}_f = \begin{cases} (p_f+1)/2 & \text{if the wallet bought YES,} \\ p_f/2 & \text{if the wallet sold YES,} \\ (1-p_f)/2 & \text{if the wallet bought NO,} \\ (2-p_f)/2 & \text{if the wallet sold NO.} \end{cases} \tag{5}$$

The midpoint is a deliberately conservative summary: it does not require us to commit to a specific assumption about where in the revealed interval the wallet's belief actually sits. A wallet that buys YES at 0.10 reveals a belief above 0.10, and the midpoint (0.55) reflects neither the boldest nor the most cautious interpretation. A side advantage of the midpoint construction is that it is robust to small departures from no-arbitrage in the realised price.

Within each wallet-market-hour triple (i,m,t), we aggregate all of wallet i's fills into a single dollar-volume-weighted revealed midpoint:

$$m_{i,m,t} = \frac{\sum_{f\in(i,m,t)} v_f\,\bar{q}_f}{\sum_{f\in(i,m,t)} v_f}, \qquad v_f = \text{shares}_f \times p_f. \tag{6}$$

This compresses a possibly large number of fills by the same wallet in the same market-hour into a single wallet-level belief proxy. The disagreement measure is then the dollar-volume-weighted cross-sectional variance of these per-wallet beliefs:

$$\mu_{m,t} = \sum_i w_{i,m,t}\, m_{i,m,t}, \qquad \sigma^2_{m,t} = \sum_i w_{i,m,t}\,(m_{i,m,t}-\mu_{m,t})^2, \qquad w_{i,m,t} = \frac{v_{i,m,t}}{V_{m,t}}, \tag{7}$$

where $V_{m,t} = \sum_i v_{i,m,t}$. We require at least two active wallets in (m,t) for $\sigma_{m,t}$ to be defined.

σ has four distinct properties. First, it is bounded in $[0,0.5]$: the maximum variance of a $[0,1]$-valued random variable under any distribution is attained at a two-point mass at the endpoints, giving variance 0.25 and standard deviation 0.5. Second, it is large precisely when wallets with substantial capital at stake are on opposing sides of the trade — the dollar-volume weighting deliberately upweights wallets whose revealed disagreement is backed by money, which we view as a feature rather than a bug. Third, it can be strictly positive even when the contract price is flat and realised price volatility is zero: steady two-sided flow at a price p still reveals wallets on opposite sides of the cutoff, a YES buyer with revealed midpoint $(p+1)/2$ and a NO buyer with midpoint $p/2$, so the cross-section of beliefs is dispersed even though the price never moves. This should be read carefully. A flat price does not make $\sigma_{m,t}$ *move*; with stable two-sided flow it simply sits at a constant positive level, and it changes

only when the composition of revealed beliefs changes, through the balance of buyers and sellers or the prices at which they trade. The point is that this level is informative about disagreement precisely where the price and its volatility, both flat here, are silent. This is the key sense in which σ picks up information that neither the price nor Ψ captures. At the same time, $\sigma_{m,t}$ *does* rise with price dynamics when disagreement is expressed through trading: when wallets with opposing views push the price back and forth, they transact at a range of prices, which widens the spread of revealed midpoints and raises $\sigma_{m,t}$ at the same time as realised volatility. The two channels are complementary, so σ is high at a flat price whenever flow is two-sided and higher still when active disagreement moves the price around, which makes it positively related to volatility without being reducible to it. Fourth, the construction is fully nonparametric: it requires no model of risk premia, no auxiliary distribution, and no calibration. It is a direct function of the trade tape.

To aggregate the market-level dispersions into a single time series, we use the same cumulative-volume weighting as for Ψ_t. Let $\mathcal{M}_t$ denote the set of markets active at t and $V^{\text{cum}}_{m,t} = \sum_{s \le t} V_{m,s}$ the cumulative dollar volume of market m. The disagreement index is the volume-weighted average of the market-level dispersions:

$$\sigma_t \;=\; \frac{\sum_{m \in \mathcal{M}_t} V^{\text{cum}}_{m,t}\, \sigma_{m,t}}{\sum_{m \in \mathcal{M}_t} V^{\text{cum}}_{m,t}}. \tag{8}$$

This is the object that enters the asset-pricing regressions of Section 6.

4.3 Validating the uncertainty index Ψ

We now subject Ψ to a battery of validity tests, asking whether it behaves the way a genuine measure of uncertainty should. We examine in turn its relation to the realised outcome variance once markets resolve, to the magnitude of future price movements, to contemporaneous trading volume, and to the arrival of macro news.

Uncertainty matches the realised outcome variance. This test exploits a feature unique to prediction markets: contracts resolve to a ground-truth outcome $o_m \in \{0, 1\}$, so we can compare the *ex-ante* uncertainty embedded in the price against the *ex-post* forecast error actually realised. It is the prediction-market analogue of forecast calibration, and no survey- or option-based uncertainty proxy can run it, because those forecasts are never settled against a binary truth. The construction is as follows. For a contract priced at $\widehat{p}_{m,t}$, the realised squared forecast error — the Brier score of that forecast — is

$$\text{Brier}_{m,t} \;=\; \left(o_m - \widehat{p}_{m,t}\right)^2. \tag{9}$$

If the price is well calibrated (i.e. $\widehat{p}_{m,t}$ is the true probability), the expected Brier error conditional on the price is exactly the Bernoulli variance,

$$\mathbb{E}\big[\,(o_m - \widehat{p}_{m,t})^2 \,\big|\, \widehat{p}_{m,t}\,\big] \;=\; \widehat{p}_{m,t}\big(1-\widehat{p}_{m,t}\big), \tag{10}$$

the irreducible "uncertainty" term in the Murphy (1973) decomposition of the Brier score; our index Ψ is, by construction, a monotone transform of that same $\widehat{p}_{m,t}(1-\widehat{p}_{m,t})$. Two predictions follow: (i) the realised Brier error should load positively on Ψ, and (ii) regressed on $\widehat{p}_{m,t}(1-\widehat{p}_{m,t})$ directly it should load with a coefficient of one whenever the market is calibrated. Table 3 reports pooled regressions of the realised Brier error across all resolved market-days, with HAC standard errors. Both predictions hold: the Brier error loads strongly on Ψ (+0.235), and on $p(1-p)$ with a coefficient of 0.97 $\approx$ 1, so the markets are almost perfectly calibrated and Ψ inherits that calibrated outcome variance. Figure 6 shows the same result non-parametrically: sorting market-days into deciles of Ψ and plotting, for each decile, the average realised Brier error against the average price-implied variance, the two curves lie almost on top of one another and climb monotonically from near zero — in the low-Ψ decile, where prices sit near 0 or 1 and outcomes are nearly certain — to about 0.24 in the top decile, where prices sit near 0.5 and outcomes are maximally uncertain. Ψ is therefore not merely *labelled* uncertainty: it equals the realised variance of the outcomes it describes.

Variable	(1)	(2)
Ψ (entropy)	+0.235***	
	(0.001)	
Price-implied var. $p(1-p)$		+0.970***
		(0.004)
R^2	0.238	0.246
N	873,842	873,842

Table 3: Outcome-based validity: Ψ matches the realised outcome variance. Dependent variable is the realised squared forecast error $(o_m - \widehat{p}_{m,t})^2$ at the market-day level (resolved markets). Pooled OLS with HAC(10) standard errors.***p<0.01, **p<0.05, *p<0.10.

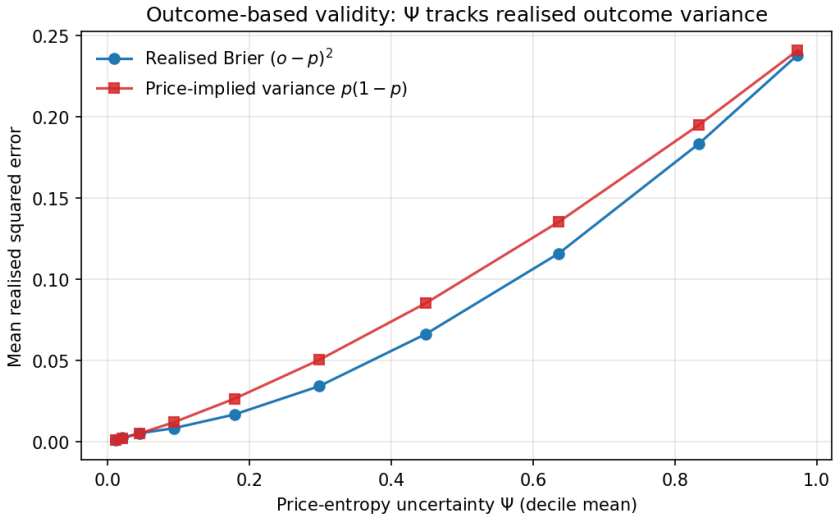


Figure 6: Outcome-based validity. Markets are binned into deciles of Ψ; for each bin the figure plots the realised Brier error $(o-p)^2$ and the price-implied variance $p(1-p)$ against mean Ψ. The two lie almost on top of each other and rise monotonically in Ψ, confirming that Ψ measures the markets' realised outcome uncertainty.

Uncertainty predicts future realised volatility. If Ψ truly measures how unresolved the state is, it should forecast the *magnitude* of subsequent price movement. The reasoning is almost mechanical: a contract priced near 0.5 carries the most pending uncertainty and has the furthest to travel before it settles at its eventual 0/1 outcome, so its price must still move a great deal as the state resolves, whereas a contract already near 0 or 1 has little left to resolve and should barely move. High Ψ should therefore be followed by high realised volatility of the contract's own price. Table 4 tests this by regressing the market's realised volatility over the next h days (the forward sum of squared daily logit returns) on $\Psi_{m,t}$. The loading is positive and significant at every horizon, and largest at the shortest horizon where the most resolution is still to come, exactly as the prediction implies.

	h=1d	h=2d	h=5d	h=10d
$\Psi_{m,t}$	$+1.162^{***}$	$+0.846^{***}$	$+0.833^{***}$	$+0.751^{***}$
	(0.061)	(0.081)	(0.127)	(0.217)
Lag return	-0.043^{***}	-0.061^{***}	-0.054^{***}	-0.060^{**}
	(0.012)	(0.014)	(0.019)	(0.028)
Lag RV_5	-0.010^{***}	-0.022^{***}	-0.047^{***}	-0.097^{***}
	(0.003)	(0.006)	(0.013)	(0.019)
Lag volume	$+0.030^{***}$	$+0.024^{***}$	$+0.002$	-0.032^{***}
	(0.002)	(0.003)	(0.004)	(0.006)
log TTR	-0.677^{***}	-0.871^{***}	-1.176^{***}	-1.589^{***}
	(0.018)	(0.024)	(0.043)	(0.076)
log Age	$+0.011$	-0.005	-0.054^{**}	-0.102^{***}
	(0.008)	(0.011)	(0.022)	(0.036)
Market FE	Yes	Yes	Yes	Yes
R^2	0.018	0.021	0.024	0.031
N	798,291	766,560	689,521	596,047

Table 4: Ψ predicts subsequent realised volatility of the Polymarket price. Dependent variable is RV_h, the forward sum of squared daily logit returns over the next h days. Market FE; SE clustered by market. Additional controls (lag price-logit, lag wallets, lag HHI) are included but omitted from display. ***p<0.01, **p<0.05, *p<0.10.

Uncertainty does *not* generate volume. This is the key discriminant test, the one that most sharply separates uncertainty from disagreement. Trade requires two parties who value the contract differently, so it is *disagreement* that should generate volume. Uncertainty is the opposite kind of object: when Ψ is high the market broadly agrees that the outcome is hard to call, a shared "we do not know" rather than a clash of views, and a shared assessment gives no one a reason to trade against anyone else. We should therefore expect Ψ to carry *no* positive volume loading, and perhaps a negative one, even though disagreement co-moves strongly with volume. Table 5 regresses market-day log dollar volume on Ψ and bears this out: the contemporaneous loading is *negative* (-0.245, and -0.730 once the lag of Ψ is included), the opposite of what we will find for disagreement (Section 4.4). The negative sign also lays to rest a mechanical concern: although Ψ is built from a volume-weighted price aggregate, it does not simply inherit the volume series; if anything it moves opposite to contemporaneous volume.

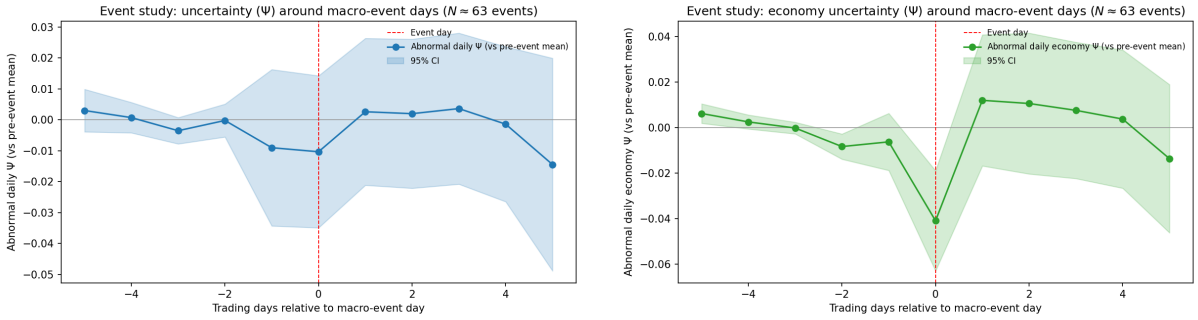
Variable	(1)	(2)
$\Psi_{m,t}$	-0.245^{***}	-0.730^{***}
	(0.029)	(0.026)
$\Psi_{m,t-1}$		$+0.793^{***}$
		(0.032)
Lag volume	$+0.329^{***}$	$+0.328^{***}$
	(0.004)	(0.004)
Lag RV_5	-0.011^{***}	-0.009^{***}
	(0.001)	(0.001)
Lag return	-0.053^{***}	-0.059^{***}
	(0.005)	(0.005)
Lag HHI	-0.896^{***}	-0.862^{***}
	(0.018)	(0.018)
log TTR	-0.232^{***}	-0.236^{***}
	(0.010)	(0.010)
log Age	$+0.129^{***}$	$+0.139^{***}$
	(0.007)	(0.007)
Market FE	Yes	Yes
R^2	0.182	0.183
N	834,367	834,367

Table 5: Ψ does not co-move with trading volume (discriminant test). Dependent variable is market-day log dollar volume. Market FE; SE clustered by market. Additional controls (lag price-logit, lag wallets) are included but omitted from display. ***p<0.01, **p<0.05, *p<0.10.

Uncertainty resolves around news. The final check asks whether uncertainty falls when fresh information arrives. If Ψ measures how unresolved the state is, a scheduled macro release should *lower* it as the state distribution collapses toward the realised outcome. We test this with an event study around U.S. macro-event days (FOMC announcements and the major GDP, CPI, and non-farm-payroll releases). Let $\mathscr{E}$ be the set of event days and τ_e the date of event e; for each event we subtract a pre-event baseline (the mean over the five days before the event) and average the resulting abnormal values across events at each lag k:

$$\overline{A\Psi}_k = \frac{1}{|\mathscr{E}|}\sum_{e\in\mathscr{E}}\left(\Psi_{\tau_e+k} - \overline{\Psi}_e^{\text{pre}}\right), \qquad \overline{\Psi}_e^{\text{pre}} = \frac{1}{5}\sum_{j=-5}^{-1}\Psi_{\tau_e+j}, \qquad k\in\{-5,\dots,+5\}. \tag{11}$$

Figure 7 bears the prediction out. In economy markets (Panel b) abnormal Ψ drops sharply and significantly on the release day ($k = 0$), the signature of uncertainty being resolved as the news arrives, before partially rebuilding; the overall index (Panel a) shows no systematic move, consistent with the events being specifically macroeconomic.



(a) Overall index

(b) Economy markets

Figure 7: Resolution event study: abnormal daily Ψ around U.S. macro-event days (±5 trading days, baselined to the pre-event mean). Economy-market uncertainty drops sharply on the release day, as uncertainty is resolved.

Ψ thus passes the battery designed for an uncertainty measure: it predicts the magnitude of future price moves, matches the realised outcome variance market by market, and resolves on news, while, crucially, not generating volume. We now put the disagreement index through its own battery.

4.4 Validating the disagreement index σ

A genuine disagreement measure should display four properties within the platform itself: it should be highest when trading is two-sided (differences of opinion, not one-way order flow); it should rise when fresh information arrives; and — following the classic differences-of-opinion literature — it should predict subsequent realised volatility (Hong and Stein, 1999; Banerjee and Kremer, 2010b) and co-move with trading volume (the Karpoff/Harris–Raviv channel). We document each in turn, working on the market-day panel with market fixed effects and standard errors clustered by market.

Disagreement reflects two-sided trading. The first check asks whether σ responds to the *structure* of order flow the way a differences-of-opinion measure must. For each market-day we compute the order imbalance

$$\text{Imb}_{m,t} = \frac{\left|\text{Buy}_{m,t} - \text{Sell}_{m,t}\right|}{\text{Buy}_{m,t} + \text{Sell}_{m,t}} \in [0,1], \tag{12}$$

the absolute net signed dollar flow scaled by gross dollar flow, where $\text{Buy}_{m,t}$ and $\text{Sell}_{m,t}$ are buyer- and seller-initiated dollar volume. It equals 0 when buying and selling are perfectly balanced (the two sides trade against each other in equal size) and 1 when flow is entirely one-directional (everyone is on the same side). The logic of the test is that genuine disagreement requires two willing counterparties with opposing views: when traders truly disagree, order flow is two-sided and the imbalance is low; when flow is one-sided — herding, or a

common reassessment that everyone shares — there is little disagreement to measure. A valid σ should therefore *fall* as the imbalance rises. Table 6 regresses market-day $\sigma_{m,t}$ on the imbalance with market fixed effects, so identification is within-market and day-to-day, and SE clustered by market and day. The coefficient is negative and highly significant both bivariate (-0.0586) and with the full control set (-0.0531): σ is highest exactly when trade is two-sided. The uncertainty index adds nothing to this margin ($\Psi = +0.0022$, insignificant), so two-sidedness of trade is a property σ captures and Ψ does not.

Variable	(1)	(2)
Buy/sell imbalance	-0.0586^{***}	-0.0531^{***}
	(0.0006)	(0.0006)
Log daily volume		$+0.0045^{***}$
		(0.0001)
Price dist. from 0.5		$+0.0425^{***}$
		(0.0053)
Log(1+age)		$+0.0009^{***}$
		(0.0002)
Log(1+days to res.)		-0.0017^{***}
		(0.0003)
Ψ		$+0.0022$
		(0.0022)
Market FE	Yes	Yes
R^2	0.1421	0.1976
N	970,793	938,997

Table 6: Disagreement is highest when order flow is two-sided. Dependent variable is market-day $\sigma_{m,t}$. Column (1) is bivariate; column (2) adds activity, price level, age, time-to-resolution, and Ψ. Market fixed effects; SE clustered by market and day. ***p<0.01, **p<0.05, *p<0.10.

Disagreement predicts future realised volatility. Differences of opinion are worked out through trading: when wallets hold opposing views they take each other on, and the price keeps moving as the disagreement is resolved, so high disagreement today should be followed by larger subsequent price movement. This is the classic differences-of-opinion prediction, documented empirically by Carlin et al. (2014). Table 7 tests it by regressing the market's realised volatility over the next h days (the forward sum of squared daily logit returns) on $\sigma_{m,t}$. The loading is positive and significant at every horizon, and decays as the horizon lengthens and the forward window averages out, exactly as predicted. Note that Ψ forecasts future volatility too (Section 4.3), but for a different reason: there it is pending uncertainty that must resolve, here it is disagreement that must be traded away.

	h=1d	h=2d	h=5d	h=10d
$\sigma_{m,t}$	$+3.780^{***}$	$+2.906^{***}$	$+1.483^{***}$	$+0.548^{**}$
	(0.128)	(0.132)	(0.165)	(0.221)
Lag return	-0.056^{***}	-0.072^{***}	-0.066^{***}	-0.073^{***}
	(0.012)	(0.014)	(0.019)	(0.028)
Lag RV_5	-0.013^{***}	-0.024^{***}	-0.049^{***}	-0.098^{***}
	(0.003)	(0.006)	(0.014)	(0.019)
Lag volume	$+0.017^{***}$	$+0.014^{***}$	-0.003	-0.033^{***}
	(0.002)	(0.003)	(0.004)	(0.006)
log TTR	-0.605^{***}	-0.820^{***}	-1.128^{***}	-1.547^{***}
	(0.016)	(0.023)	(0.041)	(0.074)
log Age	-0.041^{***}	-0.041^{***}	-0.080^{***}	-0.121^{***}
	(0.007)	(0.011)	(0.022)	(0.036)
Market FE	Yes	Yes	Yes	Yes
R^2	0.018	0.021	0.024	0.030
N	797,649	765,991	689,105	595,767

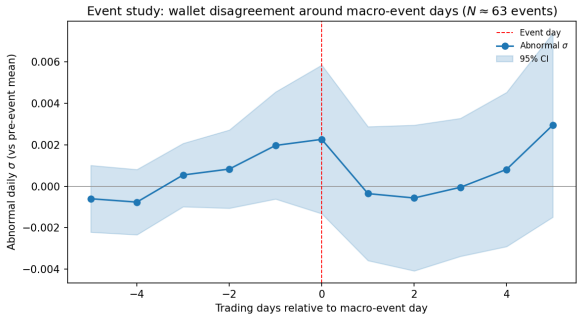
Table 7: σ predicts subsequent realised volatility of the Polymarket price. Dependent variable is RV_h, the forward sum of squared daily logit returns over the next h days. Market FE; SE clustered by market.Additional controls (lag price-logit, lag wallets, lag HHI) are included but omitted from display. ***p<0.01, **p<0.05, *p<0.10.

Disagreement co-moves with trading volume. The final and most discriminating check is volume. Trade requires two parties who value the contract differently, so differences of opinion mechanically generate turnover: the more wallets disagree, the more they trade against one another. This is the differences-of-opinion prediction of Harris and Raviv (1993b), also borne out empirically by Carlin et al. (2014), and it is exactly where disagreement parts company with uncertainty, since a shared "we do not know" (high Ψ) gives no one a reason to trade. Table 8 regresses market-day trading volume on contemporaneous $\sigma_{m,t}$. The loading is large and strongly positive (+9.21, and +9.02 with the lag of σ included), so high-disagreement markets are high-volume markets. The *negative* Ψ loading documented earlier (Section 4.3) goes the other way, so the two measures load on trading volume with opposite signs: σ positively and Ψ negatively.

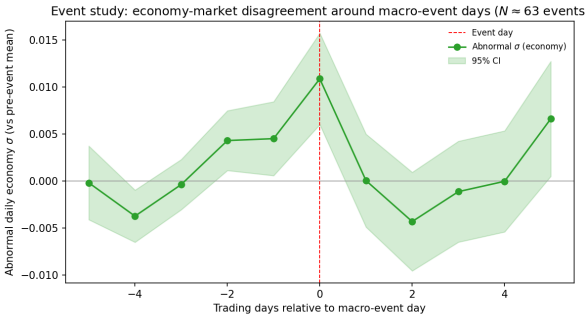
Variable	(1)	(2)
$\sigma_{m,t}$	$+9.214^{***}$	$+9.017^{***}$
	(0.102)	(0.094)
$\sigma_{m,t-1}$		$+1.284^{***}$
		(0.091)
Lag volume	$+0.303^{***}$	$+0.300^{***}$
	(0.003)	(0.003)
Lag RV_5	-0.006^{***}	-0.006^{***}
	(0.001)	(0.001)
Lag return	-0.043^{***}	-0.043^{***}
	(0.004)	(0.004)
Lag HHI	-0.671^{***}	-0.606^{***}
	(0.015)	(0.014)
log TTR	-0.231^{***}	-0.230^{***}
	(0.009)	(0.009)
log Age	$+0.099^{***}$	$+0.095^{***}$
	(0.007)	(0.007)
Market FE	Yes	Yes
R^2	0.218	0.219
N	833,585	832,927

Table 8: σ co-moves with contemporaneous trading volume. Dependent variable is market-day log dollar volume. Market FE; SE clustered by market. Additional controls (lag price-logit, lag wallets) are included but omitted from display. ***p<0.01, **p<0.05, *p<0.10.

Disagreement rises on macro-news days. The last check asks whether σ responds to the *arrival* of information. We run the same event study as for uncertainty (Section 4.3), with the identical macro-event calendar, ±5-day window, and pre-event baseline, now applied to the daily topic-level σ series. A genuine disagreement measure should *rise* when fresh, contestable information arrives, the mirror image of the drop we documented for Ψ. Figure 8 shows that it does, but *only where it should*: the response is concentrated in economy markets (Panel b), where abnormal disagreement jumps on the release day ($k = 0$), while the overall index (Panel a) shows no systematic shift. That disagreement *rises* while uncertainty *falls* on the very same dates is the discriminant we are after: a macro release splits traders into opposing positions even as it collapses the consensus state distribution, so σ and Ψ respond to identical news with opposite sign.



(a) Overall index



(b) Economy markets

Figure 8: Event study: abnormal daily σ around U.S. macro-event (news-release) days (± 5 trading days, baselined to the pre-event mean). The disagreement response is concentrated in economy markets, as expected when the events are macroeconomic releases.

Across all four checks, σ behaves exactly as a differences-of-opinion measure should: it peaks in two-sided trade, predicts both future volatility and contemporaneous volume, and rises on news. With both indices now validated against what they claim to measure, Ψ as uncertainty and σ as disagreement, we can ask whether they are, after all, the same object.

4.5 Do Ψ and σ measure the same object?

The two measures are constructed from different primitives: Ψ_t uses the cross-section of prices across the mutually exclusive states within each event (the contracts of a grouped event, or the two sides of a binary market), while σ_t uses the cross-section of revealed wallet beliefs within contracts. They may nevertheless move together mechanically if both are just proxies for general Polymarket activity. We therefore test directly whether one measure explains the other once we control for lagged platform activity.

We estimate two time-series specifications:

$$\Psi_t = \alpha + \beta \sigma_t + \gamma' A_{t-1} + \varepsilon_t, \tag{13}$$

$$\sigma_t = \alpha + \beta \Psi_t + \gamma' A_{t-1} + \varepsilon_t, \tag{14}$$

where the PM-only activity control vector is

$$A_{t-1} = \{\log(\text{PMVolume})_{t-1}, \log(\text{PMActiveWallets})_{t-1}\}.$$

The key coefficient is β: if Ψ_t and σ_t are measuring the same object, then σ_t should explain Ψ_t, and Ψ_t should explain σ_t, even after conditioning on Polymarket activity.

	Ψ_t on σ_t		σ_t on Ψ_t	
Variable	(1)	(2)	(3)	(4)
σ_t	+0.8456	+0.6443	–	–
	(0.8089)	(0.6477)	–	–
Ψ_t	–	–	+0.0161	+0.0131
	–	–	(0.0137)	(0.0122)
Lag log PM volume	–	+0.0787**	–	+0.0173***
	–	(0.0343)	–	(0.0028)
Lag log PM active wallets	–	−0.1981***	–	−0.0048
	–	(0.0327)	–	(0.0051)
PM activity controls	No	Yes	No	Yes
HAC lags	5	5	5	5
R^2	0.0136	0.3481	0.0136	0.3021
N	420	420	420	420

Table 9: Orthogonality of uncertainty and disagreement. Daily time-series OLS with Newey–West/HAC (5-lag) standard errors. Columns (1)–(2) regress Ψ_t on σ_t; columns (3)–(4) regress σ_t on Ψ_t. σ_t and Ψ_t are contemporaneous daily levels; the PM activity controls (lag log PM volume and lag log PM active wallets) are lagged one day. ***p<0.01, **p<0.05, *p<0.10.

Table 9 shows that the cross-loading coefficients are small and statistically insignificant in both directions. In the Ψ_t regression, the coefficient on σ_t is insignificant both without controls (column 1) and after adding lagged PM volume and active wallets (column 2). In the σ_t regression, the coefficient on Ψ_t is also insignificant in both specifications (columns 3 and 4). The signal-only R^2 is essentially zero in both directions (0.014), and the higher R^2 in the PM-activity specifications comes from the activity controls rather than from the other belief measure. Thus, after controlling for basic Polymarket activity, Ψ_t and σ_t are not meaningfully correlated. We interpret this as evidence that the two variables capture distinct dimensions of Polymarket prices: Ψ_t measures uncertainty about which state will realise, while σ_t measures disagreement within traded contracts.

A second way to see that the two are distinct is to follow them over the market life cycle. Table 10 regresses each measure on (standardized) market age and age-squared, with market fixed effects; age is standardized over the 0–365-day window so the coefficients are read per standard deviation of age. The two measures trace out different shapes. Uncertainty Ψ falls with age and does so *convexly*: the linear term is strongly negative (−0.13) and the quadratic term is positive (+0.012, column 2), so Ψ drops steeply when a market is young and the decline decelerates as the event nears resolution — the natural profile of uncertainty being progressively resolved. Disagreement σ instead traces a mid-life *hump*: a positive linear and negative quadratic term (column 4) mean it builds up early and unwinds late, peaking when the market is most actively contested rather than at birth or resolution. The two curvatures are thus opposite — Ψ convex-decreasing, σ concave-humped. The two objects therefore

not only load on different margins contemporaneously but also follow different — and differently sourced — life cycles.

	$\Psi_{m,t}$		$\sigma_{m,t}$	
Variable	(1)	(2)	(3)	(4)
Age (SD)	-0.1095^{***}	-0.1338^{***}	$+0.0066^{***}$	$+0.0090^{***}$
	(0.0028)	(0.0033)	(0.0004)	(0.0004)
Age^2 (SD)		$+0.0119^{***}$		-0.0012^{***}
		(0.0015)		(0.0002)
Market FE	Yes	Yes	Yes	Yes
R^2	0.149	0.155	0.010	0.011
N	877,597	877,597	876,762	876,762

Table 10: Life cycle: Ψ and σ over market age. Dependent variable is $\Psi_{m,t}$ (columns 1–2) and $\sigma_{m,t}$ (columns 3–4). Age is standardized over the 0–365-day window; columns 2 and 4 add age^2. Market fixed effects; SE clustered by market and day. ***p<0.01, **p<0.05, *p<0.10.

5 How are Polymarket and the Equity Market Connected?

This section establishes how Polymarket and the U.S. equity market are connected. This connection is the fundamental nexus on which the rest of the paper rests. The two venues are populated by very different participants: retail-skewed, globally dispersed traders on Polymarket, and the largely institutional investor base of U.S. equities. It is therefore natural to ask how anything that happens on Polymarket could matter for stock prices at all. Our answer is that the link is *informational*: we do not claim that the same agents trade both markets, or that capital flows directly between them, only that the two price the same underlying states and that information about those states is shared and transmitted across them.

Our hypothesis is that information about these states flows between the two venues in *both* directions, so that each price reflects information that also reaches the other. The two subsections that follow examine the two directions in turn.

5.1 Polymarket reacts to stocks

We begin with the more natural of the two directions: that Polymarket prices *react* to the equity market. The hypothesis is that Polymarket traders actively track real-world developments, the equity tape among them, so that any move in an underlying stock or index is reflected almost immediately in the Polymarket-implied probability of an event tied to it.

We run the test on two families of prediction markets: the S&P 500 daily Up-or-Down markets and the firm-level price-target markets on individual S&P 500 constituents.

The first family is the set of 124 S&P 500 daily Up-or-Down markets, each of which resolves YES if the S&P 500 closes up on a specified date and NO otherwise. For each market

we construct an hourly panel of the volume-weighted Polymarket price between market creation and resolution and merge it with the corresponding hourly close of the S&P 500. The sample window is the same as the rest of the paper.

The specification is an hourly panel OLS with contract fixed effects:

$$\Delta p^{\mathrm{PM}}_{m,t} = \alpha_m + \beta_0\, r^{\mathrm{stk}}_{m,t} + \beta_1\, r^{\mathrm{stk}}_{m,t-1} + \beta_{+24h}\, r^{\mathrm{stk}}_{m,t+24h} + \varepsilon_{m,t}, \tag{15}$$

where $\Delta p^{\mathrm{PM}}_{m,t}$ is the hourly first difference of the volume-weighted, YES-aligned Polymarket price and $r^{\mathrm{stk}}_{m,t}$ is the contemporaneous hourly log return on the S&P 500. Because YES pays off when the index closes up, a positive $\Delta p^{\mathrm{PM}}_{m,t}$ signals that the market has grown more optimistic that the index will rise. The right-hand side is the same S&P 500 return for every contract active in a given hour, so identification comes entirely from the within-contract, over-time covariation between the index return and the change in a contract's price. The contract fixed effect α_m is not there to absorb differences in the regressor across contracts; it removes each contract's average hourly price drift, such as the mechanical drift toward 0 or 1 as a daily contract ages to its resolution date, so that β_0 and β_1 are identified from genuine within-contract responses to the index. Time fixed effects, by contrast, are not an option here, as they would absorb the common index return that is the object of interest. Standard errors are clustered by contract. The lead term $r^{\mathrm{stk}}_{m,t+24h}$ is the return twenty-four hours ahead at the same hour of day; we use it as a placebo on the subsample where $t+24h$ falls strictly after market resolution, so that the contract is already settled and no economic mechanism could let the future return move the current price. This placebo is the sharpest available test for spurious panel correlation: if β_{+24h} is statistically indistinguishable from zero, the contemporaneous and lagged loadings cannot be artefacts of unobserved common factors operating at the daily frequency.

Table 11 reports the results for the S&P 500 daily Up-or-Down markets.

	(1)	(2)	(3)
	Δp_t^{PM}	Δp_t^{PM}	Δp_t^{PM}
$r_t^{\mathrm{S\&P}}$	8.21^{***}	9.96^{***}	11.35^{***}
	(1.49)	(1.68)	(2.10)
$r_{t-1}^{\mathrm{S\&P}}$		7.01^{***}	10.08^{***}
		(1.67)	(1.99)
$r_{t+24h}^{\mathrm{S\&P}}$			-1.02
			(1.74)
Observations	992	992	691
Within-R^2	0.059	0.106	0.147
Contract FE	Yes	Yes	Yes

Table 11: Polymarket reacts to the S&P 500 index. Hourly panel of S&P 500 Up-or-Down prediction-market prices. The dependent variable is the hourly first difference of the volume-weighted YES-probability. Column (2) adds the lagged S&P return; column (3) adds the 24h-ahead placebo and is estimated on the subsample where that placebo falls strictly after market resolution (HTR < 24h). Contract fixed effects throughout. Standard errors clustered by contract in parentheses.

The contemporaneous coefficient $\hat{\beta}_0 = 8.21$ in column (1) says that a 1% hourly S&P 500 return shifts the implied probability of the underlying S&P daily Up-or-Down prediction market by approximately 8 percentage points within the same hour. Adding the one-hour lag in column (2) brings the contemporaneous coefficient up to $\hat{\beta}_0 = 9.96$ and reveals an additional lagged effect of $\hat{\beta}_1 = 7.01$ percentage points, so the full two-hour response to a 1% S&P move is on the order of 17 percentage points of YES-probability. Both coefficients are significant at the 1% level. In column (3) the placebo $\hat{\beta}_{+24h} = -1.02$ is statistically indistinguishable from zero ($t = -0.59$). The lead twenty-four hours of return, conditional on the contemporaneous and lagged hour, contains no information about today's Polymarket-price change. The contemporaneous and lagged loadings, then, are not artefacts of daily-frequency unobserved factors.

The second family is the set of firm-level price-target markets keyed on individual S&P 500 constituents (e.g. *Will NVIDIA (NVDA) close above $170 on April 10?*). We restrict to the 4,615 direction-cleaned firm markets and, as before, build an hourly panel of the volume-weighted YES-aligned Polymarket price and merge it with the matched ticker's hourly close.

The specification mirrors the index-level test, now with the underlying stock's return on the right-hand side:

$$\Delta p_{m,t}^{\mathrm{PM}} = \alpha_m + \beta_0 r_{m,t}^{\mathrm{stk}} + \beta_1 r_{m,t-1}^{\mathrm{stk}} + \beta_{+24h} r_{m,t+24h}^{\mathrm{stk}} + \varepsilon_{m,t}, \tag{16}$$

where $\Delta p_{m,t}^{\mathrm{PM}}$ is the hourly first difference of the volume-weighted, YES-aligned price of firm

market m and $r^{\text{stk}}_{m,t}$ is the contemporaneous hourly log return on the stock that market m is keyed on. The direction-cleaning orients each contract so that its YES leg pays off when the stock rises, so a positive $\Delta p^{\text{PM}}_{m,t}$ again signals a more optimistic view of the underlying. As in the index-level test, α_m is a contract fixed effect that removes each contract's average hourly price drift, standard errors are clustered by contract, and $r^{\text{stk}}_{m,t+24h}$ is a twenty-four-hour-ahead placebo evaluated on the subsample where it falls strictly after resolution. Table 12 reports the results.

	(1)	(2)	(3)
	Δp^{PM}_t	Δp^{PM}_t	Δp^{PM}_t
r^{stk}_t	2.26***	2.44***	4.27***
	(0.096)	(0.100)	(0.216)
r^{stk}_{t-1}		2.05***	3.38***
		(0.098)	(0.192)
r^{stk}_{t+24h}			0.06
			(0.153)
Observations	30,491	30,491	14,748
Within-R^2	0.040	0.067	0.097
Contract FE	Yes	Yes	Yes

Table 12: Polymarket reacts to the matched stock (firm markets). Hourly panel of direction-cleaned firm prediction markets. The dependent variable is the hourly first difference of the YES-aligned Polymarket price. Column structure mirrors Table 11. Contract fixed effects throughout; standard errors clustered by contract in parentheses.

The firm-level magnitudes are smaller than the index-level magnitudes, which is expected: a 1% stock return moves the matched firm-market YES probability by ~ 2.4 percentage points contemporaneously and ~ 2.1 percentage points at the one-hour lag, against ~ 10 and ~ 7 percentage points respectively in the S&P daily sample. The reason is sample composition: most firm markets are weekly price-target contracts that resolve days away, so a single 1% return moves the probability of a still-far target only modestly. The S&P daily markets, by contrast, resolve hours away from any given trading observation, so the same 1% return is a substantial fraction of the cumulative move that determines the outcome. The placebo here too is statistically insignificant.

Overall, Tables 11 and 12 deliver a consistent message: Polymarket prices actively incorporate equity-market movements, at both the index and the individual-stock level. The economic logic is straightforward. For these mechanically linked contracts, the price of the underlying is a continuously updated, highly liquid public signal about precisely the state the contract is written on, so each move in the stock is direct news about the probability that the contract resolves YES. Profit-motivated traders have every incentive to act on that news

at once, which is why the implied probability adjusts within the hour rather than with a delay. That the response is contemporaneous and one-hour-lagged, scales with how close the contract is to resolution, and disappears for the future-return placebo is exactly the pattern we would expect if Polymarket is efficiently aggregating public financial information rather than drifting on noise.

5.2 Stocks react to Polymarket

We now ask the reverse question. Does the equity market react to Polymarket price drift? That such a channel plausibly exists is no longer hypothetical: as noted in Section 3.2, by late 2025 both Google Finance and the Bloomberg Terminal had begun displaying Polymarket-implied probabilities alongside conventional asset quotes (Hitchcock, 2025). Equity investors can therefore watch Polymarket move in real time on the same screens they use to trade, which gives a concrete conduit through which information first revealed on Polymarket could feed into their decisions and, in turn, into stock prices.

Establishing this relation empirically is harder because, within US trading hours, the two prices update simultaneously: any specification that regresses an intraday stock return on a contemporaneous or lagged Polymarket price change runs into the obvious endogeneity problem that both sides are reacting to the same common information.

The cleanest way around this problem is to exploit the gap in equity trading hours. The U.S. equity market trades from 13:30 to 20:00 UTC each weekday, while Polymarket trades around the clock 7 days a week. In the overnight window, from the previous trading day's close at 20:00 UTC to the next trading day's open at 13:30 UTC, the equity market is mechanically closed but Polymarket continues to absorb news. Any Polymarket price movement during this window is, by construction, generated by information events that the cash equity market has not yet had the chance to price. If the cash market subsequently responds at the next-morning open, that response cannot be running through reverse causation from the cash price itself, but rather through the Polymarket information that arrived overnight.

We test this on the same firm-level panel as above, restricted to (market × trading-day) pairs for which we observe both a Polymarket trade in the prior-day close-window (19:30–23:30 UTC) and another in the next-day open-window (12:30–15:30 UTC). For each such pair we define the overnight Polymarket move $\Delta p^{\text{PM, overnight}}_{m,d}$ as the price change between those two trades, sign-aligned with the direction in which the underlying stock would have to move for the PM contract to resolve YES. The corresponding overnight stock return $r^{\text{stk, overnight}}_{m,d+1}$ is the log return on the underlying stock from the 19:30 UTC close-hour bar on day d to the 13:30 UTC open-hour bar on day $d+1$. We regress the latter on the former:

$$r^{\text{stk, overnight}}_{m,d+1} = \alpha_m + \gamma \Delta p^{\text{PM, overnight}}_{m,d} + \delta_1 r^{\text{stk, intraday}}_{m,d} + \delta_2 r^{\text{stk, overnight}}_{m,d} + \varepsilon_{m,d+1}, \qquad (17)$$

where the coefficient of interest is γ, α_m is a contract fixed effect, and standard errors are

clustered by contract. Column (1) sets $\delta_1 = \delta_2 = 0$; column (2) adds the previous trading day's last intraday hourly return $r^{\text{stk, intraday}}_{m,d}$ and the previous overnight return $r^{\text{stk, overnight}}_{m,d}$, the latter absorbing the AR(1) of the dependent variable. The two restrictions together leave us with 2,242 market-day pairs in column (1) and 893 in column (2).

	(1)	(2)
	$r^{\text{stk, overnight}}_{t+1}$	$r^{\text{stk, overnight}}_{t+1}$
$\Delta p^{\text{PM, overnight}}$	0.076^{***}	0.076^{***}
	(0.013)	(0.021)
$r^{\text{stk, intraday}}_{d}$		-0.476^{*}
		(0.245)
$r^{\text{stk, overnight}}_{t}$		-0.219^{***}
		(0.049)
Observations	2,242	893
Within-R^2	0.071	0.096
Contract FE	Yes	Yes

Table 13: Stocks react to Polymarket in the overnight window. The dependent variable is the log return on the underlying stock from the 19:30 UTC close-hour bar on day t to the 13:30 UTC open-hour bar on day $t+1$. The regressor of interest is the Polymarket YES-probability change over the same overnight window, sign-aligned with the direction in which the underlying stock would have to move for the contract to resolve YES. Column (2) adds the prior day's last intraday hourly return and the prior overnight return. Contract fixed effects throughout; standard errors clustered by contract in parentheses.

The coefficient on the overnight Polymarket move is $\hat{\gamma} = 0.076$ and is essentially identical across columns ($t = 5.65$ unconditionally, $t = 3.55$ after controlling for the autocorrelation of overnight returns and the prior intraday signal). In economic terms, a 10-percentage-point overnight rise in the Polymarket-implied probability is followed by an approximately 0.76% open-to-prior-close gap return the next morning, and the implied within-contract R^2 of 7–10% is large by overnight-predictability standards.

Three caveats temper the causal reading. The closed-market window shuts down contemporaneous reverse causation but not common overnight news: the Polymarket move and the next-morning gap may both respond to the same information, so what the test establishes cleanly is that the news reaches Polymarket *before* the equity open, the informational-connection claim we make, rather than a mechanical channel in which the cash market reads the contract price. Relatedly, we do not control for other overnight venues such as index and single-stock futures or ADRs, so we claim only that Polymarket carries information the cash market has not yet priced, not that it is incrementally informative beyond them. Finally, conditioning on a trade in both the close and open windows selects toward active, news-driven contracts and leaves a modest sample, so the magnitude should be read with that selection

in mind.

With these qualifications, the two directions together show that the information flow between Polymarket and the U.S. equity market is two-sided: Section 5.1 showed that Polymarket absorbs intraday stock returns within an hour or two, and the overnight result here shows that the equity market in turn prices information that appears first on Polymarket. Both directions are present in the data, the precondition for what we do next.

6 Beliefs and the Equity Market

Having established the integration of Polymarket and equity markets in Section 5, we now examine how the two belief objects constructed in Section 4, investor disagreement σ_t and consensus uncertainty Ψ_t, enter equity prices. We organise the analysis around the two margins of the return distribution. We first ask what forecasts the *first moment*, next-day stock returns, and then ask what forecasts the *second moment and trading activity*, next-day realised volatility and volume.

6.1 Disagreement and next-day returns

6.1.1 Main result: σ predicts next-day stock returns

In this section, we examine whether investor beliefs are priced in stock returns. The headline specification is a daily panel regression of the daily return on each individual U.S. stock on the previous day's aggregate Polymarket disagreement and uncertainty indices:

$$r_{i,t+1} = \alpha_i + \beta\sigma_t + \gamma\Psi_t + X'_{i,t}\delta + \varepsilon_{i,t+1}, \tag{18}$$

where i indexes stocks, t indexes trading days, α_i is a ticker fixed effect, and σ_t and Ψ_t are time-series variables matched to every stock on that date. Standard errors are double-clustered by ticker and by date.

The vector $X_{i,t}$ collects a comprehensive suite of controls designed to isolate the predictive power of Polymarket disagreement from known risk factors and technical price dynamics. We first account for firm-level characteristics and momentum by including lagged own returns ($r_{i,t}$) to absorb short-horizon technical effects, alongside log market capitalization to control for the size effect. To ensure that the disagreement signal is not proxying for exposure to standard priced risk factors, we control for the Fama and French (2015) five factors: the market excess return (Mkt – RF), size (SMB), value (HML), profitability (RMW), and investment (CMA), each entered with a one-day lag in decimal units. The market factor Mkt – RF serves as our market-return control. To capture the broader macro-financial environment, we include the one-year risk-free rate, the 10y–3m term spread, and the BAA–AAA credit

spread, which proxy for the opportunity cost of capital and the market's pricing of default risk.

We further address the potential for our measure to overlap with existing proxies for uncertainty and systematic shocks. We include the VIX and the EPU index of Baker et al. (2016b) to ensure that σ captures a unique "state-risk" signal rather than merely mirroring equity-market fear or news-based indices. We additionally control for analyst forecast dispersion, the canonical firm-level disagreement proxy, constructed following Diether et al. (2002). At monthly snapshots we take, for each firm, the most recent outstanding I/B/E/S FY1 earnings-per-share estimate per analyst (requiring at least two analysts within a 180-day staleness window), compute the standard deviation of those estimates scaled by the share price, and winsorise at the 1st and 99th percentiles. We then aggregate to a market-wide daily series as the cross-firm median and enter it with a one-day lag. Including it ensures that σ is not merely repackaging the dispersion of professional analysts' earnings forecasts. We also include day-of-week dummies to remove persistent calendar effects and regular weekday patterns in return distributions.

Finally, we include Polymarket-specific microstructure controls, namely the log of dollar volume and the log of active wallets. Since σ is a dollar-weighted dispersion measure, it could mechanically co-vary with trading intensity. By absorbing these margins, we ensure that the coefficient β identifies the intensity of disagreement—the actual divergence in informed belief—rather than simply the breadth of market activity or a retail-driven hype cycle.

The main result is reported in Table 14. A one-standard-deviation rise in σ_t is associated with a 0.07 percentage point higher next-day return on the underlying stock. The coefficient is positive in column (1), remains positive in the horse race in column (3), and is statistically significant at the five-percent level in both cases. The coefficient on the entropy index Ψ_t is also positive, but it is smaller and statistically indistinguishable from zero. Thus both Polymarket uncertainty measures point in the same economic direction, but the precise and robust return-predictive content comes from the disagreement measure σ. Most significantly for our identification of state risk, the σ coefficient remains positive and statistically significant even after directly controlling for the VIX, the EPU index of Baker et al. (2016b), and analyst forecast disagreement of Diether et al. (2002). This suggests that Polymarket disagreement offers a distinct window into the pricing of states that is not subsumed by traditional measures.

The positive sign differs from the canonical negative disagreement-return relation in Miller (1977) and Diether et al. (2002). In that literature, disagreement is usually disagreement about the stock itself, and it is often studied together with trading frictions such as short-sale constraints. When pessimistic investors cannot fully express their views, prices mainly reflect the beliefs of the more optimistic side. The stock is then expensive today, so its subsequent return is lower. Our measure, however, is different. Polymarket disagreement is not primarily dispersion about the stock's own cash flow. It is disagreement about macro

states, such as politics, policy, geopolitical risk, and other aggregate events, that investors may price in equities. This makes our interpretation closer to Atmaz and Basak (2018): higher belief dispersion can represent extra state uncertainty faced by risk-averse investors, and investors require compensation for bearing that risk.

The out-of-sample results confirm that this predictability is not an artefact of in-sample over-fitting. Following Clark and West (2006), we evaluate each forecasting model out of sample against the controls-only benchmark. We split the sample chronologically, estimating each model on the first 60% of trading days and evaluating it on the remaining 40%. Relative to the controls-only forecast $\hat{y}^b$, the Clark–West measure replaces the squared forecast error with the MSPE-adjusted loss

$$f_{i,t} = (y_{i,t} - \hat{y}^b_{i,t})^2 - (y_{i,t} - \hat{y}_{i,t})^2 + (\hat{y}^b_{i,t} - \hat{y}_{i,t})^2,$$

which corrects for the additional estimation noise the larger model carries under the null, and reports $R^2_{\text{CW}} = \overline{f} \,/\, \overline{(y - \hat{y}^b)^2}$. Because σ_t and Ψ_t are common daily series, the same daily forecast error is replicated across the cross-section of stocks, and treating those replicas as independent would mechanically inflate the t-statistic. We therefore collapse $f_{i,t}$ to one equal-weighted observation per trading day and regress this daily loss-differential on a constant with Newey–West (HAC) standard errors. The reported one-sided t-statistic thus has an effective sample size equal to the number of out-of-sample trading days ($N_{\text{eff}} \approx 160$) rather than the far larger stock-day count, a conservative choice that keeps the statistics modest even where the predictive improvement is strong. The large positive Clark–West t-statistics in the disagreement and joint specifications (columns (1) and (3)) show that adding σ_t delivers a statistically significant gain in forecast accuracy, whereas the uncertainty-only specification (column (2)) carries a negative statistic: Ψ_t does not improve return forecasts on its own. It is the disagreement signal that carries genuine, exploitable predictive content for next-day returns outside the estimation window.

The control variables behave broadly in line with standard empirical regularities. Within our price dynamics and firm characteristic categories, the own-return lags absorb short-horizon reversal and momentum, while the negative loading on market capitalization reflects the higher average returns typically earned by smaller firms in high-frequency panels. The macro-financial controls together with the Fama–French five factors enter with conventional signs: the value factor (HML) loads positively and significantly, while the remaining factors are imprecisely estimated at the daily frequency. Furthermore, by including Polymarket microstructure controls, we ensure that β reflects the actual intensity of belief divergence rather than mechanical fluctuations in market activity. Finally, day-of-week dummies remove regular calendar effects in return distributions. Conditional on this full set of controls, Table 14 establishes that only the disagreement measure, σ, possesses precise and robust predictive content for next-day equity returns.

Variable	(1)	(2)	(3)
σ	$+0.0718^{**}$	–	$+0.0724^{**}$
	(0.0290)	–	(0.0291)
Ψ	–	+0.0044	+0.0046
	–	(0.0033)	(0.0032)
Own return (1)	−0.0277	−0.0277	−0.0277
	(0.0178)	(0.0178)	(0.0178)
Own return (2)	+0.0111	+0.0109	+0.0103
	(0.0141)	(0.0143)	(0.0141)
Log market cap	-0.0061^{***}	-0.0061^{***}	-0.0061^{***}
	(0.0015)	(0.0016)	(0.0015)
VIX	+0.0004	+0.0003	+0.0003
	(0.0003)	(0.0003)	(0.0003)
Policy uncertainty	+0.0000	+0.0000	+0.0000
	(0.0000)	(0.0000)	(0.0000)
Analyst dispersion	+0.8965	+1.0493	+0.8095
	(1.7271)	(1.7063)	(1.7228)
Term spread	$+0.0105^{**}$	$+0.0082^{*}$	$+0.0110^{**}$
	(0.0049)	(0.0044)	(0.0048)
Credit spread	−0.0099	−0.0080	−0.0102
	(0.0062)	(0.0060)	(0.0062)
Log PM volume	+0.0006	+0.0013	+0.0003
	(0.0009)	(0.0008)	(0.0008)
Log PM wallets	-0.0039^{**}	−0.0024	-0.0033^{*}
	(0.0019)	(0.0018)	(0.0019)
Mkt−RF	+0.0239	+0.0201	+0.0221
	(0.0992)	(0.0991)	(0.0984)
SMB	+0.0205	+0.0288	+0.0182
	(0.0907)	(0.0907)	(0.0897)
HML	$+0.1650^{*}$	$+0.1771^{**}$	$+0.1679^{**}$
	(0.0847)	(0.0847)	(0.0844)
RMW	+0.0149	+0.0196	+0.0160
	(0.1086)	(0.1094)	(0.1085)
CMA	−0.0413	−0.0424	−0.0429
	(0.0835)	(0.0816)	(0.0826)
Fixed effects	Ticker	Ticker	Ticker
DOW dummies	Yes	Yes	Yes
R^2 (in-sample)	0.0172	0.0156	0.0179
N	260,249	260,249	260,249
$R^2_{OS,CW}$ (vs ctrls-only)	+0.0344	−0.0207	+0.0131
CW t-stat (vs ctrls-only)	$+3.936^{***}$	−3.658	$+2.762^{***}$

Table 14: Disagreement, entropy, and next-day stock returns. The dependent variable is the next-day $t+1$ log return. All right-hand-side variables are measured at t. In-sample specifications are estimated with ticker fixed effects and two-way clustered standard errors by ticker and date. All specifications include day-of-week dummies, with Monday as the base. Out of sample tests are 60/40 chronological train/test split; OLS with entity dummies. CW t-statistics are one-sided HAC statistics. ***p<0.01, **p<0.05, *p<0.10.

6.1.2 The political sub-index drives the result

The aggregate disagreement index pools events of very different economic character, so a natural next step is to ask which categories of events drive its predictive power. We decompose the aggregate index into the non-overlapping topic categories introduced in Section 3.1: political (σ_t^{pol}), economic (σ_t^{econ}), and geopolitical (σ_t^{geo}). We then examine whether the return premium is common to all of them or concentrated in particular event types. This heterogeneity is of independent interest: it reveals which macro states investors actually price and sharpens the interpretation of σ as a state-risk signal.

These sub-indices are constructed by restricting the transaction tape to events of the relevant type and applying the same dollar-weighted dispersion calculation used for the aggregate measure. We apply an identical topic-specific split to the entropy index (Ψ_t), allowing us to run a "horse race" between the different sources of macro-uncertainty.

Table 15 reports these results. The findings confirm that the return-predictive content is heavily concentrated in political disagreement. In the topic-specific regressions (Columns 1–4), the coefficient on σ_t^{pol} is positive and significant at the five-percent level; economic disagreement is positive and only marginally significant, while financial and geopolitical disagreement are indistinguishable from zero. This hierarchy persists in the joint specification (Column 5).

We identify two primary reasons why political disagreement serves as the dominant predictive channel.

Mechanically, the result reflects the current maturity of the Polymarket ecosystem. During our sample period, political markets were the most liquid and heavily traded segments of the platform. Because our index is dollar-volume weighted, the political component is inherently measured with the highest precision. While Polymarket has recently expanded its coverage of explicitly financial contracts, such as gold prices, crude oil, and corporate earnings, these markets are still developing the liquidity depth required to match the signal-to-noise ratio of the political vertical.

Economically, the dominance of σ^{pol} suggests that prediction markets are most valuable when they capture *unspanned risks*. Standard macroeconomic risks, such as interest-rate shifts or inflation prints, are already traded in deep and relatively complete financial markets (e.g., Treasuries, Fed Funds futures, and index options). Following the logic of Grossman and Stiglitz (1980b), information regarding these events is likely impounded into equity prices via traditional derivatives before it manifests in prediction market volume.

Political risk, which comprises election outcomes, regulatory pivots, and policy implementation, is fundamentally different. These states are not cleanly spanned by standard financial assets, creating an informational void that Polymarket is uniquely positioned to fill. Consequently, disagreement in political markets provides the clearest incremental signal for the pricing of state risk, as it aggregates information that has not yet been fully digested by

the broader financial complex.

Variable	Politics	Finance	Geopolitics	Economy	Horse race
σ^{pol}	+0.0521**	–	–	–	+0.0646**
	(0.0265)	–	–	–	(0.0290)
Ψ^{pol}	+0.0052*	–	–	–	+0.0050
	(0.0031)	–	–	–	(0.0035)
σ^{fin}	–	+0.0055	–	–	+0.0074
	–	(0.0194)	–	–	(0.0190)
Ψ^{fin}	–	−0.0003	–	–	+0.0001
	–	(0.0032)	–	–	(0.0029)
σ^{geo}	–	–	+0.0143	–	+0.0104
	–	–	(0.0135)	–	(0.0136)
Ψ^{geo}	–	–	−0.0022	–	−0.0051
	–	–	(0.0030)	–	(0.0032)
σ^{econ}	–	–	–	+0.0436*	+0.0426*
	–	–	–	(0.0223)	(0.0230)
Ψ^{econ}	–	–	–	+0.0023	+0.0007
	–	–	–	(0.0051)	(0.0056)
Controls	Yes	Yes	Yes	Yes	Yes
Fixed effects	Ticker	Ticker	Ticker	Ticker	Ticker
R^2	0.0173	0.0150	0.0155	0.0164	0.0198
N	260,249	260,249	260,249	260,249	260,249

Table 15: Disagreement, uncertainty, and next-day stock returns by event topic. The dependent variable is the next-day log return. Columns (1)–(4) include the topic-specific σ and Ψ measures separately; column (5) includes all topic-specific measures jointly. All specifications include the same control set as Table 14, ticker fixed effects, and day-of-week dummies. Two-way clustered standard errors are reported in parentheses. ***p<0.01, **p<0.05, *p<0.10.

6.1.3 Placebo: price dispersion is not disagreement

Our measure is built from the cross-*wallet* distribution of transaction prices, mathematically it is closely related to the volatility of Polymarket prices. A natural concern is that σ_t might capture nothing more than how much Polymarket prices move around. Any effect then becomes a mechanical price-dispersion effect rather than genuine disagreement among traders. We test this directly with a placebo: we replace σ_t in the headline slot with the within-market standard deviation of the traded price, $\text{SD}(p)_t$, a pure price-dispersion measure, and re-estimate the headline return specification with the identical control set.

Table 16 repeats the headline return specification of Table 14 with $\text{SD}(p)_t$ entered in place of σ_t. The contrast is stark. Where wallet disagreement loaded at +0.0724 ($p < 0.05$), price dispersion is statistically indistinguishable from zero, both on its own (+0.0647, column (1)) and alongside Ψ_t (+0.1442, column (3)), and it does not improve in-sample fit. The return

predictability therefore comes from disagreement across traders, not from price variation per se.

Variable	(1)	(2)	(3)
SD(p)	+0.0647	–	+0.1442
	(0.2492)	–	(0.2551)
Ψ	–	+0.0044	+0.0047
	–	(0.0033)	(0.0034)
Fixed effects	Ticker	Ticker	Ticker
Controls (V2)	Yes	Yes	Yes
DOW dummies	Yes	Yes	Yes
R^2 (in-sample)	0.0150	0.0156	0.0157
N	260,249	260,249	260,249

Table 16: Placebo: price dispersion in σ's slot. The headline next-day return specification of Table 14 is re-estimated with the within-market price standard deviation SD$(p)_t$ replacing wallet disagreement σ_t. Column (1) includes price dispersion only, column (2) Ψ_t only, and column (3) both. All specifications include the same control set as Table 14, ticker fixed effects, and day-of-week dummies. Two-way clustered standard errors are reported in parentheses. ***p<0.01, **p<0.05, *p<0.10.

6.1.4 Robustness: a Fama–French portfolio panel

The headline panel uses a cross-section of individual U.S. S&P 500 stocks. A natural concern is that these are all large, highly liquid, heavily arbitraged firms, so that a disagreement premium estimated on them need not extend to the broader cross-section of smaller, less liquid, or value and low-profitability stocks. To show that the premium is not specific to this particular sample, we re-estimate the return regression on a panel of forty Fama–French value-weighted decile portfolios. These portfolios are built from the full CRSP universe of NYSE, AMEX, and NASDAQ common stocks, and so span segments of the market, small caps in particular, that are absent from the S&P 500. The sample contains the ten deciles of each of the size (ME), book-to-market (BE-ME), operating-profitability (OP), and investment (INV) univariate sorts, with portfolio fixed effects and Driscoll and Kraay (1998) standard errors (which are robust to the cross-sectional dependence induced by the common daily regressors σ_t and Ψ_t).

Table 17 shows that the result survives intact. The loading on σ_t is +0.0870 ($p < 0.05$) in the joint specification, if anything slightly larger than the individual-stock estimate, and the out-of-sample Clark–West statistic is large and positive. Uncertainty Ψ_t again carries only a marginal loading. Moreover, the effect is pervasive across the cross-section rather than concentrated in a corner of it: Figure 9 plots the per-decile σ coefficient from separate time-series regressions and shows that it is positive across every decile of all four sorts.

Variable	(1)	(2)	(3)
σ	$+0.0859^{**}$	–	$+0.0870^{**}$
	(0.0402)	–	(0.0393)
Ψ	–	+0.0069	$+0.0070^{*}$
	–	(0.0044)	(0.0042)
Fixed effects	Portfolio	Portfolio	Portfolio
Controls (V2)	Yes	Yes	Yes
R^2 (in-sample)	0.0748	0.0704	0.0790
N	16,000	16,000	16,000
$R^2_{OS,CW}$ (vs ctrls-only)	+0.1560	−0.0916	+0.0615
CW t-stat (vs ctrls-only)	$+12.387^{***}$	−11.317	$+7.585^{***}$

Table 17: Disagreement and next-day returns in a Fama–French portfolio panel. The dependent variable is the next-day return on each value-weighted decile portfolio (ME, BE-ME, OP, and INV sorts; forty portfolios). Portfolio fixed effects; Driscoll–Kraay standard errors in parentheses. $R^2_{OS,CW}$ and the Clark–West statistics are computed relative to the controls-only benchmark. ***p<0.01, **p<0.05, *p<0.10.

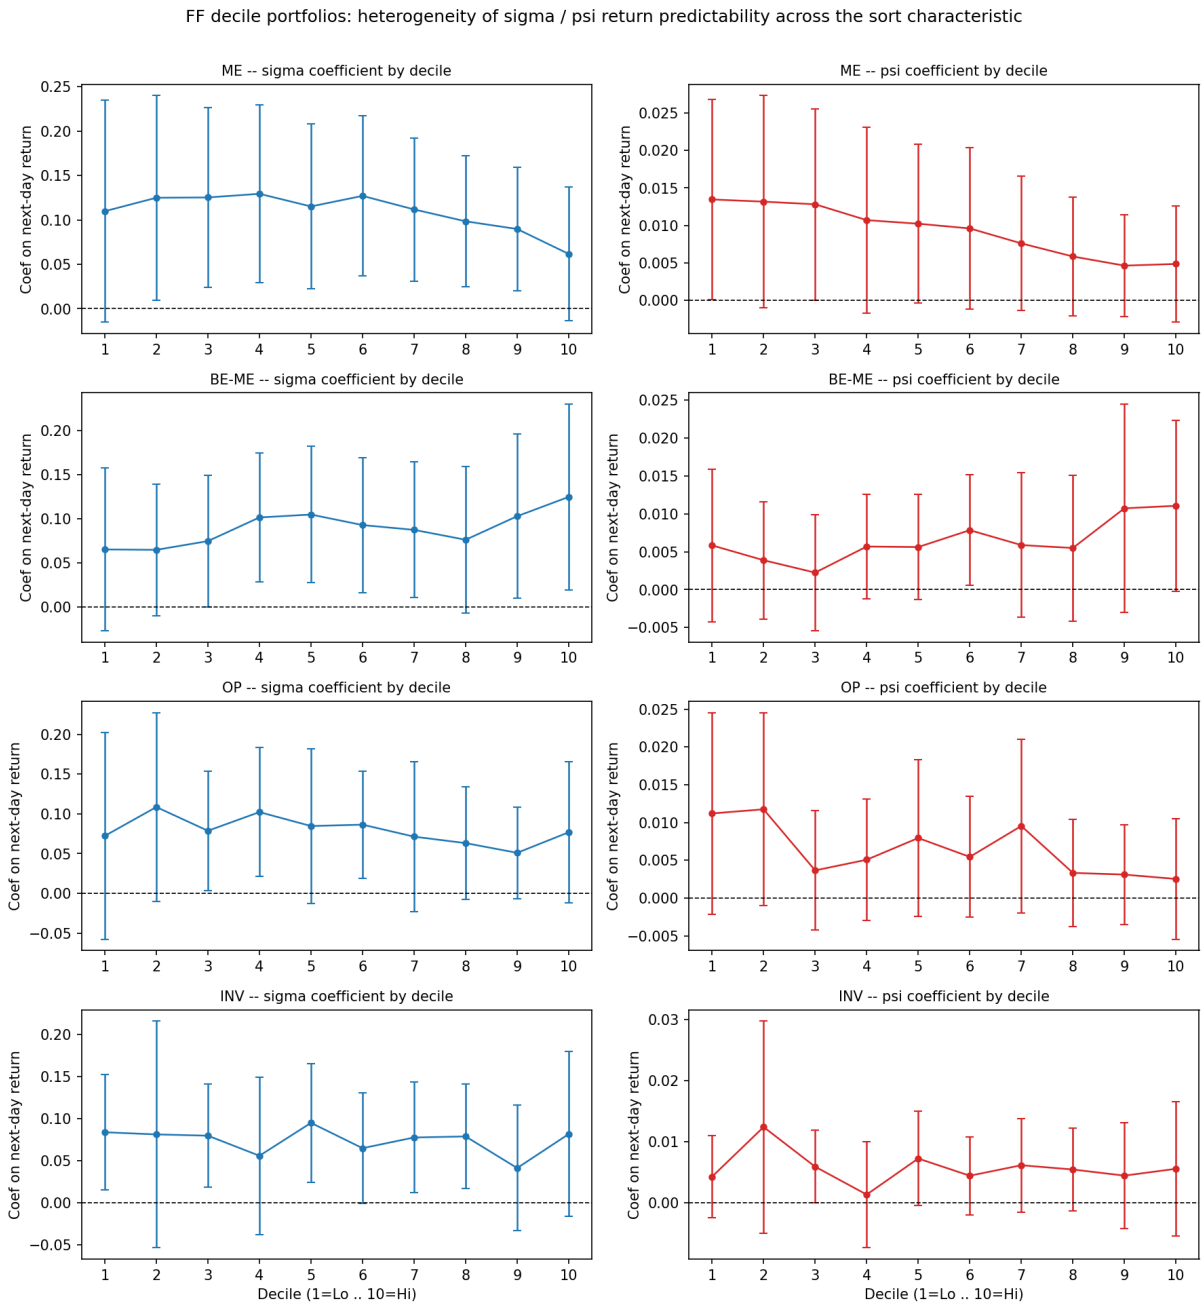


Figure 9: Per-decile disagreement loading across Fama–French sorts. Each point is the coefficient on σ_t (with Ψ_t included) from a separate time-series regression of a single value-weighted decile portfolio's next-day return, for the size (ME), book-to-market (BE-ME), operating-profitability (OP), and investment (INV) sorts. The disagreement loading is positive across the entire cross-section.

6.2 Uncertainty, volatility, and volume

We next examine whether the return premium associated with Polymarket beliefs is accompanied by immediate equity-market activity.

We first estimate the baseline specification in Columns (1) and (2) of Table 18, using the

same configuration as our core return regression:

$$y_{i,t+1} = \alpha_i + \beta \sigma_t + \gamma \Psi_t + X'_{i,t}\theta + \varepsilon_{i,t+1}. \tag{19}$$

The empirical results suggest that uncertainty, rather than investor disagreement, predicts equity-market activity. In the baseline volume and volatility specifications, high Ψ_t predicts both lower stock trading volume and lower realized volatility on the subsequent day, while the unconditional coefficient on σ_t is small and insignificant. This is already informative. If the Polymarket signal were simply spilling into equities as stock-level disagreement, one would expect more trading and more volatility. Instead, high entropy is associated with muted equity-market activity. This suggests that Ψ_t captures a state in which macropolitical information is important but not yet sufficiently resolved to be traded aggressively in the stock market.

A natural concern is that the volatility result may simply restate the volume result: if high entropy lowers trading, then realized volatility may fall mechanically. We therefore add contemporaneous stock trading volume directly to the volatility regression in Column (3). Volume enters positively, as expected, but the coefficient on Ψ_t remains negative and highly significant. Thus, Polymarket entropy is not only a proxy for lower trading intensity. It carries independent information about the timing of equity repricing. When the prediction-market state distribution is diffuse, equity investors appear to delay repricing even after controlling for the normal relation between trading activity and volatility.

We test this mechanism more directly by adding the change in uncertainty, measured as

$$\Delta\Psi_{t+1} = \Psi_{t+1} - \Psi_t.$$

As time moves closer to resolution, entropy should naturally decline. The coefficient in Column (4) is negative and significant. Since a negative $\Delta\Psi_{t+1}$ means that prediction-market entropy falls, the result implies that equity volatility rises when Polymarket uncertainty resolves. This completes the interpretation: volatility is not highest when uncertainty is most diffuse; it is released when the state distribution becomes clearer and can be incorporated into equity prices.

Consistent with the return evidence, disagreement σ_t does not predict volatility unconditionally: its coefficient is small and insignificant in every column. Disagreement matters for expected returns because it compensates for bearing aggregate state risk, whereas the timing of realized volatility is governed by the resolution of uncertainty Ψ_t, not by the level of disagreement.

Overall, entropy Ψ_t measures the unresolvedness of the state distribution and governs the timing of repricing. High entropy suppresses equity activity and realized volatility; falling entropy releases volatility. Prediction markets therefore do not only reveal whether aggregate state risk is priced, but also when that risk is likely to be incorporated into equity prices.

	(1)	(2)	(3)	(4)
Dep. var.	Volume	Volatility	Volatility	Volatility
σ_t	−0.2884	−0.0037	−0.0018	−0.0052
	(0.5285)	(0.0137)	(0.0129)	(0.0129)
Ψ_t	−0.1415*	−0.0076***	−0.0068***	−0.0079***
	(0.0752)	(0.0020)	(0.0019)	(0.0018)
$\Delta\Psi_{t+1}$	–	–	–	−0.0106**
	–	–	–	(0.0047)
Volume (contemp.)	–	–	+0.0034***	+0.0034***
	–	–	(0.0010)	(0.0010)
Fixed effects	Ticker	Ticker	Ticker	Ticker
Controls	Yes	Yes	Yes	Yes
N	260,364	260,364	260,364	260,364
R^2 (in-sample)	0.2144	0.2175	0.2397	0.2408
$R^2_{OS,CW}$	+0.0045	+0.0300	+0.0256	+0.0233
CW t-stat	+2.511***	+3.872***	+3.828***	+3.996***

Table 18: Polymarket disagreement, uncertainty, and subsequent equity market activity. The dependent variable is log trading volume in column (1) and realised volatility in columns (2)–(4). All specifications include the same control set as Table 14, ticker fixed effects, and day-of-week dummies. Two-way clustered standard errors are reported in parentheses. $R^2_{OS,CW}$ and the Clark–West statistics are computed relative to the controls-only benchmark. ***p<0.01, **p<0.05, *p<0.10.

These timing patterns invite an interpretation in terms of *ambiguity aversion*. The events that move Ψ_t, such as elections, central-bank decisions, and geopolitical flashpoints, confront investors with genuine Knightian uncertainty about which state will realise, rather than with a known probability distribution over outcomes (Ellsberg, 1961; Gilboa and Schmeidler, 1989; Fox and Tversky, 1995; Machina, 2014). An ambiguity-averse investor evaluates uncertain payoffs under a pessimistic belief drawn from a set of priors, and therefore reacts to ambiguous macro information very differently from the way she reacts to ordinary, well-specified risk. Because Ψ_t is the entropy of the price-implied state distribution, it is a direct measure of how much of this ambiguity is still unresolved on a given day, which lets us confront the predictions of that literature with data rather than inferring ambiguity from indirect proxies.

Three features of Table 18 align with this mechanism. First, high Ψ_t predicts lower equity trading volume: when the macro state is ambiguous, ambiguity-averse investors optimally refrain from trading, the portfolio-inertia result of Illeditsch (2011), so turnover in the exposed stocks falls. Second, high Ψ_t predicts lower contemporaneous realised volatility, and continues to do so after controlling for volume, because ambiguous signals are heavily discounted and prices adjust only sluggishly while the state is opaque, echoing Epstein and Schneider (2008), in whose model ambiguity about the quality of information makes prices

less responsive to news. Third, and most tellingly, a fall in Ψ_t predicts a subsequent rise in volatility: once the ambiguity resolves, the adjustment that investors had deferred is finally made and the accumulated repricing is released as a burst of realised volatility, the excess-volatility-upon-resolution prediction of Illeditsch (2011). That disagreement σ_t shows no comparable volatility effect fits the same dichotomy: σ_t is a priced risk that compensates investors in the first moment of returns, whereas Ψ_t is ambiguity that governs the timing of the second moment. Polymarket entropy thus offers a rare high-frequency, market-based gauge of the ambiguity that asset-pricing theory has long held should matter but has found difficult to observe.

6.2.1 Which uncertainty moves the second moment?

Just as the return result was concentrated in one topic, we can ask which kind of uncertainty drives the volatility effect. We split the overall Ψ_t (and σ_t) into the four topic-specific indices, estimating each separately and then jointly, exactly as for returns in Table 15. Table 19 reports the result for next-day realised volatility, and it sharpens the interpretation in two ways.

First, the suppression of equity volatility is a *political* uncertainty phenomenon. Political Ψ carries essentially all of the negative loading ($\Psi^{\text{pol}} = -0.0085$, $p < 0.01$, in the joint specification), while financial and economic uncertainty are insignificant. This fits the wait-and-see reading: when political resolution (elections, policy decisions) is pending, equity investors delay repricing until the state is known. (The same political concentration appears on the volume margin in Table 18.)

Second, geopolitical uncertainty is the exception that proves the rule. Ψ^{geo} enters with the *opposite* sign ($+0.0038$, $p < 0.05$): geopolitical risk is not something investors can wait out, so it transmits into immediate equity volatility rather than suppressing it. Finally, the disagreement loadings confirm the cross-channel separation that the topic σ terms are insignificant (the lone exception being a small σ^{econ} loading), consistent with σ being a first-moment, return-pricing signal rather than a driver of volatility.

Variable	Politics	Finance	Geopolitics	Economy	Horse race
σ^{pol}	−0.0051	–	–	–	−0.0183
	(0.0126)	–	–	–	(0.0139)
Ψ^{pol}	−0.0072***	–	–	–	−0.0085***
	(0.0019)	–	–	–	(0.0021)
σ^{fin}	–	+0.0062	–	–	+0.0027
	–	(0.0125)	–	–	(0.0122)
Ψ^{fin}	–	+0.0000	–	–	−0.0007
	–	(0.0020)	–	–	(0.0018)
σ^{geo}	–	–	−0.0025	–	+0.0021
	–	–	(0.0081)	–	(0.0082)
Ψ^{geo}	–	–	+0.0036**	–	+0.0038**
	–	–	(0.0016)	–	(0.0017)
σ^{econ}	–	–	–	+0.0181	+0.0234**
	–	–	–	(0.0111)	(0.0117)
Ψ^{econ}	–	–	–	+0.0003	+0.0043
	–	–	–	(0.0026)	(0.0028)
Controls	Yes	Yes	Yes	Yes	Yes
Fixed effects	Ticker	Ticker	Ticker	Ticker	Ticker
R^2	0.2174	0.2141	0.2153	0.2146	0.2199
N	260,364	260,364	260,364	260,364	260,364

Table 19: Uncertainty and next-day realised volatility by event topic. The dependent variable is next-day realised volatility. Columns (1)–(4) include the topic-specific σ and Ψ measures separately; column (5) includes all topic-specific measures jointly. All specifications include the same control set as Table 18, ticker fixed effects, and day-of-week dummies. Two-way clustered standard errors are reported in parentheses. ***p<0.01, **p<0.05, *p<0.10.

7 A Theoretical Framework

The empirical chapters established a set of facts that any explanation must reproduce *jointly*. We restate them compactly as the targets the model must hit:

(F1) **Disagreement premium.** Disagreement σ_t positively predicts subsequent equity returns, and the effect is *permanent*, it does not reverse (Table 14).

(F2) **Uncertainty governs the second moment.** Price-level uncertainty Ψ_t governs the magnitude and timing of volatility: high Ψ_t coincides with *low* near-term equity volume and realised volatility, with the volatility released as Ψ_t falls toward resolution (Table 18).

(F3) **Double dissociation.** The two objects do not cross over: σ_t does not forecast volatility and Ψ_t does not forecast returns.

(F4) **Lifecycle.** Ψ_t declines monotonically over a market's life while σ_t is hump-shaped; the two are negatively related within a market but nearly orthogonal in the aggregate (Table 10).

(F5) **Risk premium, not mispricing.** The σ_t premium is positive, non-reversing, and uniform across arbitrage-cost proxies, and its sign is the *opposite* of the analyst-forecast-dispersion result of Diether et al. (2002).

Facts (F1) and (F5) are enough to select the modelling lineage before any equation is written. Two classical mechanisms generate a disagreement premium, and they differ in sign. In the Miller–Harrison–Kreps–Scheinkman–Xiong channel, disagreement is priced only because binding short-sale constraints exclude pessimists, so prices reflect the optimists alone; the asset is overpriced and its subsequent return is *negative* and *reverses* (Miller, 1977; Harrison and Kreps, 1978; Scheinkman and Xiong, 2003). In the Atmaz and Basak (2018) channel, investors are risk-averse and short-unconstrained, and disagreement is priced as a genuine risk premium that does not reverse. The positive, permanent premium of (F1), and its sign reversal relative to analyst dispersion in (F5), are the signature of the second mechanism and the negation of the first. We therefore develop the leanest Atmaz–Basak economy consistent with the evidence and trace each of (F1)–(F5) to its source within it.

The construction proceeds in two passes. We first solve a static economy with a single binary event: it produces the dissociation transparently and isolates risk aversion as the necessary ingredient, but delivers a disagreement premium whose sign is *state-dependent*—not yet the unconditional premium in the data. We then change a single primitive, moving the object of disagreement from a bounded event to an unbounded growth rate, and show that this removes the state-dependence and yields a positive, unconditional premium in closed form. The lifecycle and resolution dynamics (F2, F4) follow from the learning extension of Section 7.5. A reader interested only in the destination can skip to Section 7.6: *disagreement is priced; uncertainty is traded.*

7.1 A static economy with heterogeneous beliefs

Consider two dates, $t \in \{0,1\}$, and a single binary event that resolves at $t = 1$ in one of two states, $s \in \{H, L\}$. The aggregate endowment (consumption) is $e_H > e_L > 0$, so H is the "good" (rich) state. Three assets trade at $t = 0$: a riskless bond in zero net supply with gross return R_f; a *stock* that pays the aggregate dividend $D_s = e_s$; and a pair of *Arrow securities* that pay one dollar in a single state, with prices ϕ_H and ϕ_L. With both Arrow securities the market is complete, and ϕ_H is, by construction, the price today of a dollar delivered only in state H, the exact analogue of a Polymarket contract on the event. Normalised by the bond price, $\phi_H/(\phi_H+\phi_L)$ is the market's *risk-neutral probability* of H, which is exactly what a Polymarket price reports.

A continuum of investors $i \in [0,1]$ have identical CRRA preferences over date-1 consumption,

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \qquad u'(c) = c^{-\gamma}, \qquad \gamma > 0, \tag{20}$$

and equal initial wealth, normalised to one. Investors differ *only* in their beliefs: investor i assigns probability $\pi_i \in (0,1)$ to state H. The cross-section of beliefs has mean and variance

$$\bar{q} = \int_0^1 \pi_i \, \mathrm{d}i, \qquad \sigma^2 = \int_0^1 (\pi_i - \bar{q})^2 \, \mathrm{d}i, \tag{21}$$

where $\bar{q}$ is the (objective) consensus belief and σ^2 is the *disagreement*, the population analogue of the dispersion index of Section 4.2. We also define the *consensus uncertainty*

$$\Psi \equiv \bar{q}\,(1-\bar{q}) \in [0, \tfrac{1}{4}], \tag{22}$$

the payoff variance of a unit binary claim. Like the price-entropy index of Section 4.1, Ψ is maximal when the consensus is a coin flip ($\bar{q} = \frac{1}{2}$) and zero once the state is known; the two are monotone transforms of one another, the empirical index using Shannon entropy and the theoretical object using its variance form. Crucially, Ψ describes the *first* moment of beliefs (how unsure the average trader is) while σ^2 describes the *second* (how much traders differ). There are no short-sale constraints: a pessimist simply buys the L Arrow security.

Each investor chooses state-contingent consumption to maximise expected utility under *her own* belief, subject to a budget constraint:

$$\max_{c_{iH},\, c_{iL}} \; \pi_i \, u(c_{iH}) + (1-\pi_i)\, u(c_{iL}) \qquad \text{s.t.} \qquad \phi_H c_{iH} + \phi_L c_{iL} = 1. \tag{23}$$

Taking first-order conditions and dividing the H condition by the L condition (the budget multiplier cancels) gives

$$\frac{c_{iH}}{c_{iL}} = \left(\frac{\beta_i}{\Phi}\right)^{1/\gamma}, \qquad \beta_i \equiv \frac{\pi_i}{1-\pi_i}, \qquad \Phi \equiv \frac{\phi_H}{\phi_L}. \tag{24}$$

In words: an optimist (high odds β_i on H) tilts her consumption toward the good state by buying the H claim, and the strength of that tilt is damped by risk aversion through the exponent $1/\gamma$. Imposing market clearing, $\int c_{iH}\,\mathrm{d}i = e_H$ and $\int c_{iL}\,\mathrm{d}i = e_L$, pins down the single unknown price ratio Φ through

$$\frac{e_H}{e_L} = \frac{\displaystyle\int_0^1 a_i \left(\phi_H a_i + \phi_L\right)^{-1} \mathrm{d}i}{\displaystyle\int_0^1 \left(\phi_H a_i + \phi_L\right)^{-1} \mathrm{d}i}, \qquad a_i \equiv (\beta_i/\Phi)^{1/\gamma}. \tag{25}$$

For general beliefs this condition has no elementary closed form and is solved numerically;

the special case of no disagreement ($\pi_i \equiv \bar{q}$) collapses to $\Phi_0 = \left[\bar{q}/(1-\bar{q})\right](e_H/e_L)^{-\gamma}$, which already exposes the two forces in the price of the H claim: it *rises* with the belief in H and *falls* because H is the rich state, where a risk-averse investor values a marginal dollar less. The state price is the value of a marginal dollar delivered in a state rather than the size of the payoff there, and the two move in opposite directions: since $u'(c) = c^{-\gamma}$ declines in consumption, the rich-state claim is the cheaper one. A security that pays in good states is thus a poor hedge and trades at a discount—a positive risk premium—while one that pays in bad states earns an insurance premium. This single property of marginal utility drives every premium that follows.

7.2 Disagreement, the consensus, and the equity premium

To read prices, define the *consensus belief* $\hat{q}$ as the single belief that a representative agent would have to hold to support the equilibrium prices, i.e. the $\hat{q}$ for which the clearing price ratio satisfies

$$\Phi = \frac{\hat{q}}{1-\hat{q}}\left(\frac{e_H}{e_L}\right)^{-\gamma}. \tag{26}$$

All asset prices are then functions of $\hat{q}$ alone:

$$P = \phi_H e_H + \phi_L e_L, \qquad R_f = \frac{1}{\phi_H + \phi_L}, \qquad \mathrm{RP}(\bar{q}, \hat{q}) = \frac{\bar{q}\, e_H + (1-\bar{q})\, e_L}{P} - R_f, \tag{27}$$

where the equity premium RP is evaluated under the objective measure $\bar{q}$. The structure of (27) delivers the observation that organises the rest of the analysis: the consensus $\hat{q}$ is a *sufficient statistic* for prices. Every object—the stock price P, the safe rate R_f, and the premium RP—depends on the entire cross-section of beliefs only through the single scalar $\hat{q}$. Two belief distributions with the same $\hat{q}$ are therefore priced identically, so the market cannot respond to the dispersion σ except through the imprint that dispersion leaves on the consensus. The question of whether disagreement carries a premium thus reduces to a single question: holding the objective mean $\bar{q}$ fixed, does dispersion move $\hat{q}$ away from $\bar{q}$? When investors agree, $\hat{q} = \bar{q}$ identically, so the premium is governed entirely by the gap $\hat{q} - \bar{q}$ that dispersion opens.

Whether that gap opens is a matter of curvature. Each investor's state-H demand $c_{iH}(\pi_i)$ is increasing in her optimism, and the market re-clears at the cross-sectional average demand. Writing $\pi_i = \bar{q} + \varepsilon_i$ with $\int \varepsilon_i \, \mathrm{d}i = 0$, a mean-preserving spread of beliefs leaves average demand—and hence prices—unchanged at first order, because the optimists' excess demand and the pessimists' shortfall cancel ($\int \varepsilon_i \, \mathrm{d}i = 0$). They fail to cancel only at *second* order, and only when $c_{iH}(\cdot)$ is curved: by Jensen's inequality, the average of a curved demand schedule differs from its value at the mean belief. The first surviving term is of order σ^2, which is why the premium loads on the *variance* of beliefs and why the natural expansion is to second order around no disagreement. Under log utility ($\gamma = 1$) demand aggregates

linearly in the belief odds, the curvature vanishes, and dispersion washes out completely:

$$\gamma = 1 \implies \hat{q} = \bar{q}, \qquad \frac{\partial \mathrm{RP}}{\partial \sigma^2} = 0. \tag{28}$$

A disagreement premium is thus a *risk-aversion* phenomenon. For $\gamma > 1$, expanding the clearing condition to second order in the spread gives (Appendix A)

$$\hat{q} - \bar{q} = \tfrac{1}{2}\kappa_\gamma(\bar{q})\,\sigma^2 + o(\sigma^2), \qquad \mathrm{RP} = \mathrm{RP}_0(\bar{q}) + \Lambda(\bar{q})\,\sigma^2 + o(\sigma^2), \tag{29}$$

where $\kappa_\gamma \propto (\gamma - 1)$ (so $\kappa_1 \equiv 0$) and the disagreement loading on the premium satisfies

$$\Lambda(\bar{q}) > 0 \iff \bar{q} < \tfrac{1}{2}, \qquad \Lambda(\bar{q}) < 0 \iff \bar{q} > \tfrac{1}{2}. \tag{30}$$

In words: with risk aversion, dispersion pushes the consensus *away from* $\frac{1}{2}$, toward the more likely state, and the premium inherits a sign that depends on where the consensus sits. This is the Atmaz and Basak (2018) non-monotonicity, reproduced exactly. It already explains the *dissociation on the return side* and, importantly, why our sign is opposite to that of Diether et al. (2002). Because there is no short-sale constraint, pessimists price the L claim and the consensus is not truncated toward optimists; the disagreement effect is a genuine *risk premium* operating through $\hat{q}$, not a Miller-style overpricing that must subsequently reverse (Miller, 1977; Harrison and Kreps, 1978; Scheinkman and Xiong, 2003). Analyst-dispersion proxies capture the latter, mispricing, channel and so predict *low* subsequent returns; the Polymarket measure isolates the former, risk-premium, channel and predicts *high* ones.

A testable refinement. Equation (30) is a prediction, not merely a caveat. The *static* component of the disagreement premium should be *more positive* where the consensus is far from resolution (low $\bar{q}$, high distance-to-certainty) and can even flip sign in "euphoric" high-$\bar{q}$ states. The σ–return relation should therefore vary with the consensus level, a sharper, conditional test than the unconditional regression of Section 6.1, which one can run directly by interacting σ_t with the prevailing consensus level.

7.3 Numerical illustration

We solve the equilibrium (25) exactly, with no approximation, for a continuum of investors with Beta-distributed beliefs and endowments $e_H = 1.10$, $e_L = 0.90$. Figure 10 displays the two halves of the dissociation. In the left panel, the stock's return variance is a clean increasing function of the consensus uncertainty Ψ and is *invariant* to the disagreement σ: the three σ curves lie on top of one another, the numerical counterpart of the identity (34). In the right panel, the consensus bias $\hat{q} - \bar{q}$ moves with disagreement *only* under risk aversion, it is flat at zero under log utility ($\gamma = 1$) and shifts visibly for $\gamma = 5$, the numerical counterpart of (28)–(29). Figure 11 then traces the state-dependent loading $\Lambda(\bar{q})$ of (30): positive when

the good state is unlikely, crossing zero near $\bar{q} \approx \frac{1}{2}$, and steeper for higher risk aversion, the non-monotonicity that the growth-rate formulation of Section 7.4 subsequently removes.

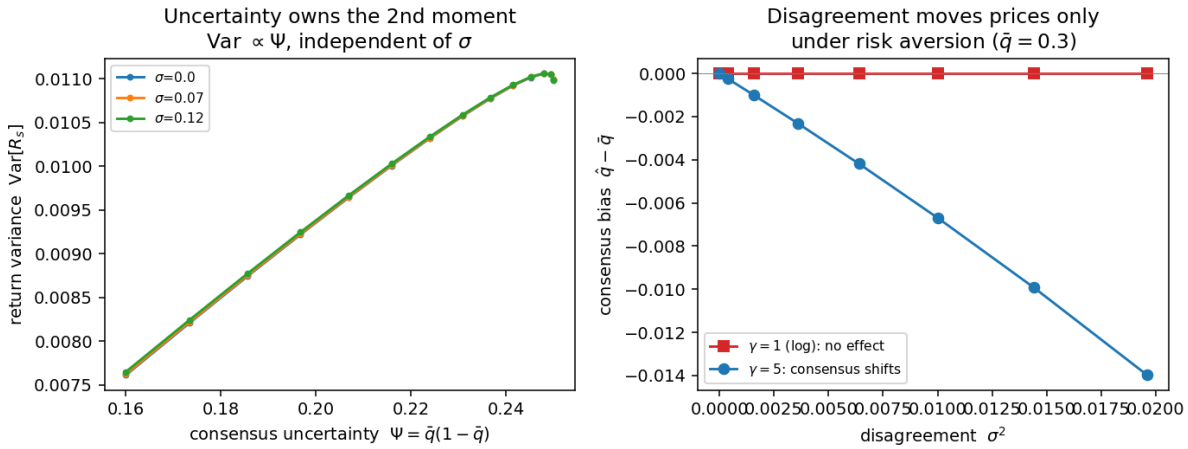


Figure 10: The double dissociation in the binomial economy. Left: stock return variance $\mathrm{Var}[R_s]$ against consensus uncertainty Ψ; the three disagreement curves coincide (variance owns the second moment, $\perp \sigma$). Right: consensus bias $\hat{q} - \bar{q}$ against disagreement σ^2 at $\bar{q} = 0.3$; flat under log utility ($\gamma = 1$) and nonzero under risk aversion ($\gamma = 5$).

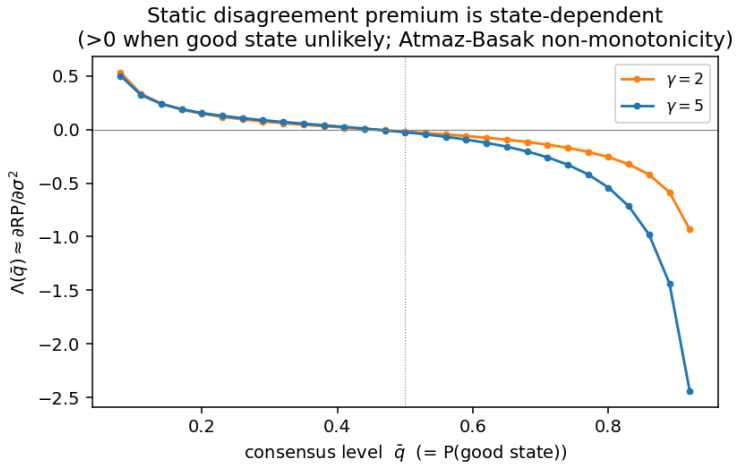


Figure 11: The static disagreement premium is state-dependent. The loading $\Lambda(\bar{q}) \approx \partial \mathrm{RP}/\partial\sigma^2$ is positive when the good state is unlikely and negative when it is likely, the Atmaz and Basak (2018) non-monotonicity. The robust, unconditional premium (F1) is delivered by the growth-rate formulation of Section 7.4.

7.4 From state-dependent to unconditional: growth-rate disagreement

The static loading (30) is the one feature that does not match the data: it changes sign at $\bar{q} = \frac{1}{2}$, whereas the estimated σ_t premium is positive and does not reverse. The reason for the sign flip is that a binary event is *bounded*, "pessimism" is directional, raising the premium

only when the good state is the cheap one. Disagreement about an *unbounded* object is not directional, and this is exactly the fix.

Let aggregate dividends follow a geometric Brownian motion under the objective measure, $\mathrm{d}D_t/D_t = \mu\,\mathrm{d}t + \sigma_D\,\mathrm{d}W_t$, and let the same continuum of CRRA investors now disagree about the *growth rate*: over a horizon the log dividend $X = \log(D_T/D_0)$ is objectively $N(m, s^2)$, while investor i believes $X \sim N(m_i, s^2)$ with the cross-section of perceived means having mean $\bar{m} = m$ and variance v. We identify the disagreement with v (equivalently its standard deviation σ). With complete markets the economy prices through a representative agent holding the wealth-weighted average belief, which is the *mixture* of the individual Gaussians,

$$\bar{f}(X) = \int f_i(X)\,\mathrm{d}i = N\big(m,\ s^2 + v\big). \tag{31}$$

Two forces act on the consensus, and both are *translation-invariant* in m, this is the crucial difference from the binary case. First, dispersed means make the average belief more diffuse: the consensus variance is inflated by the disagreement v, as in (31). Second, under risk aversion ($\gamma > 1$) the wealth-weighted aggregator leans *pessimistic*, shifting the consensus mean growth rate *down*. Neither force references the level of m, unlike the bounded binary event whose entropy peaks at $\frac{1}{2}$ and makes "pessimism" directional. Under log utility the equity premium is then available in closed form (Appendix B),

$$\boxed{\mathrm{RP}_{\log}(v) = e^{s^2 + v/2} - 1}, \qquad \frac{\partial\,\mathrm{RP}_{\log}}{\partial v} = \tfrac{1}{2}\,e^{s^2 + v/2} > 0, \tag{32}$$

which reduces to the homogeneous-belief premium $e^{s^2} - 1$ at $v = 0$ and rises with disagreement for *every* growth state m. The mechanism is most transparent in the safe asset. The bond pays one unit in every state, but that unit is most valuable in a disaster—low realised growth X, where the investor is poor and the pricing kernel e^{-X} is large. A more dispersed consensus (31) carries a fatter left tail, placing more mass exactly where the marginal unit is most prized, so the bond becomes dearer and the riskless rate falls through a precautionary channel, $r_f = \rho + m - \frac{1}{2}(s^2 + v)$; a lower safe rate is, mechanically, a higher premium. Under CRRA the consensus is Gaussian with variance inflated to $\hat{s}^2 = s^2 + v/\gamma$ to leading order, and the premium becomes $\mathrm{RP}(v) = e^{\gamma s^2 + \Theta(\gamma)\,v} - 1 + O(v^2)$ with a loading $\Theta(\gamma) > 0$ for every γ, $\Theta(1) = \frac{1}{2}$, and Θ decreasing in γ. (For $\gamma > 1$ the exact Θ also absorbs the belief-driven redistribution of wealth across agents, beyond the naive inflation v/γ; we pin it numerically below.) In all cases the loading is independent of m:

$$\frac{\partial\,\mathrm{RP}}{\partial v} > 0 \qquad \text{for all } m,\ \gamma \geq 1. \tag{33}$$

The premium is now *unconditional and positive* (the empirical pattern of facts (F1) and (F5)) because the object of disagreement is unbounded and the consensus shift no longer flips

with the state. The static binary model of Section 7.2 is the second-order, sign-dependent shadow of this same risk-premium channel; the growth formulation makes it first-order and unconditional.

From the measured dispersion to the model's object. A natural objection is that the empirical σ_t is built from *bounded* prediction-market contracts, exactly the case in which the static model warns that the premium is state-dependent, while the unconditional result requires disagreement about an *unbounded* growth rate. The reconciliation is that σ_t is not the dispersion of beliefs about any single bounded contract: it is the volume-weighted average of within-market dispersions across the whole cross-section of macro-relevant events active at t. A single event's disagreement is the bounded, directional object that the binary economy prices state-dependently. But the events that enter σ_t—elections, policy decisions, data releases, geopolitical episodes—are jointly the macro states that determine aggregate cash-flow growth, and aggregation across them is precisely the operation that converts a collection of bounded, event-specific disagreements into disagreement about the unbounded aggregate. Event-specific directionality averages out, leaving the common component—how much traders differ about the macro fundamentals as a whole—which is the object v that the growth formulation prices. The headline regressor is therefore the empirical counterpart of v, and the growth model, not the single-event binary model, is the one whose prediction it should match.

This mapping has testable content that sharpens the distinction between the two passes. The binary model predicts that a *single* event's disagreement premium varies with that event's consensus level and flips sign across $\bar{q} = \frac{1}{2}$ (the refinement of Section 7.2); the growth model predicts that the *aggregate* σ_t premium is unconditional, with no such interaction. Interacting σ_t with the prevailing consensus level therefore discriminates between them: a flat, sign-stable loading is evidence that the priced object is aggregate growth disagreement rather than bounded single-event dispersion. We regard this conditional test as the sharp empirical content of the bounded-versus-unbounded distinction and a natural next step.

Numerical confirmation. We solve the complete-market equilibrium exactly, a continuum of Gaussian-dispersed beliefs, no approximation, and verify both the unconditionality and the closed form. Figure 12 plots the disagreement premium $\mathrm{RP}(v) - \mathrm{RP}(0)$: the curves for three growth states $m \in \{-0.05, 0, +0.05\}$ coincide (left panel), so the premium rises with disagreement *regardless of the state*, and it is positive for every risk-aversion level including log utility (right panel), in sharp contrast to the sign-flipping static loading $\Lambda(\bar{q})$ of Figure 11. Figure 13 confirms the analytics: the log-utility closed form $e^{s^2+v/2} - 1$ matches the exact equilibrium to within 10^{-9} (left), and the CRRA loading $\Theta(\gamma)$ is positive for every γ with $\Theta(1) = \frac{1}{2}$ (right). The double dissociation survives the change of object: aggregate σ_t continues to price the first moment while consensus uncertainty Ψ_t governs the timing of the second

(Section 7.5).

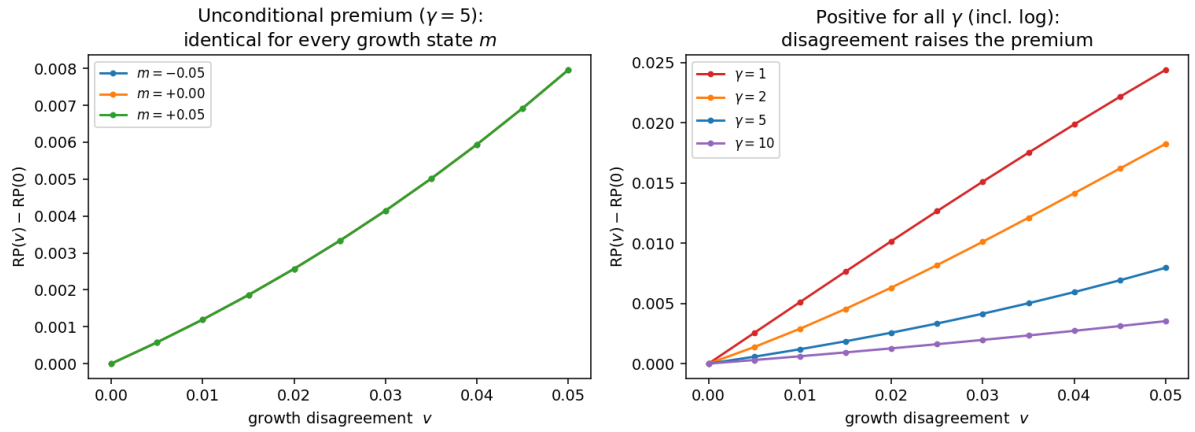


Figure 12: The unconditional disagreement premium (exact equilibrium, $s = 0.15$). Left: $\mathrm{RP}(v) - \mathrm{RP}(0)$ for three growth states m ($\gamma = 5$); the curves coincide, so the premium rises with disagreement *regardless of the state*. Right: the premium is positive for every γ, including log utility. Contrast the binomial loading $\Lambda(\bar{q})$, which is positive only for $\bar{q} < \frac{1}{2}$.

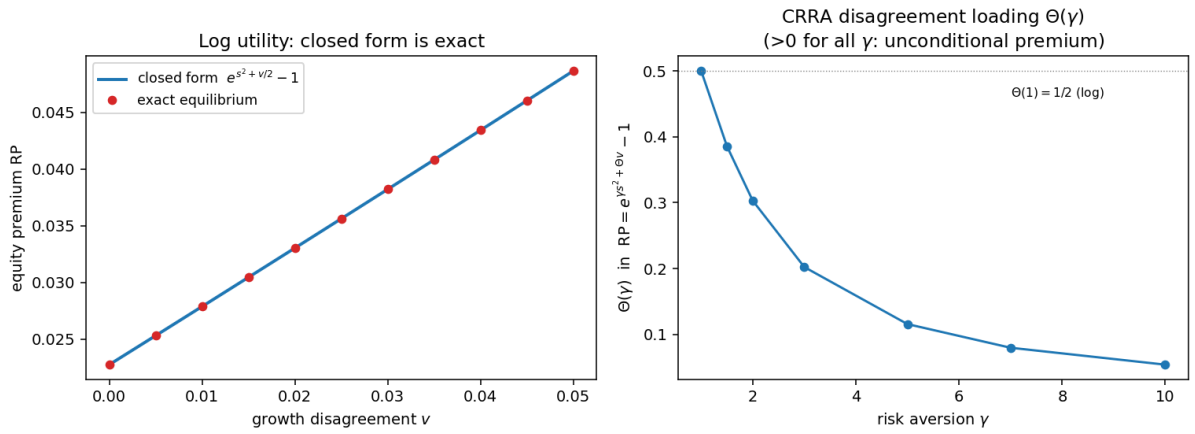


Figure 13: Closed form. Left: the log-utility premium $e^{s^2+v/2} - 1$ (line) versus the exact equilibrium (dots), they coincide to 10^{-9}. Right: the CRRA disagreement loading $\Theta(\gamma)$ in $\mathrm{RP} = e^{\gamma s^2 + \Theta v} - 1$, positive for every γ with $\Theta(1) = \frac{1}{2}$ and decreasing in γ.

7.5 Uncertainty and the timing of the second moment

It remains to place Ψ. In the static economy $\Psi = \bar{q}(1-\bar{q})$ is simply the quantity of payoff risk left to be resolved: the terminal variance of the stock is

$$\mathrm{Var}_{\bar{q}}[R_{\mathrm{stock}}] = \left(\frac{e_H - e_L}{P}\right)^2 \bar{q}(1-\bar{q}) = \left(\frac{e_H - e_L}{P}\right)^2 \Psi, \tag{34}$$

which depends on the consensus uncertainty Ψ and *not* on the dispersion σ. This is the other half of the dissociation: the amount of outcome risk is governed by Ψ, independently

of how much traders disagree. Collecting the static results, the dissociation is a clean set of partial-derivative statements:

$$\frac{\partial \mathrm{RP}}{\partial \sigma^2} = \Lambda(\bar{q}) \neq 0, \qquad \frac{\partial \mathrm{RP}}{\partial \Psi}\Big|_{\text{idiosyncratic}} = 0, \qquad \frac{\partial \mathrm{Var}_{\bar{q}}[R_{\text{stock}}]}{\partial \sigma^2} = 0, \tag{35}$$

i.e. disagreement prices the first moment, consensus uncertainty does not, and the outcome variance does not depend on disagreement. One caveat is worth stating sharply: (34) is a *definitional identity* (Ψ *is* the outcome variance, the quantity of state risk at stake), not a claim that high-Ψ markets display high *realised* stock volatility. A two-date economy collapses "before resolution" and "at resolution" into a single period, so the static model makes *no* prediction about the sign of realised volatility; that result, negative in the level of Ψ, with a burst at resolution, requires the dynamics below.

Two further points complete the mapping. First, why Ψ earns *no* return premium. Our events are *idiosyncratic* macro-political states, individual elections, rate decisions, geopolitical episodes, whose payoff risk is symmetric and, diversified across many independent events, carries no aggregate loading. It therefore commands no risk premium even though it drives each market's volatility: $\partial \mathrm{RP}/\partial \Psi = 0$, which is fact (F3) on the uncertainty side. By contrast aggregate disagreement σ_t is a *systematic* state variable that shifts the economy-wide consensus and hence the price of risk, so it is priced on the common return component.

Second, why Ψ governs the *timing* of volatility with a *negative* sign (Appendix C). A two-date economy cannot speak to this, it collapses "before resolution" and "at resolution" into one period, so consider the natural dynamic extension in which investors learn about the state from a public signal flow and update by Bayes' rule. Each posterior is a martingale under its own filtration, so the consensus belief is (approximately) a bounded martingale and $\Psi_t = \bar{q}_t(1-\bar{q}_t)$ is a *supermartingale* that drifts down to zero as the state resolves. The prediction-market price tracks the belief, so its revision variance is *increasing* in Ψ_t. The stock, however, incorporates the event's cash-flow impact only *as the state resolves*: while Ψ_t is high the event has not yet happened and equity is calm, and the variance is released as a *burst* when $\bar{q}_t$ jumps toward $\{0,1\}$ and Ψ_t falls. Hence near-term realised equity volatility is *decreasing* in the level of Ψ_t and rises as Ψ_t declines, precisely the negative coefficients on Ψ_t and on $\Delta\Psi_{t+1}$ in Table 18. This wait-and-see mechanism is the asset-pricing-with-learning channel of Pástor and Veronesi (2012) specialised to a resolving binary state.

We are deliberate about the status of this argument. The supermartingale decline of Ψ_t follows from Bayesian updating (Appendix C), but the step from that decline to the sign of *realised equity* volatility relies on the auxiliary assumption that the event's cash-flow impact is impounded into equities only as the state resolves. We therefore present the second-moment timing (F2) and the lifecycle hump (F4) as predictions of this learning extension that the data support, not as results of a fully solved dynamic equilibrium. The first-moment results (F1, F3, F5) stand on a different footing: they are closed-form implications of the

solved static and growth economies. A complete dynamic model that endogenises both indices, and with them the lifecycle dynamics of σ_t and Ψ_t, is left to companion work.

7.6 Mapping the model to the evidence

Table 20 collects the correspondence. The architecture is a *double dissociation*: the first moment of returns loads on disagreement, the second-moment timing loads on uncertainty, and neither loads on the other. Disagreement σ is priced because, aggregated across markets, it is a systematic shifter of the consensus and hence of the price of risk; once the object of disagreement is an unbounded growth rate, that price is positive in every state and does not reverse, matching the return evidence and explaining the sign reversal relative to firm-level analyst dispersion. Uncertainty Ψ is *traded* rather than priced: it is the quantity of idiosyncratic, diversifiable state risk, so it earns no premium, but it governs when that risk is incorporated into equity prices, generating the volume and volatility timing patterns we document. The slogan that organises the evidence is therefore *disagreement is priced; uncertainty is traded.*

Empirical fact	Model object	Where
σ predicts returns, positive, permanent (F1)	systematic disagreement shifts the consensus $\hat{q}$ and the price of risk; unbounded growth $\Rightarrow$ unconditional *[derived]*	Eqs. (29)–(33); Table 14
Ψ predicts volume / volatility timing, negative in level (F2)	Ψ = quantity of pending state risk; released as it resolves *[timing conjectured]*	Eqs. (34); Table 18
$\sigma \not\rightarrow$ volatility, $\Psi \not\rightarrow$ returns (F3)	premium loads on σ (systematic), variance loads on Ψ (idiosyncratic, diversifiable) *[derived]*	Eqs. (30)–(35), (34)
Ψ falls, σ humps; within-market < 0, aggregate $\perp$ (F4)	learning lowers Ψ over the lifecycle; differential interpretation generates the σ hump *[conjectured]*	Sec. 7.5; Table 10
risk premium, not mispricing; opposite sign to analyst dispersion (F5)	risk aversion, no short-sale constraint *[derived]*	Sec. 7.2; Atmaz and Basak (2018) vs. Miller (1977)

Table 20: Mapping the theoretical framework to the empirical facts. Each row links one of the five headline empirical facts (F1–F5) to the model object that generates it and to the equations, sections, and result tables where it is derived or tested. Facts F1, F3, and F5 are closed-form implications of the solved static and growth economies (marked *[derived]*), whereas the volatility-timing sign in F2 and the σ life-cycle hump in F4 are *[conjectured]* predictions of the dynamic learning extension that the data support but that we do not derive from a fully solved dynamic equilibrium. The organising slogan is *disagreement is priced; uncertainty is traded.*

Three ingredients are load-bearing and minimal. Risk aversion with $\gamma > 1$ is necessary for any disagreement premium at all (under log utility it vanishes, Eq. (28)). A *cross-section* of idiosyncratic events is necessary for the dissociation, so that consensus uncertainty diversifies and is not priced while still driving each market's volatility. And the object of disagreement must be *unbounded* (a growth rate, not a single bounded event) for the premium to be unconditional rather than sign-dependent. The model is intentionally stylised, it takes the disagreement and uncertainty indices as exogenous state variables rather than deriving their lifecycle dynamics, and it abstracts from the information-leadership channel through which Polymarket and equities are linked at high frequency (Section 5). Both are natural directions for the companion theoretical work; here the framework's role is to show that a single, standard risk-premium mechanism rationalises the full pattern of estimates without appeal to mispricing.

8 Conclusion

How investor beliefs, and the disagreement among them, shape asset prices is a question the asset-pricing literature has asked for decades but struggled to answer, because neither the consensus belief nor its dispersion is directly observable in conventional data. This paper shows that the universe of transactions on Polymarket, a money-backed prediction market with persistent wallet identifiers, recovers both objects directly, in a common probability metric and across a broad menu of macro, political, financial, and corporate events, and uses them to revisit a set of classical asset-pricing questions.

The evidence is threefold. First, the belief object is informationally integrated with the U.S. equity market in both directions: stocks lead Polymarket within the trading day, while overnight Polymarket innovations predict the next morning's equity open. Second, investor disagreement σ positively predicts next-day stock returns, driven almost entirely by the political sub-index, with a sign opposite to that of analyst-forecast dispersion and consistent with compensation for unspanned aggregate state risk rather than mispricing. Third, consensus uncertainty Ψ governs the timing of repricing: high Ψ depresses equity volume and realised volatility, which is released as the uncertainty resolves. Disagreement is priced, while uncertainty is traded.

Two implications follow. For the disagreement literature, the paper supplies a measurement infrastructure — direct, money-backed, high-frequency, wallet-resolved — that addresses the leading interpretive critiques of existing proxies and that uncovers a state-risk channel of disagreement complementary to the firm-level dispersion that the empirical literature has focused on. The same infrastructure can be brought to bear on other empirical questions in which the cross-section of beliefs matters: asset bubbles, momentum, post-earnings-announcement drift, the formation of risk premia in episodes of macro-financial stress. For the prediction-market literature, the paper points to a new use of these markets that goes beyond their forecasting function: the transaction-level cross-section of revealed beliefs is itself a measurement object, and both its dispersion (σ) and the entropy of price-implied state probabilities (Ψ) contain information about external asset prices that the price-level forecast alone does not capture.

Looking forward, as Polymarket continues to grow and as similar money-backed prediction venues come online in other jurisdictions, the measurement infrastructure developed here should become richer, more granular, and applicable to a wider range of asset-pricing questions, including the extensions to safe-asset demand and informedness-based decompositions of σ that we develop in companion work.

References

Antweiler, W. and Frank, M. Z. (2004). Is all that talk just noise? the information content of internet stock message boards. *Journal of Finance*, 59(3):1259–1294.

Atmaz, A. and Basak, S. (2018). Belief dispersion in the stock market. *Journal of Finance*, 73(3):1225–1279.

Avramov, D., Chordia, T., Jostova, G., and Philipov, A. (2009). Credit ratings and the cross-section of stock returns. *Journal of Financial Markets*, 12(3):469–499.

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

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

Bakshi, G., Kapadia, N., and Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143.

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

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

Ben-Rephael, A., Da, Z., and Israelsen, R. D. (2017). It depends on where you search: Institutional investor attention and underreaction to news. *Review of Financial Studies*, 30(9):3009–3047.

Black, F. (1986). Noise. *Journal of Finance*, 41(3):529–543.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Bollen, J., Mao, H., and Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8.

Caldara, D. and Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225.

Carlin, B. I., Longstaff, F. A., and Matoba, K. (2014). Disagreement and asset prices. *Journal of Financial Economics*, 114(2):226–238.

Carr, P. and Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341.

Chen, H., De, P., Hu, Y. J., and Hwang, B.-H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *Review of Financial Studies*, 27(5):1367–1403.

Clark, T. E. and West, K. D. (2006). Using out-of-sample mean squared prediction errors to test the martingale difference hypothesis. *Journal of Econometrics*, 135(1):155–186.

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

Cookson, J. A., Fos, V., and Niessner, M. (2026). Does disagreement facilitate informed trading? *Journal of Financial and Quantitative Analysis*, 61(2):612–639. Published online 4 July 2025.

Cookson, J. A. and Niessner, M. (2020). Why don't we agree? evidence from a social network of investors. *Journal of Finance*, 75(1):173–228.

Da, Z., Engelberg, J., and Gao, P. (2011). In search of attention. *Journal of Finance*, 66(5):1461–1499.

De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–738.

Diercks, A. M., Katz, J. D., and Wright, J. H. (2026). Kalshi and the rise of macro markets. NBER Working Paper 34702, National Bureau of Economic Research.

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

Drake, M. S., Roulstone, D. T., and Thornock, J. R. (2012). Investor information demand: Evidence from google searches around earnings announcements. *Journal of Accounting Research*, 50(4):1001–1040.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *The Review of Economics and Statistics*, 80(4):549–560.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75(4):643–669.

Epstein, L. G. and Schneider, M. (2008). Ambiguity, information quality, and asset pricing. *Journal of Finance*, 63(1):197–228.

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22.

Fox, C. R. and Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics*, 110(3):585–603.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

Gjerstad, S. (2004). Risk Aversion, Beliefs, and Prediction Market Equilibrium. Technical report, University Library of Munich, Germany.

Golez, B. and Goyenko, R. (2022). Disagreement in the equity options market and stock returns. *Review of Financial Studies*, 35(3):1443–1479.

Gómez-Cram, R., Guo, Y., Jensen, T. I., and Kung, H. (2026a). Financial prediction markets: A new measure of earnings expectations.

Gómez-Cram, R., Guo, Y., Jensen, T. I., and Kung, H. (2026b). Prediction market accuracy: Crowd wisdom or informed minority?

Greenwood, R. and Shleifer, A. (2014). Expectations of returns and expected returns. *Review of Financial Studies*, 27(3):714–746.

Grossman, S. J. and Stiglitz, J. E. (1980a). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408.

Grossman, S. J. and Stiglitz, J. E. (1980b). On the impossibility of informationally efficient markets. *The American Economic Review*, 70(3):393–408.

Hansen, L. P. and Singleton, K. J. (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica*, 50(5):1269–1286.

Harris, M. and Raviv, A. (1993a). Differences of opinion make a horse race. *Review of Financial Studies*, 6(3):473–506.

Harris, M. and Raviv, A. (1993b). Differences of opinion make a horse race. *Review of Financial Studies*, 6(3):473–506.

Harrison, J. M. and Kreps, D. M. (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *Quarterly Journal of Economics*, 92(2):323–336.

Hitchcock, L. (2025). Google finance integrates polymarket, kalshi prediction market data. *Yahoo Finance*. Accessed: 2026-04-14.

Hong, H. and Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance*, 54(6):2143–2184.

Hong, H. and Stein, J. C. (2003). Differences of opinion, short-sales constraints, and market crashes. *Review of Financial Studies*, 16(2):487–525.

Illeditsch, P. K. (2011). Ambiguous information, portfolio inertia, and excess volatility. *Journal of Finance*, 66(6):2213–2247.

Johnson, T. C. (2004). Forecast dispersion and the cross section of expected returns. *Journal of Finance*, 59(5):1957–1978.

Kandel, E. and Pearson, N. D. (1995). Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy*, 103(4):831–872.

Karpoff, J. M. (1987). The relation between price changes and trading volume: A survey. *Journal of Financial and Quantitative Analysis*, 22(1):109–126.

Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin, Boston.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.

Lo, A. W. and Wang, J. (2000). Trading volume: Definitions, data analysis, and implications of portfolio theory. *Review of Financial Studies*, 13(2):257–300.

Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica*, 46(6):1429–1445.

Machina, M. J. (2014). Ambiguity aversion with three or more outcomes. *American Economic Review*, 104(12):3814–3840.

Manela, A. and Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of Financial Economics*, 123(1):137–162.

Manski, C. F. (2004). Interpreting the predictions of prediction markets. NBER Working Paper 10359, National Bureau of Economic Research. Published in Economics Letters, 2006, 91(3), 425–429.

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

Murphy, A. H. (1973). A new vector partition of the probability score. *Journal of Applied Meteorology*, 12(4):595–600.

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

Sadka, R. and Scherbina, A. (2007). Analyst disagreement, mispricing, and liquidity*. *Journal of Finance*, 62(5):2367–2403.

Salome, G., Tauchen, G., and Li, J. (2024). Disagreement in market index options. *Journal of Financial Econometrics*, 22(4):1006–1041.

Scheinkman, J. A. and Xiong, W. (2003). Overconfidence and speculative bubbles. *Journal of Political Economy*, 111(6):1183–1219.

Snowberg, E., Wolfers, J., and Zitzewitz, E. (2011). How prediction markets can save event studies. NBER Working Paper 16949, National Bureau of Economic Research.

Snowberg, E., Wolfers, J., and Zitzewitz, E. (2012). Prediction markets for economic forecasting. NBER Working Paper 18222, National Bureau of Economic Research. Published in Handbook of Economic Forecasting, 2013, Vol. 2A, 657–684.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance*, 62(3):1139–1168.

Vlastakis, N. and Markellos, R. N. (2012). Information demand and stock market volatility. *Journal of Banking and Finance*, 36(6):1808–1821.

Whaley, R. E. (2000). The investor fear gauge. *Journal of Portfolio Management*, 26(3):12–17.

Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. NBER Working Paper 10504, National Bureau of Economic Research. Published in Journal of Economic Perspectives, 2004, 18(2), 107–126.

Wolfers, J. and Zitzewitz, E. (2006). Prediction markets in theory and practice. NBER Working Paper 12083, National Bureau of Economic Research. Published in The New Palgrave Dictionary of Economics, 2008.

A Derivations for the static economy

This appendix solves the static economy of Section 7.1 and derives the log-utility neutrality result (28) and the second-order disagreement loading (29)–(30).

Reduction to a single price. Combining the budget constraint $\phi_H c_{iH} + \phi_L c_{iL} = W$ with the first-order condition $c_{iH}/c_{iL} = a_i$ from (24) gives each investor's state consumption as

$$c_{iH} = \frac{W\, a_i}{\phi_H a_i + \phi_L}, \qquad c_{iL} = \frac{W}{\phi_H a_i + \phi_L}, \qquad a_i = \left(\frac{\beta_i}{\Phi}\right)^{1/\gamma}, \quad \Phi = \frac{\phi_H}{\phi_L}. \tag{36}$$

Imposing the two market-clearing conditions $\int c_{iH}\,\mathrm{d}i = e_H$ and $\int c_{iL}\,\mathrm{d}i = e_L$ and taking their ratio, the common wealth W and the price level cancel, leaving a single equation in the single unknown Φ:

$$\frac{\displaystyle\int_0^1 a_i\,(\phi_H a_i + \phi_L)^{-1}\,\mathrm{d}i}{\displaystyle\int_0^1 (\phi_H a_i + \phi_L)^{-1}\,\mathrm{d}i} = \frac{e_H}{e_L} \equiv R. \tag{37}$$

Define $x_i \equiv \Phi a_i = \Phi^{1-1/\gamma}\beta_i^{1/\gamma}$ and the increasing map $\varphi(x) \equiv x/(1+x)$. Dividing numerator and denominator of (37) by ϕ_L and using $c_{iH} \propto \varphi(x_i)$, $c_{iL} \propto 1-\varphi(x_i)$, the clearing condition becomes

$$\int_0^1 \varphi(x_i)\,\mathrm{d}i \;=\; \frac{\Phi R}{1+\Phi R} \;\equiv\; S(\Phi). \tag{38}$$

The equilibrium Φ solves (38), and the consensus $\hat{q}$ is read off from (26), $\Phi = [\hat{q}/(1-\hat{q})]\,R^{-\gamma}$. At the homogeneous point $\pi_i \equiv \bar{q}$ every x_i equals $x_0 = \Phi_0 R = [\bar{q}/(1-\bar{q})]\,R^{1-\gamma}$, and (38) is satisfied with $\hat{q} = \bar{q}$.

Log utility: disagreement washes out. When $\gamma = 1$ the exponent $1/\gamma = 1$, so $x_i = \beta_i = \pi_i/(1-\pi_i)$ is *independent of* Φ and $\varphi(x_i) = \pi_i$. The clearing condition (38) then reads

$$\int_0^1 \pi_i\,\mathrm{d}i = \bar{q} = \frac{\Phi R}{1+\Phi R} \quad\Longrightarrow\quad \Phi = \frac{\bar{q}}{1-\bar{q}}\,R^{-1}, \tag{39}$$

which by (26) gives $\hat{q}/(1-\hat{q}) = \bar{q}/(1-\bar{q})$, i.e. $\hat{q} = \bar{q}$ regardless of the spread. Prices depend only on the mean belief, so $\partial \mathrm{RP}/\partial\sigma^2 = 0$, which is (28).

Risk aversion: the second-order loading. For $\gamma \neq 1$ write $\pi_i = \bar{q} + \varepsilon_i$ with $\int \varepsilon_i\,\mathrm{d}i = 0$ and $\int \varepsilon_i^2\,\mathrm{d}i = \sigma^2$, and let the equilibrium price shift be $\Phi = \Phi_0(1+\omega)$ with $\omega = O(\sigma^2)$. Using $\log\beta(\pi)$ with $\mathrm{d}\log\beta/\mathrm{d}\pi = 1/[\pi(1-\pi)] = 1/\Psi$ and $\mathrm{d}^2\log\beta/\mathrm{d}\pi^2 = (2\bar{q}-1)/\Psi^2$ (evaluated at $\bar{q}$, with $\Psi = \bar{q}(1-\bar{q})$), a second-order expansion of x_i gives

$$\frac{x_i}{x_0} = 1 + \frac{1}{\gamma}\frac{\varepsilon_i}{\Psi} + \left(1-\tfrac{1}{\gamma}\right)\omega + \frac{1}{2}\frac{1}{\gamma}\left[\frac{2\bar{q}-1}{\Psi^2} + \frac{1}{\gamma\Psi^2}\right]\varepsilon_i^2 + o(\sigma^2). \tag{40}$$

Substituting into (38), expanding φ to second order and $S(\Phi)$ to first order in ω, the linear-in-ε terms integrate to zero (because $\int \varepsilon_i \, \mathrm{d}i = 0$), and collecting the surviving $O(\sigma^2)$ terms yields

$$\omega = \frac{\sigma^2}{\Psi^2}\Big[\big(\bar{q} - \tfrac{1}{2}\big) + \tfrac{1}{\gamma}\big(\tfrac{1}{2} - \varphi(x_0)\big)\Big], \qquad \varphi(x_0) = \frac{x_0}{1+x_0}, \quad x_0 = \frac{\bar{q}}{1-\bar{q}}\Big(\frac{e_H}{e_L}\Big)^{1-\gamma}. \tag{41}$$

Since $\hat{q}$ and Φ are linked through (26) by $\mathrm{d}\log\Phi = \mathrm{d}\hat{q}/[\hat{q}(1-\hat{q})]$, the consensus shift is $\hat{q} - \bar{q} = \Psi\,\omega$, i.e.

$$\hat{q} - \bar{q} = \tfrac{1}{2}\kappa_\gamma(\bar{q})\,\sigma^2, \qquad \kappa_\gamma(\bar{q}) = \frac{2}{\Psi}\Big[\big(\bar{q} - \tfrac{1}{2}\big) + \tfrac{1}{\gamma}\big(\tfrac{1}{2} - \varphi(x_0)\big)\Big]. \tag{42}$$

At $\gamma = 1$ we have $\varphi(x_0) = \bar{q}$, so the bracket is $(\bar{q} - \frac{1}{2}) + (\frac{1}{2} - \bar{q}) = 0$ and $\kappa_1 \equiv 0$, consistent with the log case; to leading order near $\gamma = 1$, $\kappa_\gamma \propto (\gamma - 1)$. For $\gamma > 1$ one has $(e_H/e_L)^{1-\gamma} < 1$, hence $\varphi(x_0) < \bar{q}$, and (42) pushes the consensus *away from* $\frac{1}{2}$, toward the more likely state.

Finally, the premium loads on the consensus through the stock price alone: from (27) the objective expected payoff $\bar{q}e_H + (1-\bar{q})e_L$ is fixed at given $\bar{q}$, while a more optimistic consensus bids up the price of the rich-state claim and raises P, so $\partial \mathrm{RP}/\partial\hat{q} < 0$. By the chain rule,

$$\Lambda(\bar{q}) \equiv \frac{\partial \mathrm{RP}}{\partial \sigma^2} = \frac{\partial \mathrm{RP}}{\partial \hat{q}} \cdot \tfrac{1}{2}\kappa_\gamma(\bar{q}), \tag{43}$$

so Λ carries the sign opposite to κ_γ: $\Lambda(\bar{q}) > 0 \Leftrightarrow \bar{q} < \frac{1}{2}$ and < 0 for $\bar{q} > \frac{1}{2}$, which is (30). The magnitude is $O(\sigma^2)$ and is confirmed by solving (38) numerically without approximation.

B The unconditional premium under growth-rate disagreement

This appendix derives the closed-form premium (32) and its CRRA generalisation. Normalise $D_0 = 1$ and let $X = \log D_T \sim N(m, s^2)$ under the objective measure. With complete markets and heterogeneous Gaussian beliefs m_i (cross-sectional mean m, variance v), the economy prices through a representative agent holding the wealth-weighted *consensus* belief; the stochastic discount factor is $M_T = e^{-\rho}(D_T/D_0)^{-\gamma} = e^{-\rho}e^{-\gamma X}$, and asset prices are consensus expectations $\mathbb{E}^c[\cdot]$ of M_T times payoffs.

Log utility ($\gamma = 1$). Here the consensus belief is exactly the *mixture* of the individual Gaussians, $X \sim N(m,\ s^2 + v)$ (Eq. (31)), and $M_T = e^{-\rho}e^{-X}$. The riskless rate follows from $R_f^{-1} = \mathbb{E}^c[M_T] = e^{-\rho}\,\mathbb{E}^c[e^{-X}] = e^{-\rho}e^{-m+(s^2+v)/2}$, i.e.

$$r_f = \rho + m - \tfrac{1}{2}\big(s^2 + v\big), \tag{44}$$

so disagreement v *lowers* the safe rate through the diffuse consensus. The stock is a claim to D_T; its price is $P_0 = \mathbb{E}^c[M_T D_T] = e^{-\rho}\,\mathbb{E}^c[e^{-X}e^{X}] = e^{-\rho}$, so the gross return is $R_{\text{stock}} = D_T/P_0 = e^{\rho}e^{X}$ and its objective expectation is $\mathbb{E}^P[R_{\text{stock}}] = e^{\rho}\,\mathbb{E}^P[e^X] = e^{\rho}e^{m+s^2/2}$ (objective

$X \sim N(m, s^2)$). The equity premium is therefore

$$\mathrm{RP}_{\log}(v) = \frac{\mathbb{E}^P[R_{\text{stock}}]}{R_f} - 1 = \exp\left\{\tfrac{s^2}{2} + \tfrac{s^2+v}{2}\right\} - 1 = e^{s^2+v/2} - 1, \tag{45}$$

which is (32). It is independent of the growth state m (unconditional), reduces to $e^{s^2} - 1$ at $v = 0$, and satisfies $\partial \mathrm{RP}_{\log}/\partial v = \frac{1}{2} e^{s^2+v/2} > 0$.

CRRA ($\gamma > 1$). For a Gaussian consensus $X \sim N(m, \hat{s}^2)$ and $M_T = e^{-\rho} e^{-\gamma X}$, the same three expectations give $R_f^{-1} = e^{-\rho} e^{-\gamma m + \gamma^2 \hat{s}^2/2}$, $P_0 = e^{-\rho} \mathbb{E}^c[e^{(1-\gamma)X}] = e^{-\rho} e^{(1-\gamma)m + (1-\gamma)^2 \hat{s}^2/2}$, and $\mathbb{E}^P[R_{\text{stock}}] = e^{m+s^2/2}/P_0$, so that

$$\mathrm{RP}(v) + 1 = \frac{\mathbb{E}^P[R_{\text{stock}}]}{R_f} = \exp\left\{\tfrac{s^2}{2} + (2\gamma - 1)\tfrac{\hat{s}^2}{2}\right\}. \tag{46}$$

The consensus aggregation of Atmaz and Basak (2018) inflates the consensus variance by the disagreement scaled by risk tolerance, $\hat{s}^2 = s^2 + v/\gamma$ to leading order, which substituted into (46) gives

$$\mathrm{RP}(v) = e^{\gamma s^2 + \Theta(\gamma)\, v} - 1 + O(v^2), \qquad \Theta(\gamma) = 1 - \frac{1}{2\gamma} > 0, \quad \Theta(1) = \tfrac{1}{2}. \tag{47}$$

The loading $\Theta(\gamma)$ is strictly positive for every $\gamma \geq 1$ and *independent of m*, so the premium is unconditional, this is (33). Equation (46) is exact given a Gaussian consensus; the only approximation is the leading-order variance inflation $\hat{s}^2 = s^2 + v/\gamma$, which omits higher-order belief-driven wealth-redistribution terms. The exact loading is pinned by solving the complete-market equilibrium numerically and coincides with (47) at $\gamma = 1$, where the log derivation above is exact.

C Uncertainty and the timing of volatility

Outcome variance. In the static binary economy the stock return takes the value e_H/P in state H (objective probability $\bar{q}$) and e_L/P in state L. Its variance is therefore

$$\mathrm{Var}_{\bar{q}}[R_{\text{stock}}] = \bar{q}(1-\bar{q})\left(\frac{e_H - e_L}{P}\right)^2 = \left(\frac{e_H - e_L}{P}\right)^2 \Psi, \tag{48}$$

which is (34): the outcome variance is governed by the consensus uncertainty $\Psi = \bar{q}(1-\bar{q})$ and does not depend on the dispersion σ. Across a continuum of independent idiosyncratic events each carrying symmetric risk Ψ_j, the law of large numbers removes the common loading, so this risk is diversifiable and earns no premium, $\partial \mathrm{RP}/\partial \Psi = 0$, while still driving each market's own volatility.

Ψ is a supermartingale, and volatility is released at resolution. Add a continuous-time learning stage in which investors observe a public signal and update by Bayes' rule. Each

posterior is a martingale under its own filtration, so the consensus belief is (approximately) a bounded martingale, $\mathrm{d}\bar{q}_t = \varsigma_t \, \mathrm{d}B_t$ with zero drift. Applying Itô's lemma to $\Psi_t = g(\bar{q}_t)$ with $g(q) = q(1-q)$, $g' = 1-2q$, $g'' = -2$,

$$\mathrm{d}\Psi_t = (1-2\bar{q}_t)\,\mathrm{d}\bar{q}_t - \varsigma_t^2\,\mathrm{d}t, \qquad \mathbb{E}_t[\mathrm{d}\Psi_t] = -\varsigma_t^2\,\mathrm{d}t \le 0, \tag{49}$$

so Ψ_t is a supermartingale that drifts down to zero as the state resolves. Two volatilities follow, with opposite signs in Ψ. The prediction-market price tracks the belief, so its instantaneous revision variance ς_t^2 is increasing in Ψ_t (the posterior responds to a signal in proportion to the uncertainty $\bar{q}_t(1-\bar{q}_t)$ still at stake). The stock, by contrast, incorporates the event's cash-flow impact only as the state resolves: while Ψ_t is high the event has not occurred and equity is calm, and the realised variance is released as a burst precisely when $\bar{q}_t$ jumps toward $\{0,1\}$ and Ψ_t falls ($\Delta\Psi < 0$). Near-term realised equity volatility is therefore *decreasing* in the level of Ψ_t and rises as Ψ_t declines, the negative coefficients on Ψ_t and $\Delta\Psi_{t+1}$ in Table 18. This is the learning-driven, wait-and-see channel of Pástor and Veronesi (2012) specialised to a resolving binary state.