Financial Prediction Markets: A New Measure of Earnings Expectations

Roberto Gómez-Cram Yunhan Guo Theis Ingerslev Jensen Howard Kung*

First version: December 10, 2025

This version: April 1, 2026

Abstract

We construct a high-frequency measure of earnings expectations using financial prediction markets. Unlike analyst forecasts, which are updated infrequently and prone to agency conflicts, prediction market prices reflect real-time, stake-backed beliefs. Each contract pays one dollar if realized earnings exceed the analyst consensus; we derive the conditions under which these prices represent subjective probabilities. Relative to analyst forecasts, we find that market-implied expectations are (i) more accurate; (ii) free of standard forecast biases; (iii) incrementally informative for earnings announcement returns; and (iv) lead in price discovery. We trace this informativeness to two forces: the relative skill of earnings-market traders, as evidenced by their profitability in non-earnings markets, and insider trading, for which we find suggestive though not conclusive evidence. We expect this measure to become increasingly informative as the market matures, liquidity deepens, and financial prediction markets become part of mainstream finance.

*Gomez Cram, Guo, and Kung are at London Business School, Jensen is at Yale University; www.tijensen.com. We are grateful for helpful comments from James Choi, Fulin Li, Lasse Pedersen, and Alp Simsek; and seminar participants at Texas A&M and Stockholm School of Economics. Jensen gratefully acknowledges support from the Center for Big Data in Finance (grant no. DNRF167).

1 Introduction

Equity analyst forecasts have become a standard measure of expectations in finance because they are observable, forward-looking, and available for a large cross-section of firms. However, a long-standing concern is that the forecasts may not reflect analysts' true beliefs, because they face incentives to maintain good relations with corporate managers (Michaely and Womack, 1999; Richardson et al., 2004), generate trading commissions (Irvine, 2004), and avoid career risk by staying close to the consensus view (Hong et al., 2000). Moreover, forecasts are updated infrequently, making it difficult to observe how expectations evolve in real time.

We construct a new measure of earnings expectations using prices from newly introduced prediction markets on company earnings on Polymarket. Prediction markets aggregate information continuously through trading, allowing prices to reflect collective beliefs about future outcomes in real time. Participants are financially motivated to acquire and act on relevant information, which encourages the aggregation of dispersed private signals. As a result, prediction markets are widely regarded as efficient mechanisms for information aggregation and have proven accurate predictors of future events across a variety of settings (Wolfers and Zitzewitz, 2004, 2009; Snowberg et al., 2007, 2011).

The first earnings contracts were launched in September 2025, and the expectation is that they, and financial prediction markets more generally, will soon become part of mainstream financial markets. Our current sample covers the period from the market's inception on September 15, 2025, through March 9, 2026 and serves as an initial dataset that we will regularly update as new contracts are resolved and the market matures.[1] Each earnings contract is structured as a binary option contingent on a specific outcome, such as "Will NVIDIA (NVDA) beat quarterly earnings estimates?" The contract pays \$1 if the firm's reported earnings per share (EPS) exceed the analyst consensus forecast—fixed at the time of the market's inception, on average 10 days prior to the announcement—and \$0 otherwise.

In these winner-take-all markets, the price can be interpreted as the market-implied

[1] In a November 2025 earnings call, Robinhood CEO Vlad Tenev stated that the company's prediction-market product had become its fastest-growing business line, with revenues on track to reach an annual run rate of approximately \$300 million. Robinhood facilitates trading on Kalshi, which together with Polymarket accounts for the vast majority of prediction market volume.

probability of an earnings beat, under the assumption that the time value of money and risk premia are negligible (Proposition 1). Further, we show that the current structure of the prediction market we study implies a violation of the law of one price, because these platforms do not pay interest on investors' collateral (Proposition 2). This suggests that market participants as a group deviate from traditional utility-maximizing rational agents.

We provide four new results regarding the informational content and efficiency of financial prediction markets. First, we document that prediction market prices are highly accurate. The Brier score of prediction markets is lower than benchmarks based on analyst forecasts, bias-adjusted analyst forecasts, and the firm-specific historical beat rate. A more intuitive way to see the accuracy is the fact that the market-implied beat probability correctly classifies 75% of earnings announcements seven days prior and 79% on the day before. By contrast, the analyst forecasts issued after the prediction market was opened but before it closed accurately classify 62%.

Second, we find that prediction market expectations are largely consistent with rational benchmarks, in contrast to the well-documented biases in analyst forecasts. While analysts systematically set expectations below subsequent realizations, prediction market prices are approximately unbiased in levels. Tests of forecast revisions, as in Coibion and Gorodnichenko (2015), reveal some evidence of overreaction at very short horizons, a pattern opposite to the underreaction typically found in analyst consensus forecasts (Bouchaud et al., 2019; Bordalo et al., 2019), but the estimates are not significant. Finally, Mincer and Zarnowitz (1969) regressions fail to reject the null hypothesis of forecast rationality, providing some evidence that these markets efficiently aggregate dispersed information despite their nascent status.

Third, we show that prediction-market expectations provide incremental information beyond analyst expectations for explaining earnings-announcement returns. This finding is important because our ability to explain stock returns, even with hindsight, is low in general (Roll, 1988), and for earnings-announcement returns specifically (Ball and Brown, 1968; Beaver, 1968; Ball and Brown, 2014; Hand et al., 2022). When regressing the earnings announcement return on the earnings surprise relative to analyst and prediction-market expectations, respectively, the associated R^2 is 0.038 and 0.046. When adding both surprises

as explanatory variables, they both enter significantly, and the R^2 is 0.063—an increase of 66% relative to the regression that only uses analyst information. Considering that financial prediction markets are still in their infancy, the fact that their informational content surpasses that of analysts is striking.

Fourth, we show that prediction markets play a leading role in price discovery for earnings-related information. New information is incorporated into prediction market prices before it is reflected in analyst forecast revisions, indicating that prediction markets convey information more quickly than professional analysts.

We then examine the drivers of this informational efficiency. Because Polymarket settles trades on a public blockchain, we are able to observe the universe of executed trades and link them directly to individual accounts ("wallets"). Using this data, we document strong positive selection into earnings markets. While these accounts generate zero net profit within earnings events by construction, they extract significant profits elsewhere, averaging $505 per account in non-earnings markets. Given that the vast majority of prediction market traders systematically lose money (Gómez-Cram et al., 2026), this cross-market outperformance suggests that earnings markets attract a disproportionately sophisticated subset of traders.

We also analyze recent public allegations of insider trading in earnings markets, specifically targeting contracts on companies audited by a single auditor. We find that prediction markets on these contracts are indeed more accurate than markets for companies with other auditors. To be clear, our evidence is not conclusive about whether there is insider trading, as other factors could be driving the results.

Finally, we explore how prediction markets can inform research on analyst expectations by creating a market-based measure of analysts' incentives-related biases. We show that past beat history is strongly associated with the market-based bias, yet substantial variation remains unexplained. Across firms, we find that the market-based bias is high for large-quality firms and low for small-junk firms. These results can inform academic research using analyst forecasts, which often faces the difficulty that predictable forecast errors can arise from both true cognitive biases and incentive-related ones. Our results suggest that forecast errors manifesting in small-junk firms are less likely to arise from analyst incentives than those manifesting in large-quality firms.

Given the nascent nature of this prediction market, our current sample is inherently limited. We intend to regularly update our dataset as future earnings seasons conclude. As more data becomes available, the resulting increase in statistical power will allow us to further validate the robustness of our current results.

A key advantage of market-based expectations is that they reveal beliefs backed by financial stakes, providing a direct measure of how investors act on information. This feature addresses an important concern raised by Stefan Nagel in "Perspectives on the Future of Asset Pricing" (Brunnermeier et al., 2021). Nagel argues that asset-pricing models require empirically grounded measures of subjective beliefs and notes that survey-based expectations may suffer from a belief–action disconnect because respondents are not required to act on their stated views. Giglio et al. (2021) demonstrate the relevance of this concern by showing that the link between survey-based expectations and portfolio choices is positive but weak. Market-based expectations inferred from prediction markets avoid this disconnect by reflecting beliefs that are backed by actual trades and therefore directly tied to investors' actions.

We contribute to the literature that uses analyst forecasts to study asset prices and investor behavior (e.g., La Porta, 1996; Engelberg et al., 2018; Berger et al., 2019; Bordalo et al., 2019; Bouchaud et al., 2019; Delao and Myers, 2021; Bordalo et al., 2024; Delao et al., 2024; Bastianello, 2025; Bordalo et al., 2025; de Silva et al., 2025; Jensen, 2025; McCarthy and Hillenbrand, 2025; Thesmar and Verner, 2025). Most of the literature relies on forecasts sourced from I/B/E/S, which are primarily from sell-side analysts employed by banks and have many of the incentive-related issues mentioned above. Another part of this literature seeks to address the issues by relying on crowdsourced earnings forecasts from Estimize (e.g., Jame et al., 2016; Da and Huang, 2020; and Jame et al., 2022). Our main contribution to this literature is to study earnings expectations generated through trading by financially motivated participants.

We also contribute to the recent and growing literature on modern prediction markets. Closest to our paper is work using contracts on macroeconomic variables to measure high-frequency, market-implied macro expectations (Swanson et al., 2025; Diercks et al., 2026; Eichengreen et al., 2025); work linking prediction-market beliefs to asset prices (Li and Xie,

2025; Hartley, 2025); work on forecast accuracy and systematic biases in prediction markets (Bürgi et al., 2025; Reichenbach and Walther, 2025); and work on price discovery across prediction-market venues (Ng et al., 2025). To the best of our knowledge, we are the first to use prediction markets to study firm-level earnings expectations.

2 Financial Prediction Markets on Earnings

2.1 Overview

Prediction markets, often termed information markets or event futures, enable participants to trade contracts with payoffs contingent on the outcomes of uncertain future events. On September 15, 2025, Polymarket[2] introduced a new category of contracts dedicated to forecasting the earnings of publicly traded companies. Our current sample covers the period from the market's inception on September 15, 2025, through March 9, 2026. Given the nascent nature of this prediction market, the available time series is inherently short. However, we intend to continuously update our dataset and analysis as new contracts are listed and resolved, thereby extending the sample period over time.

Each contract represents a binary claim contingent on a specific corporate event, such as "Will NVIDIA (NVDA) beat quarterly earnings?" Participants can buy either "YES" or "NO," and a trade occurs when a YES and a NO side are matched. We will focus our discussion on the YES contract, as the NO contract is just the exact opposite asset. The YES contract pays \$1 if the realized earnings per share are above the analyst consensus *at the contract opening* (for reference, the median contract is opened 11 days before the earnings announcement, so think of the strike price as being the analyst consensus 11 days before the earnings announcement day).[3] Following Polymarket, we interpret the contract

[2]Polymarket is a U.S.-based cryptocurrency prediction market platform and the largest of its kind. It allows users to trade on the outcomes of future events across categories such as economic indicators, weather, awards, and political or legislative developments. Participants deposit USDC through the Polygon blockchain network and trade shares that reflect the market-implied probability of specific outcomes.

[3]For example, consider the contract referencing NVIDIA's earnings release scheduled for November 19, 2025. At the time of market creation, the consensus estimate for NVIDIA's non-GAAP earnings per share (EPS) was \$1.25. Accordingly, the contract resolves to "Yes" if the non-GAAP EPS reported in NVIDIA's official earnings documents strictly exceeds \$1.25, and "No" otherwise. While the benchmark metric—GAAP versus non-GAAP EPS—varies across firms, the applicable definition and resolution source are explicitly

price as the market-implied probability of an earnings beat. Proposition 1 shows that this interpretation is not strictly speaking correct, but likely to hold approximately for the type of short-duration firm-specific earnings contracts we study.

2.2 Data

We obtain prediction market data directly from Polymarket through its public APIs. Polymarket provides two primary APIs: the Gamma API, which delivers market-level information, and the central limit order book (CLOB) API, which provides trade and price data. The Gamma API returns detailed metadata for each market, including the market identifier, question text, description, opening and closing timestamps, while the CLOB API provides real-time order book information and historical prices.

First, we retrieve the full set of earnings markets via the Gamma API. Specifically, we query the endpoint `https://gamma-api.polymarket.com/markets` with the parameters `closed=True` and `tag_id=1013` to identify all concluded earnings-related markets. For each market, the response includes, among other fields, (1) `id`, a unique market identifier; (2) `question`, the question underlying the market; (3) `slug`, a human-readable part of a web address that embeds the stock ticker, earnings announcement date, and the analyst consensus EPS estimate (the market's strike price); (4) `startDate`, the timestamp when the market starts; (5) `closedTime`, the timestamp when the market closes; (6) `volume`, the total trading volume of the market; (7) `clobTokenIds`: the identifiers of the "Yes" and "No" contracts.

Using the market information, we next obtain historical prices from the CLOB API. We query the endpoint `https://clob.polymarket.com/prices-history` with parameters `market`, `startTs` and `endTs`. For each earnings market, the `market` field is set to the first element of `clobTokenIds`, i.e. the "Yes" contract id. The time window parameters `startTs` and `endTs` correspond to the market's creation time and closing time, respectively, expressed as Unix timestamps[4]. Each request returns a minute-level time series of historical prices for that specific earnings market.

specified for each contract ex ante.

[4]For markets with long durations, we accommodate API limits by segmenting the overall time horizon into subperiods and downloading the corresponding data in batches.

To compare the prediction market with sell-side analysts, we collect historical analyst forecasts from I/B/E/S. Using the I/B/E/S Unadjusted Detail History File, we extract quarterly EPS forecasts issued within a 90-day window prior to each firm's earnings announcement. Forecasts are adjusted with CRSP cumulative adjustment factors. Consensus forecasts are computed as the mean, consistent with the strike price definition on Polymarket. As of March 2026, I/B/E/S data are available through February 19, 2026. Consequently, our historical I/B/E/S data (spanning 1983 to February 19, 2026) do not fully cover the Polymarket sample period (September 15, 2025 to March 9, 2026).

To cover the remaining days, we supplement our data with analyst forecasts from FactSet. For each earnings announcement corresponding to a Polymarket market, we manually collect the consensus forecast and realized EPS from FactSet Connect. FactSet computes consensus forecasts as the mean of all outstanding analyst estimates up to the earnings announcement date, whereas Polymarket strike prices are fixed at market creation. Differences between FactSet consensus forecasts and Polymarket strike prices, therefore, capture revisions in analysts' expectations that occur between market creation and the earnings announcement.

Table 1 reports summary statistics for our sample. Column 1 indicates that the median contract opens 11 days prior to the earnings announcement. As shown in Column 2, the average trading volume per contract is approximately $27,197. Column 3 reports an average pre-announcement contract price of 0.72, implying a market-implied probability of an earnings beat of 72%. This is slightly lower than the realized beat frequency of 73% reported in Column 4. The interquartile range for the prediction market price is 0.62 to 0.87 (see also Figure E.4 in the Appendix for the full distribution). Columns 5 and 6 reveal that prediction markets cover significantly larger firms than the broader universe of companies followed by analysts: the average market capitalization in our sample is approximately $120 billion, compared to $16 billion for the full I/B/E/S dataset. Similarly, Columns 7 and 8 indicate that firms in the prediction market sample exhibit lower book-to-market ratios.

The share of firms that beat analyst forecasts, as well as the average implied probability from prediction markets, is well above the 50% benchmark that would prevail if analyst expectations were unbiased. This pattern is well known and is often attributed to analysts' incentives to please management by "walking down" their forecasts, thereby making it easier

Table 1: Summary Statistics

	Prediction Market	Trading	Prediction Market	Actual	Stock Market Capitalization (Bn $)		B/M Ratio	
	Duration (days)	Volume ($)	Price	Beat	Prediction	Analysts	Prediction	Analysts
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean	10.59	27,197	0.72	0.73	117.06	16.28	0.41	0.68
Std	4.86	69,776	0.21	0.44	387.19	118.28	0.38	0.77
Min	0	1,336	0.07	–	0.03	0.00	0.01	0.00
25%	7	7,791	0.62	–	5.08	0.28	0.12	0.24
Median	11	13,560	0.79	–	19.18	1.30	0.28	0.51
75%	14	24,230	0.87	–	74.62	5.84	0.57	0.88
Max	36	995,130	0.97	–	3,785	3,785	2.46	16.51
Obs	639	625	637	639	623	3,803	560	3,395

Notes: This table presents summary statistics for the sample of earnings prediction markets. Column 1 represents the number of days between market creation and the earnings announcement date. Column 2 represents the prediction market trading volume. Column 3 represents the prediction market price one day prior to the earnings announcement. Column 4 represents a dummy variable that equals 1 if the actual outcome is beat and 0 otherwise. Columns 5 and 6 (7 and 8) represent the stock market capitalization (book-to-market ratio), for firms covered by prediction markets and analysts from I/B/E/S, respectively. The prediction market sample spans 15 September 2025 to 9 March 2026.

for firms to report earnings beats (Richardson et al., 2004).[5]

Table E.1 in the Appendix details the industry composition and lists the top three firms by trading activity. Figure E.2 in the Appendix shows that even though the contracts trade 24/7, the market is still relatively illiquid, in the sense that the average contract has one or more price updates in 6% of all minutes it is traded, 41% of all hours, and 92% of all days. The figure also shows that activity is higher during regular U.S. trading hours, but there is still significant activity outside these hours, and overall, the difference in activity is much smaller than for stocks. Finally, Figure E.3 in the Appendix shows that the trading activity is highest around market close and on Wednesdays, whereas there is less activity during weekends.

[5]Figure E.1 plots the yearly proportion of firms beating the I/B/E/S consensus forecast from 1983 to 2025, and it reveals a distinct upward trend. In 1983, approximately 45% of firms exceeded the benchmark; this share increased steadily to roughly 75% by the end of the sample.

3 Interpreting Prediction Market Prices

Prediction market platforms like Polymarket and Kalshi often refer to prices as reflecting physical probabilities. In this section, we show that this interpretation is not, in general, correct. Further, we show that the current structure of prediction markets for earnings implies a violation of the law of one price.

Let t be the time at which a contract is opened, and $t+K$ the time at which it is resolved (i.e., the company's earnings announcement). Let $R_{ft}^{(K)}$ be the gross risk-free return from t to $t+K$. Further, let X_{t+K} denote realized earnings per share and $\overline{F}_t$ the analyst consensus forecast. Define the event

$$A_{t+K} \equiv \{X_{t+K} > \overline{F}_t\},$$

and let $\mathbb{1}_{A_{t+K}}$ be \$1 if the event materializes and \$0 otherwise. The "YES" contract pays $\mathbb{1}_{A_{t+K}}$ at time $t+K$, while the "NO" contract pays $1 - \mathbb{1}_{A_{t+K}}$. We first characterize the pricing of the YES contract (all proofs are in Appendix A).

Proposition 1 (Prediction market prices are not generally physical probabilities)**.** *Suppose there exists a stochastic discount factor $M_{t,t+K}$. Then the price of the YES contract is*

$$p_t = \mathrm{E}_t\left[M_{t,t+K}\mathbb{1}_{A_{t+K}}\right] \tag{1}$$

$$= \frac{\mathrm{E}_t\left[\mathbb{1}_{A_{t+K}}\right]}{R_{ft}^{(K)}} + \mathrm{Cov}_t\left(M_{t,t+K}, \mathbb{1}_{A_{t+K}}\right). \tag{2}$$

Hence, the prediction market price equals the physical probability of the event only under special conditions. In particular, the prediction market price equals the physical probability if $R_{ft}^{(K)} = 1$ and $\mathrm{Cov}_t\left(M_{t,t+K}, \mathbb{1}_{A_{t+K}}\right) = 0$.

Polymarket interprets the price of the YES contract as the physical probability that the company will beat the consensus forecast.[6] For example, they interpret a price of \$0.65 as indicating a 65% probability that the firm beats the analyst consensus. Proposition 1 shows that this interpretation is not generally correct, as it ignores the effects of time discounting

[6] See for example, the following Polymarket articles: Polymarket 101 and How Prices are Calculated.

and risk premia.

That said, we argue that the interpretation is approximately correct for earnings contracts. In particular, the typical duration of the contracts is short, implying that the risk-free rate earned is likely small, and the contract outcomes represent diversifiable firm-specific risk, implying that the covariance adjustment is also likely small. Appendix B provides additional detail and discussion of the assumptions under which the market price can be interpreted as revealing investors' expectations. For the remainder of the paper, we will interpret prediction market prices as the prediction market-implied probability of an earnings beat.[7] Next, we consider the joint pricing of the YES and NO contracts.

Proposition 2 (Prediction market prices violate the law of one price)**.** *The payoff from the portfolio of one YES and one NO contract guarantees the holder a risk-free payoff of \$1. By construction on Polymarket, this portfolio is always priced at \$1. The portfolio can be replicated by buying a risk-free zero-coupon bond paying \$1 at the time of the contract resolution, which has a price of* $\frac{1}{R_{ft}^{(K)}}$*. The prediction market price thus violates the law of one price whenever* $R_{ft}^{(K)} > 1$*.*

A portfolio of one YES contract and one NO contract guarantees a risk-free payoff of \$1. The price of this portfolio is always \$1, because a trade only happens when a person wanting to buy a YES contract at p_t finds a person willing to buy a NO contract at $1-p_t$. Proposition 2 shows that this violates the law of one price, because the price of the corresponding risk-free bond is below one in the empirically relevant case of positive interest rates. The proposition shows that, collectively, prediction market traders would be better off investing in a risk-free bond.[8]

Proposition 2 suggests that there would be no trading on Polymarket if everyone had standard preferences and rational beliefs. So why do we see so much trading on prediction

[7] While we will focus on the prediction market implied beat probability, we also propose a methodology for converting these into earnings level expectations in Appendix D.

[8] Polymarket may be aware of the law of one price violation, as it seems like they are taking actions to remove it. Specifically, Polymarket recently added a 4% annualized holding reward on 13 eligible contracts, related to long-term political contracts (for a list of the current eligible contracts, see https://help.polymarket.com/en/articles/13364459-holding-rewards). At the time of writing, no earnings contracts were part of the eligible contracts, meaning that our entire sample is subject to the law of one price violation from Proposition 2.

markets? One reason could be differences in beliefs. Both the YES and NO investors may think they are making a good trade based on their subjective beliefs, and believe that the law of one price violation arises because the other side of the trade is making a particularly bad trade. Alternatively, investors may get non-pecuniary rewards from trading in prediction markets. Finally, many trades are small, so the uncollected interest may be trivial and therefore ignored. In any case, Proposition 2 shows an inherent inefficiency in the construction of modern prediction markets, which suggests that prediction market investors as a group deviate from standard utility-maximizing rational agents.

4 Informational Content

In this section, we evaluate the informational content of prediction market-implied expectations by examining their accuracy, their relation to stock returns, potential bias, and their ability to predict analyst revisions.

4.1 Accuracy

Table 2 reports the predictive accuracy of the prediction markets by comparing the market price to the subsequent realization. To align with the close of the U.S. stock market, we sample prediction market prices at 4:00 PM Eastern Time.

Panel A computes the Brier score of prediction markets as

$$\text{Brier Score} = \frac{1}{N}\sum_{i=1}^{N}\left(p_{i,t} - \mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}}\right)^2. \tag{3}$$

The overall Brier score for prediction markets is 0.152. To assess whether this is a "good" score, we compare it to several benchmarks. The first benchmark predicts 0.5 for all contracts, and we call it the "analyst" benchmark, as 0.5 is the price that would prevail if prediction markets agreed with analysts. The Brier score is 0.25, suggesting that prediction markets are more accurate than analysts. The second benchmark predicts 0.73 (the realized beat rate within our sample) for all contracts, and we call it the "analyst bias-adjusted" benchmark, as it is the price that would prevail if prediction market investors simply ad-

Table 2: Prediction Accuracy

Panel A. Prediction markets versus benchmarks

	Prediction	Benchmark		
	Market Price 1 Day Before	Analyst $p = 0.5$	Analyst bias-adjusted $p = 0.73$	Firm Beat Rate Past 20 Quarters
Brier Score	0.152	0.250	0.193	0.185
Observations	613	613	613	613

Panel B. Prediction markets versus analyst revisions

	Prediction Market						Analyst	
	Full Sample				FactSet Sample		FactSet Sample	
	7 Days Before		1 Day Before		1 Day Before			
Actual	Miss	Beat	Miss	Beat	Miss	Beat	Miss	Beat
Miss	7%	20%	11%	16%	13%	17%	21%	10%
Beat	6%	67%	6%	68%	5%	64%	28%	42%
Accuracy	75%		79%		78%		62%	
Observations	448		637		365		365	

Notes: This table evaluates the predictive accuracy of prediction markets relative to analyst-based benchmarks. Panel A reports Brier scores for prediction market prices and three analyst-based price measures: (i) a naive benchmark with $p = 0.5$, (ii) prices implied by the aggregate beat rate ($p = 0.73$), and (iii) prices implied by firm-specific historical beat rates computed over the preceding 20 quarters. Panel B reports a confusion matrix comparing the prediction markets and analyst revisions against actual outcomes. The prediction market predictions are shown for 7 and 1 day prior to the earnings announcement. The predicted outcome is "Beat" if the price is over 0.5, and "Miss" otherwise. Analyst predictions are based on forecast revisions between the earnings market creation and the actual earnings announcement. The predicted outcome is "Beat" ("Miss") if the revision is positive (negative). Accuracy is the sum of the diagonal, that is, the fraction of outcomes that were correctly classified as beat or miss. Results are reported for both the full sample and the restricted FactSet sample. The prediction market sample spans 15 September 2025 to 9 March 2026.

justed for the average analyst bias. The Brier score is 0.193, suggesting that prediction markets do more than a simple level-bias adjustment of analyst forecasts. The third and final benchmark predicts the firm-specific realized beat rate over the past 20 quarters. The Brier score is 0.185, the best among the benchmarks, but still far above the Brier score of prediction markets.

Panel B evaluates accuracy using a confusion matrix, by comparing the binary outcome implied by the market price (beat if above 0.5; miss if below) to the subsequent realiza-

tion. Accuracy is the share of correct predictions. Seven days before the announcement, the overall accuracy is approximately 75%, and this number increases to 79% on the day before the earnings announcement. This suggests that prediction market prices become more informative closer to the resolution date.

We also compare the prediction market price to analyst data sourced from FactSet. We code the FactSet-implied prediction as a beat if the FactSet consensus was revised upwards between the market creation and the market resolution, and as a miss if it was revised downwards. The resulting accuracy for FactSet is 62%, whereas the prediction market accuracy on a comparable sample is 78%.

4.2 Biases

Forecasts produced by professional forecasters, such as analysts, often exhibit systematic biases, meaning they deviate from a rational benchmark. In this subsection, we examine whether similar biases are present in the prediction market-implied expectations.

4.2.1 Level Bias

A straightforward implication of full-information rational expectations models is that the average market-implied probability of an earnings beat should coincide with the average realized outcome. As shown in Figure E.1, this implication is violated in analyst data, where analyst expectations immediately before the earnings announcement are, on average, too low. In other words, the probability of a company beating analyst expectations is above 50%. The standard explanation for this pattern in the accounting and finance literature is that the low expectations reflect analysts' strategic incentives rather than their true beliefs (Richardson et al., 2004; Berger et al., 2019).[9] An alternative explanation is that analysts are reporting their true beliefs, in which case the low expectations could reflect a form of "defensive pessimism" (Norem, 2008; Norem and Cantor, 1986). Strategic incentives do not distort market-based expectations, so they provide a clean laboratory for distinguishing

[9]Specifically, the idea is that analysts have a strategic incentive to please company management (e.g., to retain access to meetings with the corporate management) so they report expectations lower than what they truly believe to make it easier for companies to report a "beat."

between these two explanations.

Let $\mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}}$ be a dummy variable equal to one if firm i on earnings date $t+K$ beats the analysts' consensus forecast and zero otherwise, where K is the number of days prior to the earnings announcement at time t. Let $p_{i,t}$ denote the contract price for company i on day t. Define the prediction error as $e_{t|t+K} = \mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}} - p_{i,t}$. The bias in the prediction can then be estimated using the following regression:

$$e_{t|t+K} = a + u_t, \tag{4}$$

where $a = 0$ indicates that forecasts are unbiased at horizon K.

Panel A of Table 3 reports our estimates of a. The coefficients are not statistically significant at conventional levels, which supports the null hypothesis of no bias in prediction markets. However, the point estimates are consistently positive across all horizons, suggesting that market-implied expectations may be, on average, too low; a pattern similar to that observed in analyst expectations. As more data becomes available, we may gain the statistical power necessary to distinguish between the presence of systematic bias and the null hypothesis of no bias.

Another possible interpretation of the positive parameter estimate, is that it captures our approximation error. In particular, we assume that the prediction market price represents the physical market-implied probability of an earnings beat, but in reality prices also incorporate adjustments for the time value of money and risk.

The results suggest that the time adjustment is unlikely to be of first-order importance, since the parameter estimates are similar across all three horizons. The risk adjustment hypothesis is more difficult to test. The positive sign is consistent with the contracts being positively correlated with the stochastic discount factor, implying that prices are below the cash-flow expectation. That said, these events represent seemingly diversifiable firm risk, so a risk adjustment of 0.02 seems implausibly large. To see this point, note that the average contract price one day before earnings announcement is approximately 0.72. If this price reflects investors' cash-flow expectation of 0.74 plus a risk adjustment of 0.02, then it would imply a discount rate of $0.74/0.72 - 1 = 2.8\%$, which corresponds to an implausibly high

rate when annualized.[10]

4.2.2 Mincer–Zarnowitz Regressions

Another standard test of forecast rationality is the Mincer and Zarnowitz (1969) regression, which evaluates whether realized outcomes move one-for-one with forecasts. Applied to our setting, the test examines whether firms that have higher market-implied probabilities of beating earnings are indeed more likely to do so. Formally, we estimate the regression:

$$\mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}} = \alpha + \beta p_{i,t+K-1} + u_{i,t+K}, \tag{5}$$

Under full-information rational expectations, forecast rationality implies that $\alpha = 0$ and $\beta = 1$, meaning there should be no bias and realized outcomes should move one-for-one with beating probabilities.

Panel B of Table 3 reports our estimates of α and β from equation (5). We find that the intercept α is indistinguishable from zero and the slope coefficient β is close to unity. We cannot statistically reject the hypothesis that $\beta = 1$. These results are consistent with the null hypothesis of forecast rationality. We note, however, that the power of these tests is currently limited by our sample size. As the prediction market matures and our dataset expands, we will obtain sharper estimates.

To get further insights into the calibration of prediction markets, Figure 1 shows a binned scatter plot of the prediction market price against the actual outcome. Consistent with a positive level bias, most of the dots lie above the 45-degree line. However, a notable outlier is the group of prediction market prices between 0 and 0.2. These low prices happen rarely—only 13 cases in our sample—but the striking fact is that all 13 resolved as missed. While the sample is too small to draw any reliable conclusions, the apparent overestimation of low-probability events is consistent with the long-shot bias that has been documented extensively in horse betting (e.g., Snowberg and Wolfers, 2010), but which does not seem to be a feature of the typical prediction market contract (Reichenbach and Walther, 2025).

[10] Relatedly, our regression results imply that a simple strategy that buys all contracts before the event would have achieved an extraordinarily high return in our sample. That said, given the limited liquidity in the market, this high relative return does not necessarily translate into a large dollar profit. Still, we expect

Table 3: Biases in Financial Prediction Markets

	(1)	(2)	(3)
Panel A. Level Bias			
	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}} - p_{i,t+K-1}$	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}} - p_{i,t+K-3}$	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}} - p_{i,t+K-5}$
Intercept	0.02	0.01	0.02
t-statistic	[1.10]	[0.63]	[0.94]
Observations	637	626	582
Panel B. Mincer-Zarnowitz test			
	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}}$	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}}$	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}}$
Intercept	-0.00	-0.02	0.02
t-statistic	[0.16]	[0.28]	[0.00]
$p_{i,t+K-1}$	1.03	–	–
t-statistic	[0.16]		
$p_{i,t+K-3}$	–	1.05	–
t-statistic		[0.28]	
$p_{i,t+K-5}$	–	–	1.00
t-statistic			[0.00]
Observations	637	626	582
Panel C. Coibion-Gorodnichenko Test			
	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}} - p_{i,t+K-1}$	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}} - p_{i,t+K-1}$	$\mathbb{1}_{\{X^i_{t+K}>\overline{F}^i_t\}} - p_{i,t+K-3}$
Intercept	0.02	0.02	0.01
t-statistic	[1.07]	[1.25]	[0.66]
$\Delta p_{i,[t+K-3,t+K-1]}$	-0.19	–	–
t-statistic	[-1.37]		
$\Delta p_{i,[t+K-5,t+K-1]}$	–	-0.14	–
t-statistic		[-1.18]	
$\Delta p_{i,[t+K-5,t+K-3]}$	–	–	-0.21
t-statistic			[-1.24]
Observations	626	582	582

Notes: This table tests for biases in prediction market prices. Panel A tests whether prices are unbiased predictors of actual outcomes. Panel B reports the Mincer–Zarnowitz regression, where the reported t-statistics for $p_{i,t+K-h}$ are from the test of whether the coefficient equals one. Panel C implements the Coibion–Gorodnichenko test. The prediction market sample spans 15 September 2025 to 9 March 2026.

the level bias to become smaller as the market becomes more efficient.

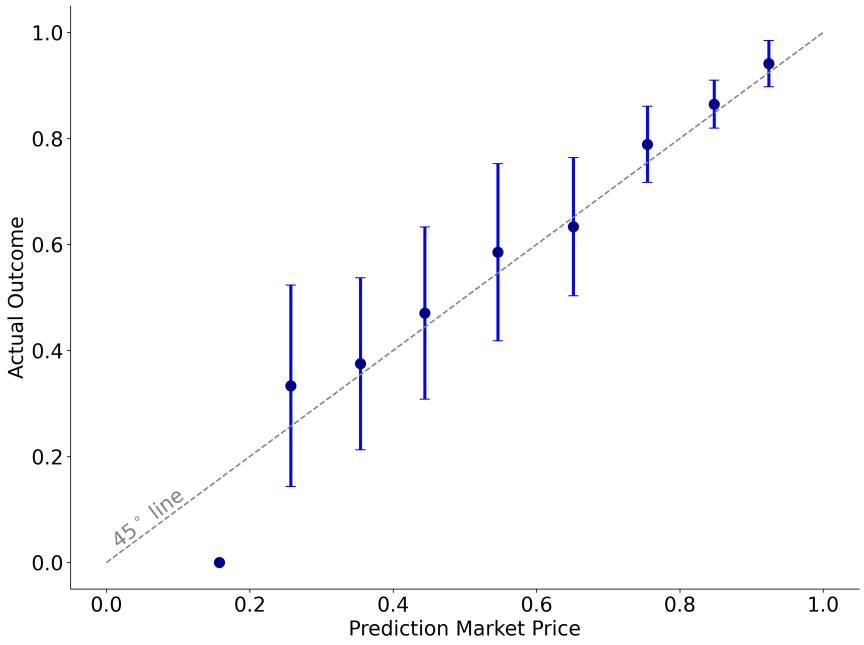


Figure 1: Actual Outcomes and Prediction Market Prices

Notes: This figure presents a binned scatter plot of prediction market prices against actual outcomes. The sample is partitioned into 9 bins of equal width (0.1) based on the prediction market price, with the two lowest bins (0–0.1 and 0.1–0.2) merged due to sparse observations. Each point plots the bin-average prediction market price against the average realized outcome. Error bars denote 95% confidence intervals based on standard errors clustered by firm and earnings announcement date. The dashed line is the 45° line representing perfect calibration. The prediction market sample spans 15 September 2025 to 9 March 2026.

4.2.3 Coibion-Gorodnichenko Test

Another standard test of forecast rationality is the test from Coibion and Gorodnichenko (2015), which assesses whether forecast revisions predict subsequent forecast errors. Intuitively, if market-based expectations are rational, revisions to an earnings announcement should not predict the subsequent error. We analyze this relationship using the regression

$$\mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}} - p_{i,t+K-1} = \alpha + \beta \Delta p_{i,[t+K-3,t+K-1]} + u_{i,t+K}, \tag{6}$$

where $\mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}}$ equals one if firm i beats the analysts' EPS consensus on date $t+K$, $p_{i,t+K-1}$ is the market-implied beat probability on the day before the announcement, and $\Delta p_{i,[t+K-3,t+K-1]}$ denotes the cumulative change in the market-implied probability over the three days preceding the announcement. Under full-information rational expectations, $\alpha = 0$ and $\beta = 0$.

Panel C of Table 3 reports the estimated version of equation (6). We find some evidence of overreaction: the coefficient on the three-day price revision, $\Delta p_{i,[t+K-3,t+K-1]}$, is negative but not statistically significant. The same is true for the other horizons we consider. The evidence of overreaction contrasts with the typical finding in the analyst expectations literature, where Coibion-Gorodnichenko regressions generally indicate underreaction in short-run forecasts (Bouchaud et al., 2019; Bordalo et al., 2019). But the lack of statistical significance, suggest that this finding could be due to random variation.

4.3 Stock Returns

Roll (1988) shows that even with hindsight, our ability to explain stock returns with observable information is limited. The limited ability is particularly noticeable for returns around earnings announcements, where a long-standing puzzle is that earnings surprise (that is, realized earnings minus analysts' expectation), revenue surprises, and other observable news still leave the vast majority of earnings announcement return variation unexplained (Ball and Brown, 1968; Beaver, 1968; Ball and Brown, 2014; Hand et al., 2022).

We replicate the earnings announcement puzzle in our sample, and then test whether financial prediction markets can help improve our understanding of stock returns by estimating three regressions:

$$r_{i,t} = a_1 + b_1 \text{SUE}^{\text{analyst}}_{i,t} + \epsilon_{i,t}, \tag{7}$$

$$r_{i,t} = a_2 + b_2 \text{SUR}^{\text{pred}}_{i,t} + \epsilon_{i,t}, \tag{8}$$

$$r_{i,t} = a_3 + b_3 \text{SUE}^{\text{analyst}}_{i,t} + b_4 \text{SUR}^{\text{pred}}_{i,t} + \epsilon_{i,t}, \tag{9}$$

where $r_{i,t}$ is the return of the stock on the day of its earnings announcement, $\text{SUE}^{\text{analyst}}$ is the analyst earnings surprise, computed as the difference between actual and analyst consensus EPS, scaled by the stock price one day before the earnings announcement, and SUR^{pred} is the prediction market surprise, computed as the difference between the realized outcome and the prediction market price.[11]

[11] Earnings announcements frequently happen outside of regular market hours. The way we compute earnings announcement returns, therefore, depends on when the firm announces. For firms that announce

Table 4: Earnings Announcement Returns and Earnings Surprises

	Earnings Announcement Day Return					
	Raw			Market-Adjusted		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	-0.05	-0.10	-0.12	0.09	0.04	0.02
t-statistic	[-0.13]	[-0.24]	[-0.30]	[0.22]	[0.11]	[0.06]
$\text{SUE}^{\text{analyst}}$	1.60	–	1.14	1.57	–	1.14
t-statistic	[4.97]	–	[3.39]	[4.94]	–	[3.41]
SUR^{pred}	–	0.06	0.05	–	0.06	0.04
t-statistic	–	[5.51]	[4.13]	–	[5.35]	[3.97]
Observations	635	635	635	635	635	635
R^2	0.038	0.046	0.063	0.037	0.043	0.061

Notes: This table reports regressions of earnings announcement returns on sell-side analyst surprises and prediction market surprises. $\text{SUE}^{\text{analyst}}$ is the difference between the actual and consensus EPS, scaled by the stock price one day before the earnings announcement. SUR^{pred} is the difference between the realized outcome dummy and the prediction market price one day before the earnings announcement. Columns (1)–(3) use raw announcement returns; columns (4)–(6) use market-adjusted returns. *t*-statistics are reported in brackets. The prediction market sample spans 15 September 2025 to 9 March 2026.

Table 4 shows the results. The first column shows that the parameter estimate on the analyst earnings surprise is highly significant and positive, but the R^2 is low at 0.038, confirming the puzzle that our ability to explain earnings announcement returns is limited. The second column shows that the prediction market surprise is also highly significant and positive, and the R^2 is higher than for analyst expectations. The third column shows that when both analyst and prediction market earnings surprises are included in the same regression, they are both significant, and, importantly, the R^2 is higher than for the univariate regressions by 66% (relative to analysts only) and by 37% (relative to prediction markets only). These results show that prediction markets provide additional information beyond analyst expectations, which can help us understand why stock prices move.

To make further progress on understanding stock returns from prediction market prices, Table 5 regresses stock returns on contemporaneous changes in prediction market probabili-

earnings before market open or during regular market hours, the returns are based on prices at market close from $t-1$ to t. For firms that announce earnings after market close, the return is based on close prices from t to $t+1$.

Table 5: Stock returns and prediction market prices

	Stock Return $r_{i,[t_1,t_2]}$ Across Earnings Windows $[t_1, t_2]$				
	[-4, 1]	[-3, 1]	[-2, 1]	[-1, 1]	[0, 1]
	(1)	(2)	(3)	(4)	(5)
Intercept	-0.23	-0.69	-0.33	-0.10	-0.06
t-statistic	[-0.40]	[-1.27]	[-0.55]	[-0.18]	[-0.12]
$\Delta p_{i,[t_1,t_2]}$	0.06	0.07	0.06	0.05	0.06
t-statistic	[5.22]	[6.61]	[6.66]	[5.44]	[6.19]
Observations	529	591	624	635	637
R^2	0.050	0.079	0.057	0.041	0.043

Notes: This table reports regressions of contemporaneous stock returns and prediction market price changes across several earnings announcement windows. Event windows are defined by trading days relative to the earnings announcement date (EAD); for example, [-4, 1] denotes the period from four trading days before EAD to one trading day after. The prediction market sample spans 15 September 2025 through 9 March 2026.

ties across various periods. In Column 1, we look at the period from four days prior to one day after the earnings announcement. The parameter estimate indicates that a one–percentage-point increase in the prediction market beat probability is associated with a 6 basis point increase in stock returns, with a t-statistic of 5.22. Columns 2 to 5 consider narrower windows that all end one day after the earnings announcement date. The parameter estimates remain positive and do not change much in magnitude despite the narrower window, which suggests that the effect is stronger closer to the earnings announcement date.

4.4 Price Discovery in Prediction Markets

A key feature of prediction markets is that prices can incorporate new information in real time, whereas professional analysts may revise their forecasts with a delay. Thus, a potential benefit of prediction markets is their ability to incorporate and convey information about earnings more quickly than analysts.

We motivate this idea by a case study of the prediction market contract on Tesla's 2025Q3 earnings. Figure 2 plots the prediction market price change and individual analyst forecasts in the pre-announcement period, corresponding to the quarter ending in September 2025.

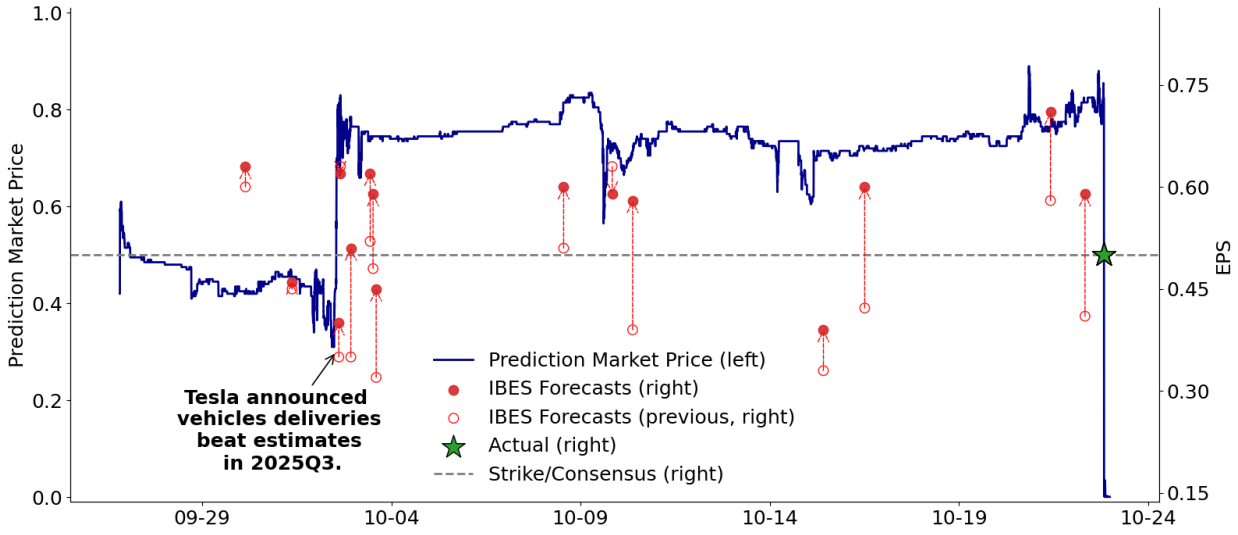


Figure 2: Case study: Tesla's 2025Q3 earnings

Notes: This figure compares the prediction market price and analyst forecasts for Tesla's 2025Q3 earnings. The dark blue line represents the prediction market price. The solid red circle represents the most recent individual analyst EPS forecasts, and the open red circle for the previous forecast issued by the same analyst. The green star represents the actual announced EPS. The dashed line represents the strike price used by the prediction market, i.e. the analyst consensus forecast to beat when the earnings market was created.

The prediction market was created on September 26, 2025, with the earnings announcement date scheduled for October 22, 2025. At approximately 9:00 a.m. on October 2, Tesla announced that its third-quarter vehicle deliveries exceeded analyst expectations. Immediately following this news, the prediction market price increased sharply, rising from 0.44 at 9:02 a.m. to 0.68 by 9:14 a.m., and further to 0.80 by 9:53 a.m. Market participants also posted comments discussing this announcement, as shown in Appendix C.2. Analyst forecast revisions followed, with analysts generally raising their earnings expectations after the price adjustment in the prediction market. Despite the positive news, Tesla failed to beat its earnings, and the contract closed as a "miss". Still, in this particular case, prediction market prices contained more timely information than analysts.[12]

To test this mechanism more generally, we collect all individual analyst revisions from I/B/E/S between each contract's creation and resolution date, and estimate the following

[12]Appendix C.3 contains an additional case study of J.P. Morgan's 2025Q4 earnings, showing how the timeliness of analyst expectations is sometimes informed by users openly spreading value-relevant information in the comments section.

Table 6: Prediction market prices lead analyst revisions

	$\Delta\overline{F}_{i,t}$ Regressed on h-hour Prediction Market Price Changes				
	$h = 6$	$h = 12$	$h = 24$	$h = 48$	$h = 120$
	(1)	(2)	(3)	(4)	(5)
Intercept	-0.001	-0.001	-0.001	-0.001	0.001
t-statistic	[-0.53]	[-0.54]	[-0.81]	[-0.91]	[1.06]
$\Delta p_{i,[t-h,t]}$	0.032	0.031	0.036	0.024	0.018
t-statistic	[2.32]	[2.62]	[3.61]	[2.44]	[4.41]
Observations	380	368	356	331	208
R^2	0.014	0.018	0.036	0.018	0.086

Notes: This table shows the relationship between changes in prediction market prices and subsequent analyst forecast revisions. Specifically, the dependent variable is the revision in analyst consensus forecast, scaled by the stock price. The independent variable is the prediction market price change in the preceding 6, 12, 24, 48, and 120 hours, with each column corresponding to a different window length. The prediction market sample spans 15 September 2025 through 9 March 2026.

regression:

$$\Delta\overline{F}_{i,t} = a + b\Delta p_{i,[t-h,t]} + \epsilon_{i,t} \tag{10}$$

where $\Delta\overline{F}_{i,t}$ is the consensus forecast revision for firm i at time t, measured at the minute level, and is computed as the change in analyst consensus EPS forecast when an individual analyst revises their forecast, scaled by the stock price. $\Delta p_{i,[t-h,t]}$ represents the change in the prediction market price over the h-hour window preceding the analyst revision. A positive b means that prediction markets lead analyst revisions.

Table 6 displays the results. The coefficients are positive and statistically significant at the 5% level for all five horizons considered. Appendix Table C.1 shows five alternative regression specifications where the lag in prediction market prices is computed over days instead of hours, and the effect is positive and significant at the 5% level for four horizons, and positive and significant at the 10% level for the last. Overall, the results suggest that prediction markets lead analyst forecast revisions, thus providing more timely information to the market.

5 Why Are Earnings Markets Informative?

In this section, we analyze what makes prediction markets on earnings informative, finding (i) that earnings market traders appear to be more skilled than traders in other prediction markets, and (ii) suggestive evidence of insider trading.

5.1 Earnings Market Traders

To explore who makes earnings market prices informative, we leverage the fact that Polymarket settles trades on a public blockchain, allowing us to observe all trades and link them to individual accounts ("wallets"). We aggregate all trades to the account level and then define earnings accounts as those with at least one trade in earnings markets. We do the same for other markets on Polymarket to understand whether earnings markets are representative or special.

Table 7 shows the account-level profit-and-loss (PnL) distribution for earnings markets relative to all Polymarket accounts and various other markets. Panel A shows that the profits and losses made in earnings markets are split roughly equally, in the sense that about 50% of accounts make money while the other 50% lose. This pattern is dramatically different from the pattern when looking across all Polymarket contracts. The profit-and-loss distributions across all contracts are highly skewed, with only about 1/3 of accounts making money while the remaining 2/3 lose. This pattern is mainly observed in the two biggest markets on Polymarket—sports and politics—and is less obvious in other markets. The skewed profit-loss distribution is consistent with a minority of skilled traders making money from a majority of unskilled ones. Gómez-Cram et al. (2026) study this pattern in detail, showing that prediction markets tend to derive their accuracy from a few informed traders, rather than from the "wisdom of the crowd." But informed traders need uninformed traders to make money, and the roughly symmetric PnL distribution in earnings markets suggests that there are fewer uninformed traders than in sports and politics markets.[13]

To study this hypothesis deeper, we look at the profits made by earnings markets accounts in non-earnings markets. If the population of earnings market traders is more skilled than

[13]Table E.2 contains additional statistics about the distribution of accounts trading in earnings markets.

Table 7: PnL Distribution and Concentration by Market Category

Panel A: Account Composition									
	Earnings	All	Sports	Politics	Crypto	Economy	Finance	Tech	Culture
N (thousands)	28.3	2,271.4	1,157.6	1,178.9	975.7	479.9	352.0	335.7	609.6
Winners (%)	49.7	31.7	31.3	36.8	48.5	44.7	47.9	51.1	39.7
Losers (%)	49.2	67.8	66.5	60.6	49.0	50.2	48.4	43.1	55.0
Panel B: Share of Total Profits/Losses (%)									
	Earnings	All	Sports	Politics	Crypto	Economy	Finance	Tech	Culture
Winners:									
Top 0.1%	26.9	52.5	62.0	69.5	38.6	60.5	47.7	53.1	49.9
Top 1%	64.4	77.5	83.5	90.3	69.7	89.1	82.9	86.1	83.4
Top 5%	86.1	92.1	94.4	97.7	93.9	97.3	97.0	97.1	95.7
Top 10%	92.6	96.6	97.5	99.0	97.7	98.8	98.9	98.9	98.1
Losers:									
Top 0.1%	14.0	44.9	56.9	57.1	22.6	44.6	39.8	29.2	29.8
Top 1%	44.5	69.2	80.6	83.6	51.1	76.1	72.7	65.2	65.6
Top 5%	72.3	87.0	94.3	95.1	81.1	92.3	92.0	88.4	88.1
Top 10%	83.2	93.9	97.4	97.7	92.2	96.6	96.6	94.6	94.2

Notes: This table reports the distribution of profits and losses in earnings markets and other markets. Panel A reports the number of unique accounts and the share of winners (PnL > 0) and losers (PnL < 0). The winner and loser percentages do not necessarily sum to one, as a minority of accounts have exactly zero profits. Panel B reports the cumulative share of aggregate profits captured by the top x% of winners (ranked by PnL), and the cumulative share of aggregate losses borne by the top x% of losers (ranked by absolute loss). The `Earnings` column restricts to earnings markets. The `All` column covers all prediction markets and topic columns (`Sports` through `Culture`) restrict to markets under specific topics. The sample spans January 1, 2023 to March 9, 2026.

the typical prediction market traders, we should expect them to make money on average in these other markets. Table 8 shows evidence consistent with this interpretation. Specifically, we group accounts into 14 groups based on their position in the earnings market PnL distribution, and compute their average PnL in non-earnings markets. Almost by construction, 7 of 14 groups lose money in earnings markets, but surprisingly, 11 of 14 groups make money in the non-earnings markets. In total, 28,291 earnings market accounts made $14.28 million USD in other markets—on average 505 dollars per account. This result is consistent with earnings market traders being a relatively skilled subset of the overall Polymarket trader population.

Table 8: Trading Activities by Earnings-Market PnL Group

		Earnings Markets				Other Markets			
PnL Group	N	Trades	Markets	Volume ($)	PnL ($)	Trades	Markets	Volume ($)	PnL ($)
Bottom 1%	283	186.02	25.94	8,794.30	−1,672.02	21,576.23	962.27	2,423,529.00	11,779.38
1–5%	1,132	34.19	10.57	2,196.81	−198.65	3,242.53	363.01	255,882.97	557.82
5–10%	1,414	20.44	7.66	395.37	−52.42	4,184.06	312.77	133,392.47	941.52
10–20%	2,829	10.68	5.36	131.74	−15.76	2,420.86	303.79	61,068.66	−42.09
20–30%	2,829	6.26	3.66	51.17	−4.88	1,725.83	246.81	38,158.67	17.52
30–40%	2,808	3.74	2.34	26.42	−1.55	1,347.12	206.73	29,256.13	−77.08
40–50%	2,611	3.80	2.17	22.41	−0.41	1,629.35	273.05	21,574.95	−138.72
50–60%	3,068	2.46	1.82	43.83	0.08	684.27	180.06	52,561.63	142.84
60–70%	2,829	2.55	1.91	23.62	0.49	530.71	152.25	62,691.70	35.23
70–80%	2,829	4.22	2.68	76.76	1.82	1,084.11	207.44	138,816.64	435.34
80–90%	2,828	7.96	4.09	202.78	6.96	1,538.92	302.71	103,597.47	531.61
90–95%	1,416	20.43	8.15	587.15	24.84	3,619.06	570.51	406,801.50	2,613.88
95–99%	1,132	47.45	13.07	3,163.80	131.40	6,313.89	861.16	404,382.78	1,133.71
Top 1%	283	299.93	38.32	20,062.89	2,210.63	19,365.77	854.59	1,209,672.38	4,839.63

Notes: This table reports average trading activity by PnL group, where groups are defined by quantiles of earnings market profits and losses across all accounts that participate in earnings markets. Rows correspond to quantile intervals of the earnings PnL distribution, ranging from the bottom 1% to the top 1%. For each group, columns report the number of accounts (N) and group-average values of total trades, number of markets, trading volume, and PnL, computed separately for earnings markets and all other prediction markets. The sample spans January 1, 2023 to March 9, 2026.

5.2 Earnings Market Insiders?

Prediction markets face frequent allegations of insider trading, and, specifically for earnings markets, several people have reported suspicious activity involving companies audited by a particular auditor. We cannot prove or disprove whether insider trading exists, but we can test whether the accuracy of earnings markets differs by auditor.

To test this possibility, we use Audit Analytics to find the primary auditor for the companies underlying the earnings contracts, finding a match for 630 of 639 contracts. After excluding contracts with missing information on controls that are necessary for our analysis (specifically, the historical beat rate), we are left with a final sample of 608 contracts. We do not name the auditor because insider trading is a serious and contentious allegation, and our analysis cannot conclusively determine whether any systematic differences in accuracy across auditors reflect insider trading or instead arise from other explanations. Instead, we will refer to the auditor under suspicion as Aud1, the other three big auditors in our sample as Aud2–Aud4, and other auditors as AudOthers. The auditors cover firms associated with 105, 190, 154, 128, and 31 contracts, respectively.

Table 9 shows the Brier score by auditor for the prediction market prices and the same three benchmarks we used to evaluate overall accuracy in Table 2 Panel A. The prediction market prices for Aud1 contracts have a Brier score of 0.115, which is the lowest of all auditors, but not substantially different from the Brier score for Aud3 (0.128). However, this raw comparison fails to account for the fact that Aud3 contracts are generally easier to predict, as evidenced by the historical beat rate achieving a Brier score of 0.191 for Aud1 contracts but 0.141 for Aud3. The difference between the Brier score of prediction markets and the historical beat rate is by far the largest for Aud1, suggesting that prediction markets are more accurate for these contracts than for contracts of companies audited by other firms. Table E.4 shows a regression of the squared difference between the outcome and the prediction market price on an Aud1 dummy with and without controls for the historical beat rate, finding that the parameter estimate on the Aud1 dummy is significantly negative. This pattern is consistent with the possibility that insider trading is more prevalent for Aud1 companies, but it is, of course, not conclusive. An alternative explanation is that another

Table 9: Brier Scores by Auditor

	Aud1	Aud2	Aud3	Aud4	AudOthers
Prediction Market	0.115	0.180	0.128	0.151	0.246
Benchmarks:					
Analyst	0.250	0.250	0.250	0.250	0.250
Analyst Bias-adjusted	0.212	0.203	0.157	0.175	0.250
Historical Beat Rate	0.191	0.214	0.141	0.175	0.263
Volume ($ Million)	2.79	4.96	5.65	2.91	0.60
Number of Markets	105	190	154	128	31

Notes: This table reports Brier scores by the firm's external auditor, compared with benchmarks. Auditors are anonymized. Contracts with suspicious activity are audited by Aud1. Aud2, Aud3, and Aud4 are the three other big auditors in our sample. AudOthers contains contracts audited by other auditors. Brier scores and benchmark definitions follow Table 2 Panel A. The prediction market sample spans 15 September 2025 to 9 March 2026.

factor correlated with auditor choice, such as industry or firm size, is driving the difference in prediction market accuracy across auditors.[14]

6 A Market-Based Measure of Analyst Bias

Finally, we explore how prediction markets can inform research on analyst expectations. A large body of work studies predictable patterns in analyst forecast errors, but a central challenge in this literature is that such patterns can arise from both true cognitive biases—such as overreaction or underreaction to information—and incentive-related ones, such as analysts strategically lowering forecasts to help firms report earnings beats. Being able to measure analysts' incentive-related biases for individual firms in real time would help researchers disentangle these two sources of predictable forecast errors. Prediction markets offer exactly this opportunity.

Specifically, we compute the analyst bias for a specific firm and quarter as the difference

[14]The appendix contains a case study of potential insider trading in the 2025Q4 Wells Fargo contract. Wells Fargo has 12 quarters of consecutive earnings beats before 2025Q4, and Figure E.6 shows that the price for most of the contract duration hovered around 0.8. However, two accounts bought large volumes of the "NO" contract, bringing its price down to 0.2. The market was subsequently resolved as "NO," and the two accounts made a large profit from these trades. Table E.5 breaks down the total trading activity of these two accounts by auditor and shows that most of their volume and profits come from trading Aud1 contracts.

between the prediction market price and the 0.5 benchmark that would arise if analysts were unbiased:

$$\text{abias}_{i,t} = p_{i,t} - 0.5. \tag{11}$$

Table 10 shows that the market-based bias is strongly associated with a firm's historical beat rate, but with an R^2 of only 0.22, much of the cross-sectional variation remains unexplained. To better understand what drives analyst bias across firms, Figure 3 shows the estimates from regressing the market-based bias on eight prominent equity factor characteristics, sourced from the data set in Jensen et al. (2023). The characteristics have been signed such that a high value predicts high future stock returns and rescaled cross-sectionally to lie between zero and one. In univariate regressions, several characteristics predict analyst biases, but in a multivariate regression with all eight characteristics, only size and quality are statistically significant (size only at the 10% level, whereas quality is significant at the 5% level). Specifically, smaller firms exhibit lower bias than their larger counterparts, and high-quality firms exhibit higher bias than their junk-firm counterparts.

These results can help researchers using analyst forecasts separate cognitive biases from incentive-related ones. Specifically, our results suggest that forecast errors manifesting in small-junk firms are less likely to arise from analyst incentives than those manifesting in large-quality firms.

7 Conclusion

We introduce a high-frequency measure of earnings expectations derived from financial prediction markets. Our analysis demonstrates that prediction markets are efficient aggregators of information: they significantly outperform sell-side analysts in predictive accuracy, show less bias, possess significant incremental explanatory power for earnings announcement returns, and lead in price discovery. Earnings market traders are a relatively skilled subset of prediction market traders, making money from their trades in non-earnings markets, which could explain their informativeness. Finally, prediction markets can be used to inform re-

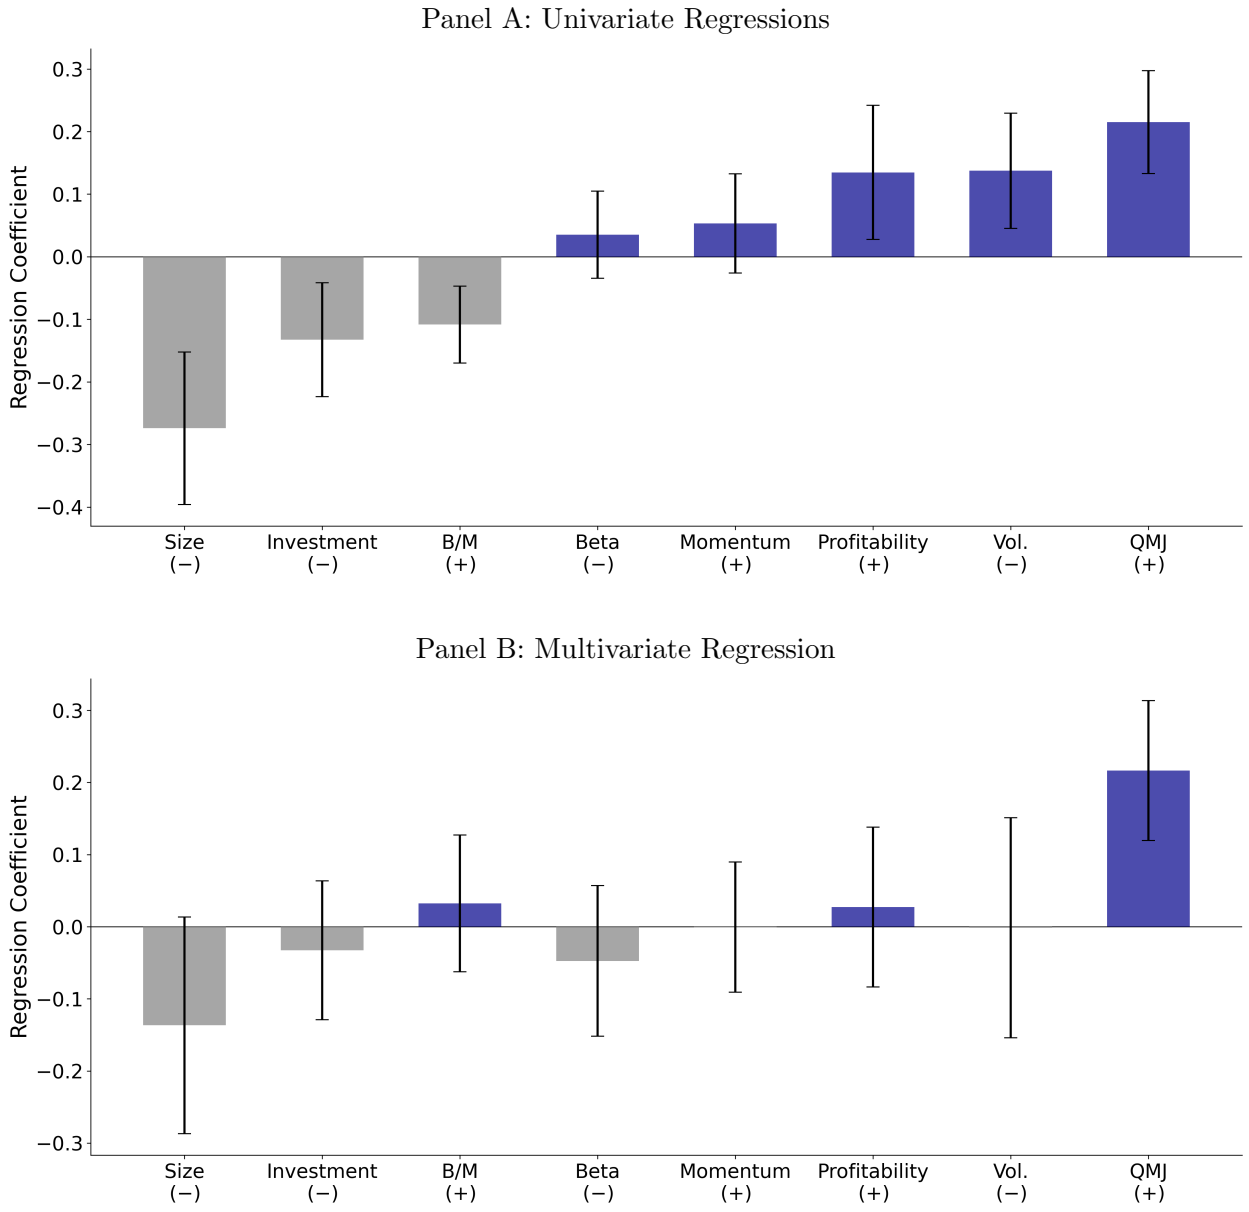


Figure 3: Analyst bias and equity factor characteristics

Notes: This figure presents coefficient estimates from regressions of the prediction market price one day before the earnings announcement on firm characteristics. Panel A reports estimates from univariate regressions, one for each characteristic. Panel B reports estimates from a single multivariate regression including all characteristics simultaneously. All characteristics are ranked cross-sectionally and normalized to $[0, 1]$ and sign-adjusted as shown in parentheses below each label. Dark blue (gray) bars indicate positive (negative) coefficient estimates. Error bars represent 95% confidence intervals based on standard errors clustered by firm and earnings announcement date. The prediction market sample spans 15 September 2025 to 9 March 2026.

Table 10: Analyst Bias, Historical Beat Rates, and Earnings Guidance

	Prediction Market Price - 0.5	
	(1)	(2)
Intercept	-0.16	0.20
t-statistic	[-4.94]	[16.02]
Historical Beat Rate	0.50	
t-statistic	[11.87]	
Guidance Dummy		0.10
t-statistic		[4.56]
Observations	613	636
R^2	0.222	0.039

Notes: This table reports results of regressing analyst bias, computed as in (11), on historical beat rates and an earnings guidance dummy. Specifically, the dependent variable is the prediction market price one day before the earnings announcement date minus 0.5. `Historical Beat Rate` is the percentage of the firm beating the analyst consensus EPS estimate in the prior 20 quarters. `Guidance Dummy` equals one if the firm issued earnings guidance prior to the announcement, and zero otherwise. t-statistics are reported in brackets. The prediction market sample spans 15 September 2025 to 9 March 2026.

search using analyst expectations, and we show one application by developing a market-based measure of analysts' incentives-related biases.

Our approach addresses the "belief-action disconnect" inherent in survey data by recovering expectations directly tied to financial incentives. Unlike analysts, who are not required to act on their stated views, our measure is derived from prices formed through actual trading, yielding earnings expectations that are backed by financial stakes. We will regularly update our earnings measure as new contracts resolve. We expect this measure to become increasingly informative as the market matures, liquidity deepens, and prediction markets become part of mainstream finance.

References

Ball, R. and P. Brown (1968). An empirical evaluation of accounting income numbers. *Journal of Accounting Research*, 159–178.

Ball, R. and P. R. Brown (2014). Ball and brown (1968): A retrospective. *The Accounting Review 89*(1), 1–26.

Bastianello, F. (2025). Time-series and cross-section of risk premia expectations: A bottom-up approach. *Available at SSRN 4204968*.

Beaver, W. H. (1968). The information content of annual earnings announcements. *Journal of accounting research*, 67–92.

Berger, P. G., C. G. Ham, and Z. R. Kaplan (2019). Do analysts say anything about earnings without revising their earnings forecasts? *The Accounting Review 94*(2), 29–52.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2025). Finance without exotic risk. *Journal of Financial Economics 173*, 104145.

Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer (2019). Diagnostic expectations and stock returns. *The Journal of Finance 74*(6), 2839–2874.

Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer (2024). Belief overreaction and stock market puzzles. *Journal of Political Economy 132*(5), 1450–1484.

Bouchaud, J.-P., P. Krueger, A. Landier, and D. Thesmar (2019). Sticky expectations and the profitability anomaly. *The Journal of Finance 74*(2), 639–674.

Brunnermeier, M., E. Farhi, R. S. Koijen, A. Krishnamurthy, S. C. Ludvigson, H. Lustig, S. Nagel, and M. Piazzesi (2021). Perspectives on the future of asset pricing. *The review of financial studies 34*(4), 2126–2160.

Bürgi, C., W. Deng, and K. Whelan (2025). Makers and takers: The economics of the kalshi prediction market.

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

Da, Z. and X. Huang (2020). Harnessing the wisdom of crowds. *Management Science 66*(5), 1847–1867.

de Silva, T., E. Larsen-Hallock, A. Rej, and D. Thesmar (2025). Expectations formation with fat-tailed processes: Evidence and theory.

Delao, R., X. Han, and S. Myers (2024). The cross-section of subjective expectations: Understanding prices and anomalies. *Working Paper*.

Delao, R. and S. Myers (2021). Subjective cash flow and discount rate expectations. *The Journal of Finance 76*(3), 1339–1387.

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

Eichengreen, B., G. Viswanath-Natraj, J. Wang, and Z. Wang (2025). Under pressure? central bank independence meets blockchain prediction markets. *Central Bank Independence Meets Blockchain Prediction Markets (July 26, 2025)*.

Engelberg, J., R. D. McLean, and J. Pontiff (2018). Anomalies and news. *The Journal of Finance 73*(5), 1971–2001.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). Five facts about beliefs and portfolios. *American Economic Review 111*(5), 1481–1522.

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

Hand, J. R., H. Laurion, A. Lawrence, and N. Martin (2022). Explaining firms' earnings announcement stock returns using factset and i/b/e/s data feeds. *Review of Accounting Studies 27*(4), 1389–1420.

Hartley, J. (2025). Recession prediction markets and rare macroeconomic disaster risk in asset prices. *Available at SSRN 3524686*.

Hong, H., J. D. Kubik, and A. Solomon (2000). Security analysts' career concerns and herding of earnings forecasts. *The Rand journal of economics*, 121–144.

Irvine, P. J. (2004). Analysts' forecasts and brokerage-firm trading. *The Accounting Review 79*(1), 125–149.

Jame, R., R. Johnston, S. Markov, and M. C. Wolfe (2016). The value of crowdsourced earnings forecasts. *Journal of Accounting Research 54*(4), 1077–1110.

Jame, R., S. Markov, and M. C. Wolfe (2022). Can fintech competition improve sell-side research quality? *The Accounting Review 97*(4), 287–316.

Jensen, T. I. (2025). Subjective risk and return.

Jensen, T. I., B. T. Kelly, and L. H. Pedersen (2023). Is there a replication crisis in finance? *The Journal of Finance 78*(5), 2465–2518.

La Porta, R. (1996). Expectations and the cross-section of stock returns. *The Journal of Finance 51*(5), 1715–1742.

Li, F. and Y. Xie (2025). Trade policy surprises and us convenience yields: Evidence from prediction markets. *Available at SSRN 5882222*.

McCarthy, O. and S. Hillenbrand (2025). Heterogeneous beliefs and stock market fluctuations. *Available at SSRN 3944887*.

Michaely, R. and K. L. Womack (1999). Conflict of interest and the credibility of underwriter analyst recommendations. *The Review of Financial Studies 12*(4), 653–686.

Mincer, J. A. and V. Zarnowitz (1969). The evaluation of economic forecasts. In *Economic forecasts and expectations: Analysis of forecasting behavior and performance*, pp. 3–46. NBER.

Ng, H., L. Peng, Y. Tao, and D. Zhou (2025). Price discovery and trading in prediction markets. *Available at SSRN 5331995*.

Norem, J. K. (2008). Defensive pessimism, anxiety, and the complexity of evaluating self-regulation. *Social and Personality Psychology Compass 2*(1), 121–134.

Norem, J. K. and N. Cantor (1986). Defensive pessimism: Harnessing anxiety as motivation. *Journal of personality and social psychology 51*(6), 1208.

Reichenbach, F. and M. Walther (2025). Exploring decentralized prediction markets: Accuracy, skill, and bias on polymarket. *Skill, and Bias on Polymarket (December 12, 2025)*.

Richardson, S., S. H. Teoh, and P. D. Wysocki (2004). The walk-down to beatable analyst forecasts: The role of equity issuance and insider trading incentives. *Contemporary accounting research 21*(4), 885–924.

Roll, R. (1988). R2. *The Journal of Finance 43*, 541–566.

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

Snowberg, E., J. Wolfers, and E. Zitzewitz (2007). Partisan impacts on the economy: Evidence from prediction markets and close elections. *Quarterly Journal of Economics 122*(2), 807–829.

Snowberg, E., J. Wolfers, and E. Zitzewitz (2011). How prediction markets can save event studies. *Working Paper, University of Utah*.

Swanson, E. T., R. Wang, and Y. Wu (2025). The effects of monetary policy on macroeconomic expectations: High-frequency evidence from traded event contracts. Technical report, Working Paper.

Thesmar, D. and E. Verner (2025). Beliefs and stock market fluctuations: New evidence from the past seven decades. *Available at SSRN*.

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

Wolfers, J. and E. Zitzewitz (2009). Using markets to inform policy: The case of the Iraq war. *Economica 76*(302), 225–250.

Online Appendix

Financial Prediction Markets:

A New Measure of Earnings Expectations

Roberto Gómez-Cram Yunhan Guo Theis Ingerslev Jensen Howard Kung

Appendix Table of Contents

A. Proofs

B. Earnings Contract Prices as Implied Probabilities

C. Price Discovery in Financial Prediction Markets

D. Translating Beat Probabilities to Level Expectations

E. Additional Figures and Tables

A Proofs

Proof of Proposition 1

Proof. Equation (1) is the standard asset-pricing relation. Using

$$\mathrm{E}_t[XY] = \mathrm{E}_t[X]\mathrm{E}_t[Y] + \mathrm{Cov}_t(X, Y),$$

we obtain

$$p_t = \mathrm{E}_t[M_{t,t+K}]\,\mathrm{E}_t[\mathbb{1}_{A_{t+K}}] + \mathrm{Cov}_t(M_{t,t+K}, \mathbb{1}_{A_{t+K}}).$$

Since the price of a risk-free claim paying one at time $t+K$ is

$$\mathrm{E}_t[M_{t,t+K}] = \frac{1}{R_{ft}^{(K)}},$$

equation (2) follows immediately. The final claim follows by comparing p_t with the physical probability $\mathrm{E}_t[\mathbb{1}_{A_{t+K}}]$. □

Proof of Proposition 2

Proof. The combined payoff of one "YES" contract and one "NO" contract is

$$\mathbb{1}_{A_{t+K}} + \left(1 - \mathbb{1}_{A_{t+K}}\right) = 1.$$

Hence, the portfolio consisting of one "YES" contract and one "NO" contract is a risk-free claim paying one at time $t+K$. By standard asset pricing, its value must therefore equal the price of a risk-free zero-coupon bond with the same maturity, namely $1/R_{ft}^{(K)}$.

In reality, the structure of Polymarket ensures that the portfolio is always priced at \$1, so if $R_{ft}^{(K)} > 1$, then the same risk-free payoff has price \$1 in the prediction market and price $1/R_{ft}^{(K)} < \$1$ in the bond market. Since the same payoff has two different prices, the law of one price is violated. □

B Earnings Contract Prices as Implied Probabilities

This section shows the conditions under which the market price of an earnings contract can be interpreted as revealing investors' subjective expectations. In fact, Polymarket itself makes exactly this interpretation. We show that the interpretation is not strictly valid, because prices also reflect the time-value of money and other risk adjustments. Nevertheless, we argue that, for the earnings contracts traded on Polymarket, the interpretation is approximately valid.

The earnings contracts traded on Polymarket are essentially contingent claims on whether an event occurs. One share of "Yes" asset will generate a payoff of one if this event happens, and zero otherwise. Polymarket sets the price of the "No" asset to one minus the price of the Yes asset, so it suffices to understand the Yes asset's price. The Yes asset has the following payoff:

$$\mathrm{CF}_{t+K} = \mathbb{1}_{\{X_{t+K} > \overline{F}_t\}} = \begin{cases} 1, & \text{if } X_{t+K} > \overline{F}_t \\ 0, & \text{if } X_{t+K} \leq \overline{F}_t \end{cases}$$

where CF_{t+K} is the cash flow at the contract expiration, which is equal to the indicator variable $\mathbb{1}_{\{x_{t+K} > \overline{F}_t\}}$ that is one if the contract-specific realized earnings, X_{t+K}, is above the contract-specific strike price (the mean sell-side analyst consensus expectations at the contract's creation date t), $\overline{F}_t$.

If a stochastic discount factor (SDF) exists, then the price of this asset is:

$$\begin{aligned} p_t &= \mathbb{E}_t\Big[M_{t,t+K}\mathrm{CF}_{t+K}\Big] && \text{(B.1)} \\ &= \mathbb{E}_t\Big[M_{t,t+K}\mathbb{1}_{\{X_{t+K} > \overline{F}_t\}}\Big] && \text{(B.2)} \\ &= \mathbb{E}_t\Big[M_{t,t+K}\Big]\mathbb{E}_t\Big[\mathbb{1}_{\{X_{t+K} > \overline{F}_t\}}\Big] + \mathrm{Cov}_t\Big(M_{t,t+K}, \mathbb{1}_{\{X_{t+K} > \overline{F}_t\}}\Big) && \text{(B.3)} \\ &= \frac{\mathbb{E}_t\Big[\mathbb{1}_{\{X_{t+K} > \overline{F}_t\}}\Big]}{R_{ft}^{(K)}} + \mathrm{Cov}_t\Big(M_{t,t+K}, \mathbb{1}_{\{X_{t+K} > \overline{F}_t\}}\Big), && \text{(B.4)} \end{aligned}$$

where p_t is the contract price, $M_{t,t+K}$ is the SDF from t to $t+k$, $R_{ft}^{(K)}$ is the gross risk-free

rate over this short K period horizon, $\mathbb{E}$ is investors' subjective expectation, and Cov is the covariance under this subjective expectation.

Therefore, to interpret the price as investors' subjective expectation of an earnings beat, we need two assumptions:

A1. $R_{ft}^{(K)} = 1$,

A2. $\text{Cov}_t\Big(M_{t,t+K}, \mathbb{1}_{\{X_{t+K} > \overline{F}_t\}}\Big) = 0.$

None of these assumptions holds exactly, but we argue that they hold approximately.

Assumption A1 implies that there is a negligible time value between the investment time t and the payoff time $t + K$. At the time of writing, the risk-free rate (which we proxy for by using the 1-month U.S. treasury bill rate) is around 4% and contracts are typically opened 10 days before the earnings resolution date. The discount rate adjustment due to the time-value of money is therefore $1/(1 + 0.04)^{10/365} = 0.9989$, which is economically close to the discount rate of 1 implied by A1.

Assumption A2 implies that the payout realization is uncorrelated with investors' SDF. This is likely approximately true for multiple reasons. First, the exact earnings of the company have already been decided, in the sense that the contract resolution date is based on the earnings announcement date, which is typically at least a month after the fiscal period has ended.[15] Second, the earnings beat or miss for a specific company probably mainly reveals idiosyncratic firm-specific news, which will not covary strongly with systematic risks faced by well-diversified investors. Third, as prediction markets on earnings are still in their infancy, the amounts wagered are relatively small, meaning the bets likely represent small fractions of any investor's overall portfolio.

Even if both assumptions are violated, the effect on the interpretation is small. To illustrate this point, suppose an investor estimates a 60% beat probability and thus an expected cash-flow of $\mathbb{E}_t\Big[\text{CF}_{t+K}\Big] = 0.6$. If the contract opens ten days before the earnings announcement, the investor is willing to pay

$$p_t = \frac{0.6}{(1+r)^{10/365}}, \tag{B.5}$$

[15] For example, Apple announced their Q4/2025 earnings on October 30th, 2025, but the associated fiscal period ended on September 27, 2025.

where r is the investor's annual discount rate. Even if the investor perceives the payoff as risky and sets the discount rate at, say, $r = 0.2$, the price is 0.597. In a more realistic case where the investor perceives the payoff as safe and discounts it at a risk-free rate of 4%, the price is 0.599. In both scenarios, the price is close to the investor's estimated probability.

C Price Discovery in Financial Prediction Markets

C.1 Prediction Markets Lead Analyst Forecasts: Robustness

Table C.1: Prediction Market Prices Lead Analyst Revisions

	$\Delta\overline{F}_{i,t}$ Regressed on d-day Prediction Market Price Changes				
	$d = 1$	$d = 2$	$d = 3$	$d = 4$	$d = 5$
	(1)	(2)	(3)	(4)	(5)
Intercept	-0.001	-0.001	0.000	0.000	0.001
t-statistic	[-0.81]	[-0.91]	[0.27]	[0.85]	[1.06]
$\Delta p_{i,[t-d,t]}$	0.036	0.024	0.010	0.018	0.018
t-statistic	[3.61]	[2.44]	[1.69]	[4.05]	[4.41]
Observations	356	331	295	251	208
R^2	0.036	0.018	0.010	0.062	0.086

Notes: This table reports the predictability of prediction market price changes on analyst forecast revisions. The dependent variable is the revision in analyst consensus forecast, scaled by the stock price. The independent variable is the prediction market price change in the preceding 1, 2, 3, 4, and 5 days, with each column corresponding to a different window length. The prediction market sample spans 15 September 2025 through 9 March 2026.

C.2 Case Study: Tesla

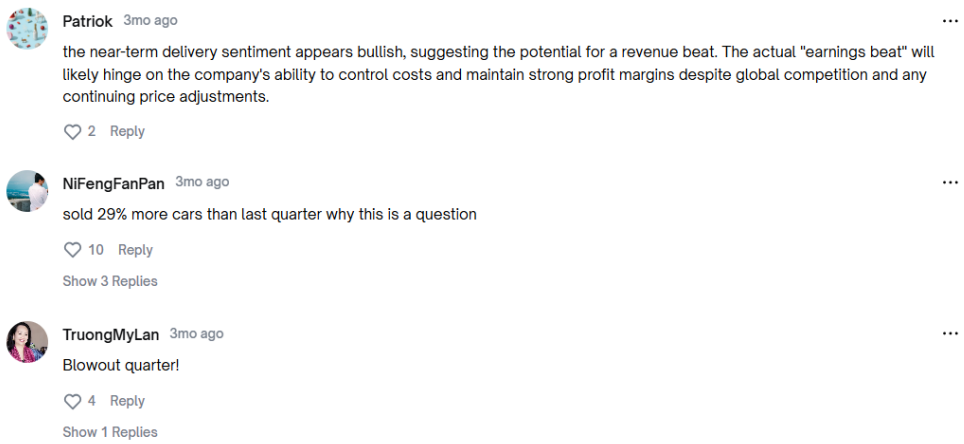


Figure C.1: Tesla 2025Q3 Earnings Prediction Market Comments

Notes: This figure shows the comments for Tesla's 2025Q3 earnings prediction market.

C.3 Case Study: JPMorgan

Figure C.2 shows the prediction market prices during the pre-announcement window for JPMorgan Chase's 2025Q4 earnings, corresponding to the fiscal quarter ending December 31, 2025. The market was created on December 31, 2025, with the earnings announcement scheduled for January 13, 2026. At the time of market creation, JPMorgan had a strong earnings track record, having exceeded analyst consensus GAAP EPS expectations in each of the previous seven quarters. Consistent with this history, the prediction market price rose rapidly after market opening and remained around 0.95 through January 7.

On January 7, JPMorgan announced on its press release website that it would replace Goldman Sachs as the issuer of the Apple Card, and that it expected to recognize a $2.2 billion provision for credit losses in 2025Q4. Following this announcement, trading activity shifted markedly. On January 8, trader *Charge2DaGame* sold 2461 shares, and the market price went down to 0.85. While the price subsequently stabilized around 0.9, the same trader sold an additional 819 shares on January 10.

On January 11, *Charge2DaGame* posted two comments referencing external news:

"Pls read about 64c per share (2.2Bn provision) headwind they announced last week . . . come on guys , easy money

https://finance.yahoo.com/news/jpm-over-apple-card-plans-171500873.html"

"The provision was announced post the market creation date so the EPS here isn't taking this into account"

These comments, shown in Figure C.3, were accompanied by the sale of an additional 282 shares. Following this information release and continued selling pressure, the market price gradually declined to approximately 0.4 ahead of the earnings announcement, during which period *Charge2DaGame* sold a further 1,320 shares.

On the morning of January 13, *Charge2DaGame* closed 1250 shares at prices around 0.5 prior to the announcement. JPMorgan subsequently reported GAAP EPS of 4.63, below the contract strike of 4.94, resulting in a miss. The market price immediately fell to nearly 0 and *Charge2DaGame* closed her remaining positions, making over $3,500 in total.

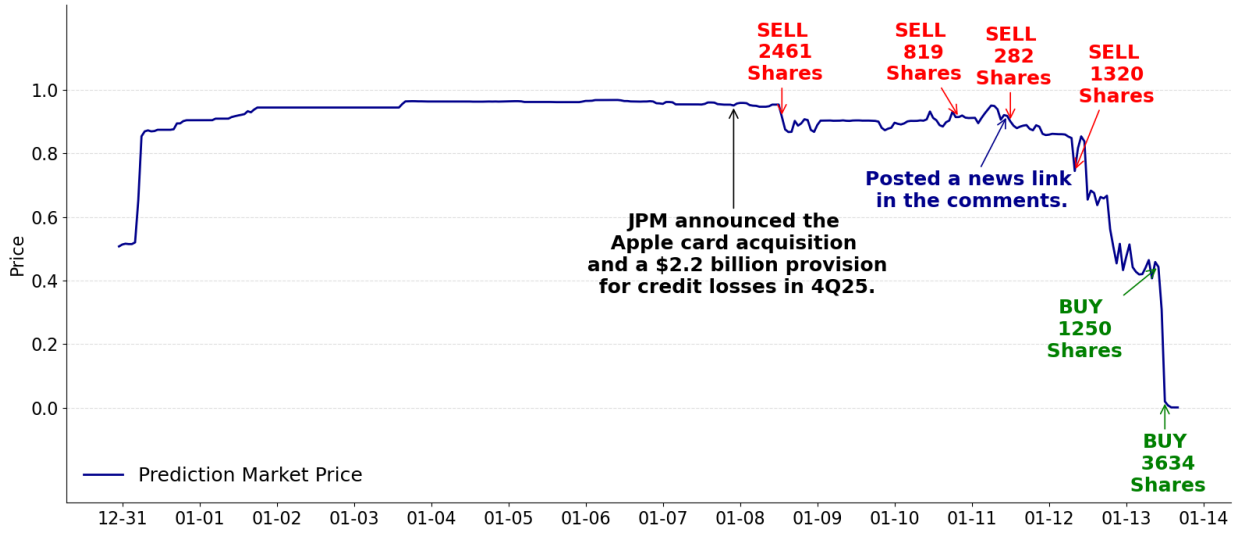


Figure C.2: JPMorgan 2025Q4 Earnings Prediction Market

Notes: This figure displays the prediction market price and trader *Charge2DaGame*'s activities for JPMorgan's 2025Q4 earnings market. The dark blue line represents the prediction market price. Annotations include JPMorgan news release (black), and trader *Charge2DaGame*'s buy (red), sell (green), and comments (dark blue).

D Translating beat probabilities to level expectations

Prediction markets provide a measure of the implied probability of an earnings beat, but many asset pricing applications need earnings level expectations. In this section, we propose a

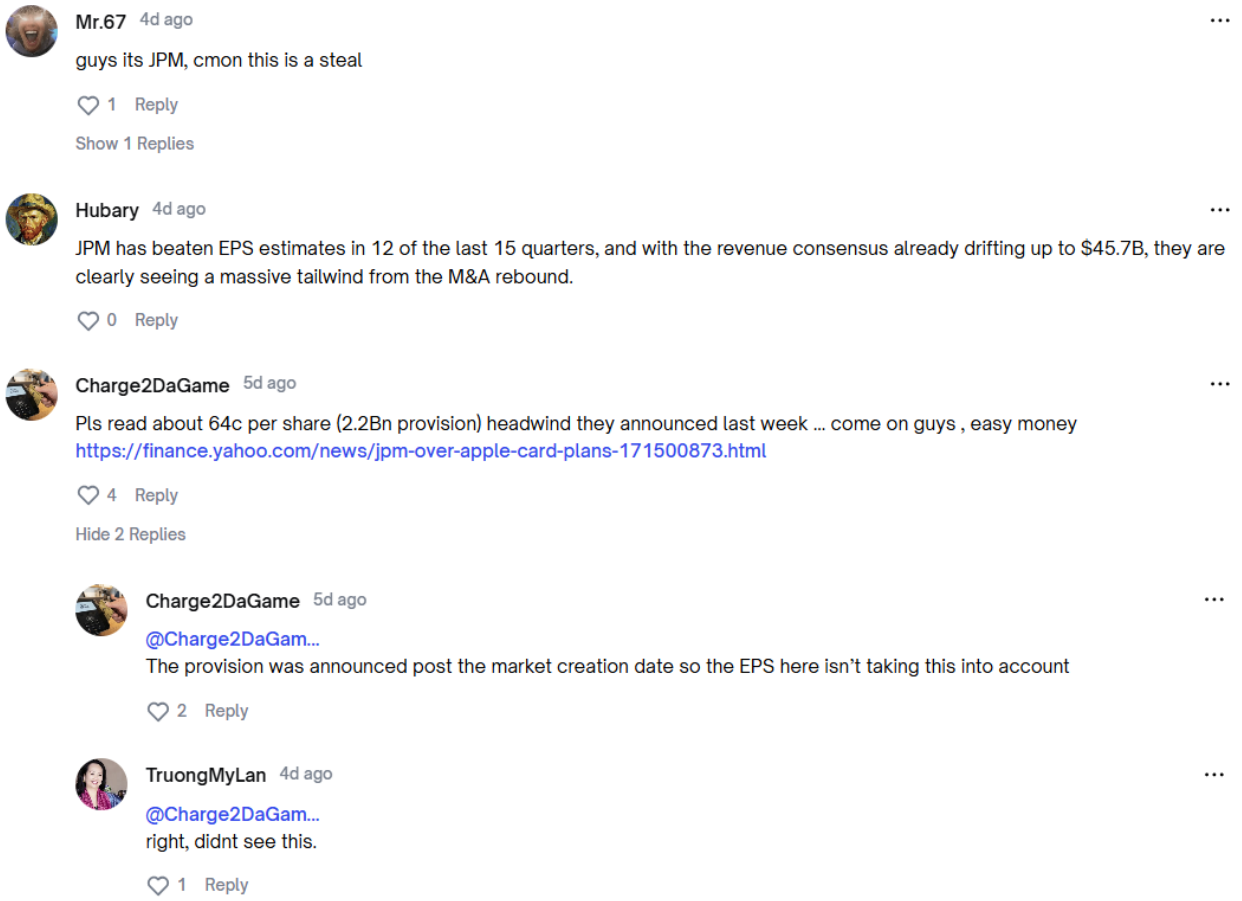


Figure C.3: JPM 20260113 Comments

Notes: This figure shows the comments for JPMorgan's 2025Q4 earnings prediction market.

methodology for making this translation. In other words, we want to translate the probability between zero and one into an unbounded level forecast, similar to those made by analysts.

We observe the prediction market price, $p_{i,t}$, which we interpret as the prediction-market probability that realized earnings exceed the analyst consensus, $\overline{F}_{i,t}$. To map a probability statement into a level statement, we impose a simple structure on the distribution of earnings surprises.

Assumption (common shape, firm-quarter specific shift). Under the subjective beliefs of prediction-market investors, earnings for firm i in quarter $t+K$, denoted $X_{i,t+K}$, satisfy:

$$\frac{X_{i,t+K} - \overline{F}_{i,t}}{p_{i,t}^{\text{stock}}} = \mu_{i,t} + \epsilon_{i,t+K}, \tag{D.1}$$

where $p_{i,t}^{\text{stock}}$ is the stock price on the day before the announcement, $\mu_{i,t}$ is a firm-quarter specific shift that captures how optimistic the market is about earnings relative to the analyst consensus, and $\epsilon_{i,t+K}$ is a random error term with a distribution that is common across firm-quarters. Let G denote the cumulative distribution function of $\epsilon_{i,t+K}$, that is, $G(x) = \text{Prob}(\epsilon_{i,t+K} \leq x)$, and let G^{-1} denote its quantile function. We normalize the error term to be mean zero, $\mathbb{E}[\epsilon_{i,t+K}] = 0$.[16]

The main assumption we impose is that heterogeneity across firm-quarters is entirely driven by a location shift of an otherwise common distribution. Importantly, scaling by price makes the assumption more plausible because raw earnings per share mechanically depend on shares outstanding, whereas the scaled version does not.

From a beat probability to a shift. Interpreting the prediction market price as

$$p_{i,t} = \text{Prob}_t \left(X_{i,t+K} > \overline{F}_{i,t} \right), \tag{D.2}$$

and using the assumed structure, we have:

$$\begin{aligned} p_{i,t} &= \text{Prob}_t \left(\frac{X_{i,t+K} - \overline{F}_{i,t}}{p_{i,t}^{\text{stock}}} > 0 \right) \\ &= \text{Prob}_t(\mu_{i,t} + \epsilon_{i,t+K} > 0) \\ &= \text{Prob}_t(\epsilon_{i,t+K} > -\mu_{i,t}) \\ &= 1 - G(-\mu_{i,t}), \end{aligned}$$

Therefore, the shift parameter implied by the market price is

$$\mu_{i,t} = -G^{-1}(1 - p_{i,t}) \tag{D.3}$$

With the shift parameter identified, we can then proceed to the main task of extracting the implied earnings-level forecast.

[16] The normalization $\mathbb{E}[\epsilon_{i,t+K}] = 0$ is without loss of generality, as we can always absorb any nonzero mean into the definition of $\mu_{i,t}$. Specifically, if $\mathbb{E}[\epsilon_{i,t}] \neq 0$, we can redefine $\mu_{i,t} \coloneqq \mu_{i,t} + \mathbb{E}[\epsilon_{i,t}]$ and $\epsilon_{i,t} \coloneqq \epsilon_{i,t} - \mathbb{E}[\epsilon_{i,t}]$, and we are back to the case where the error term has a zero mean.

Implied earnings-level forecasts. Rearranging the basic relation gives:

$$X_{i,t+K} = \overline{F}_{i,t} + p_{i,t}^{\text{stock}} \mu_{i,t} + p_{i,t}^{\text{stock}} \epsilon_{i,t+K}. \tag{D.4}$$

Taking conditional expectations given the prediction market price, plugging in the expression of $\mu_{i,t}$ from (D.3), and noting that $\mathbb{E}[\epsilon] = 0$, yields the prediction-market-implied earnings forecast:

$$\overline{F}_{i,t}^{M} = \overline{F}_{i,t} - p_{i,t}^{\text{stock}} G^{-1}(1 - p_{i,t}). \tag{D.5}$$

Implementation. We estimate G nonparametrically from the empirical distribution of realized standardized earnings surprises, $(X_{i,t+K} - \overline{F}_{i,t})/p_{i,t}^{\text{stock}}$, pooled across all firm-quarters in our prediction market sample. Figure D.1 shows the distribution. We further subtract the average earnings surprise to be consistent with the normalization we assumed in our derivation.[17]

Results. Table D.1 shows how the prediction market forecasts compare to analyst forecasts from I/B/E/S. The first two columns show that the relative precision of the two forecasts is similar. For example, the R^2 from regressing realized EPS on forecasted EPS (both scaled by prices) is 0.708 for analysts and 0.712 for prediction markets. The third and fourth columns show that, on average, prediction market forecasts are noticeably less biased than analyst forecasts. Specifically, the third column shows that the realized beat probability—that is, the percentage of cases in which realized earnings exceed forecasts—is 73% for analyst expectations but only 41% for prediction market expectations. Finally, the fourth column shows that the average standardized earnings surprise relative to analyst expectations is significantly positive, whereas the corresponding estimate for prediction markets is not significantly different from zero.

[17] An alternative approach that we explored, but ultimately decided against, is to assume that the realized earnings follow a normal distribution with a firm-quarter specific mean and standard deviation. However, this approach is problematic because (i) it requires us to estimate the conditional standard deviation for all firm-quarters—a difficult statistical task because each firm only has four earnings realizations per year and these realizations must be adjusted for seasonality—and (ii) because realized earnings surprises are far from normally distributed, as shown in Figure D.1.

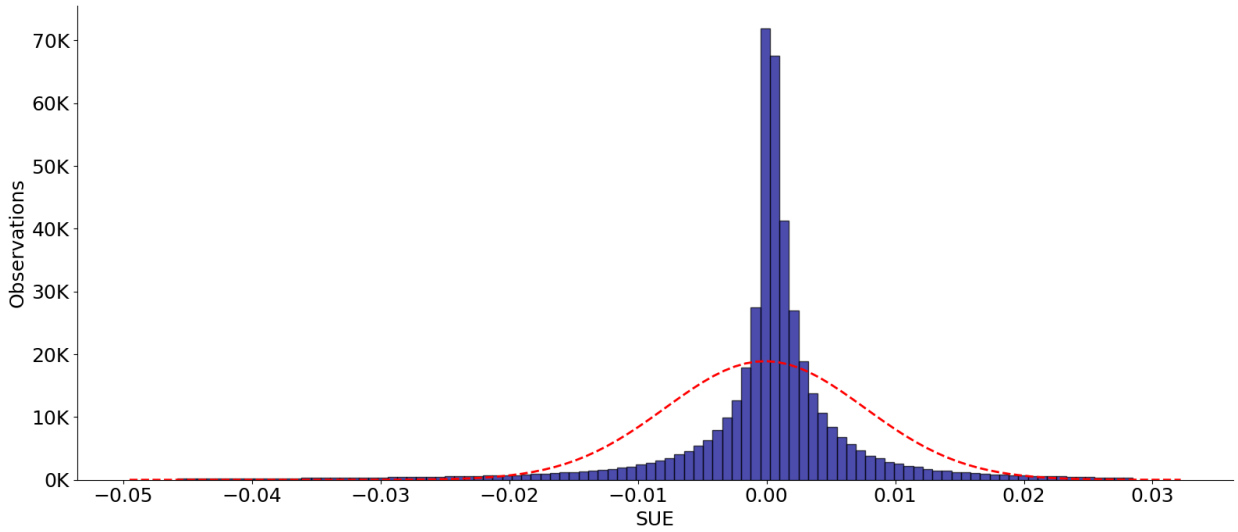


Figure D.1: Distribution of Analyst SUEs

Notes: The figure shows the distribution of standardized unexpected earnings (SUEs), defined as the difference between the realized EPS and the analyst consensus EPS scaled by the stock price on the day before the earnings announcement. The red line shows the density function of the normal distribution with the sample mean and standard deviations.

Table D.1: Precision and Bias of Prediction Market Expectations

	Precision		Bias	
Expectation	R^2	RMSE	Beat	SUE
Analyst	0.708	0.0113	73%	0.06
t-statistic				[1.97]
Prediction Market	0.712	0.0112	41%	-0.04
t-statistic				[-1.13]

Notes: This table compares analyst expectations and prediction market expectations in terms of precision and bias. For precision, we report R^2 and RMSE for both EPS expectations measures (scaled by stock prices). For bias, we report the percentage of actual EPS that beat expectations and the corresponding earnings surprises (scaled by stock prices). The prediction market sample spans 15 September 2025 to 9 March 2026.

E Additional Figures and Tables

Table E.1: Industry Distribution of Firms

	Prediction		Analysts		
NAICS-2d Industry	Obs.	%	Obs.	%	Top 3 Firms by Trading Volume
Manufacturing	173	27.07	1,675	38.41	Nvidia, Tesla, Rocket Lab
Information	136	21.28	444	10.18	Alphabet, T-Mobile, Etsy
Finance and Insurance	109	17.06	637	14.61	JPMorgan, Robinhood Markets, Coinbase Global
Retail Trade	62	9.70	127	2.91	Dick's Sporting Goods, Carmax, Amazon
Accommodation and Food Services	35	5.48	59	1.35	Domino's Pizza, Papa Johns, McDonald's
Transportation and Warehousing	28	4.38	124	2.84	Delta Air Lines, American Airlines Group, Lyft
Administrative and Support and Waste Management and Remediation Services	12	1.88	62	1.42	S&P Global, Booking Holdings, Korn Ferry
Professional, Scientific, and Technical Services	11	1.72	108	2.48	Accenture, Aecom, Omnicom Group
Arts, Entertainment, and Recreation	11	1.72	28	0.64	Draftkings, Tko Group Holdings, Flutter Entertainment Plc
Wholesale Trade	9	1.41	87	1.99	Sysco, Richardson Electronics, Fastenal

Notes: This table shows the distribution of firms covered by prediction markets and by sell-side analysts respectively, across two-digit NAICS industries. For each industry, the top 3 firms, ranked by prediction market trading volume, are also reported. The prediction market sample spans 15 September 2025 to 9 March 2026.

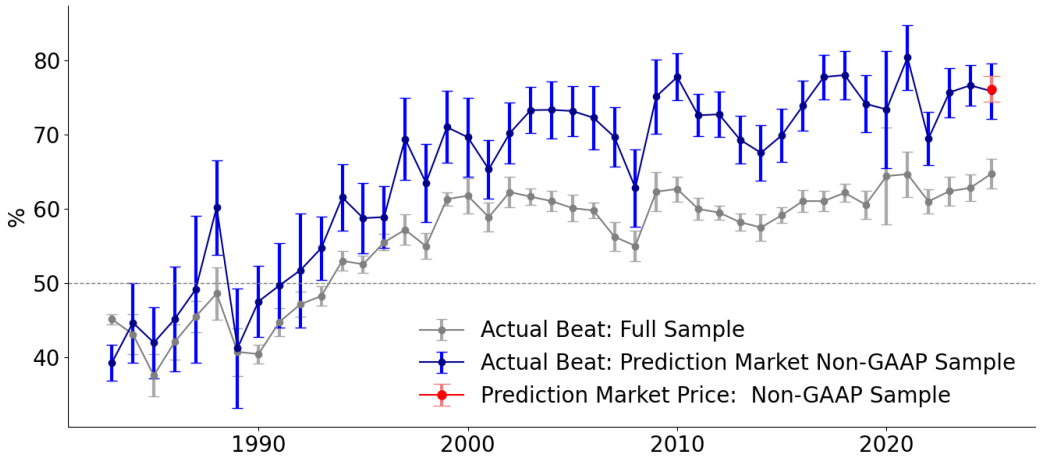


Figure E.1: Percentage of Actual Beat and Prediction Market Price

Notes: This figure plots the annual share of firms that beat the analyst consensus EPS forecast alongside the prediction market-implied probability of a beat. The dark blue line shows the beat rate for the subset of firms listed on Polymarket, while the gray line shows it for the full I/B/E/S universe. The red dot represents the average Polymarket contract price one day before the earnings announcement, with error bands denoting 95% confidence intervals. Because Polymarket contracts sometimes reference GAAP EPS and sometimes non-GAAP EPS (so-called "street earnings"), whereas I/B/E/S reports non-GAAP earnings, the gray and red data points are restricted to firms with non-GAAP contracts to ensure comparability. The prediction market sample spans 15 September 2025 to 9 March 2026. **Interpretation:** First, the figure shows that the probability of companies beating the analyst consensus has trended up over time. Second, the difference between the two lines shows that firms on Polymarket are more likely to beat than the typical firm in I/B/E/S.

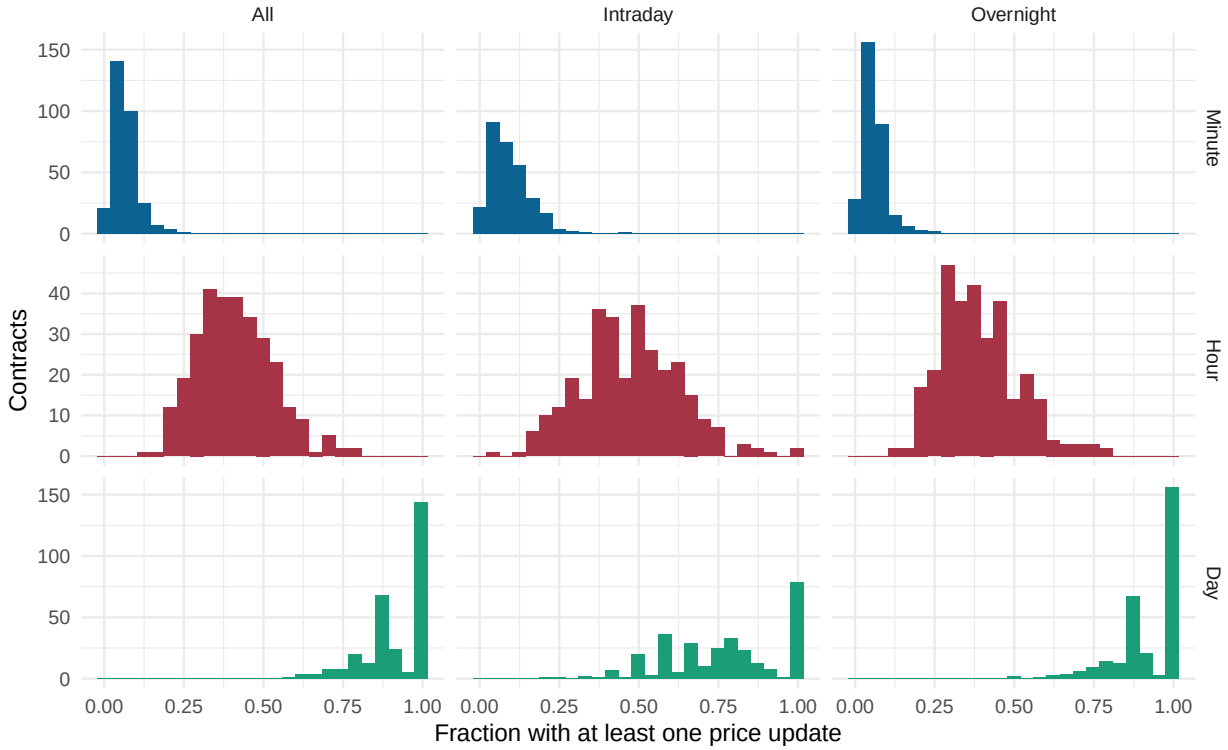


Figure E.2: Polymarket Trading Activity

Notes: This figure shows the fraction of time units with at least one price update by trading period. We consider minutes, hours, and days, and compute the fraction separately for intraday, overnight, and all observations. Intraday is defined as observations from Monday to Friday between 9:30 am and 4:00 pm in New York. The fraction is computed for each contract, that is, each observation is for a specific firm's earnings announcement.

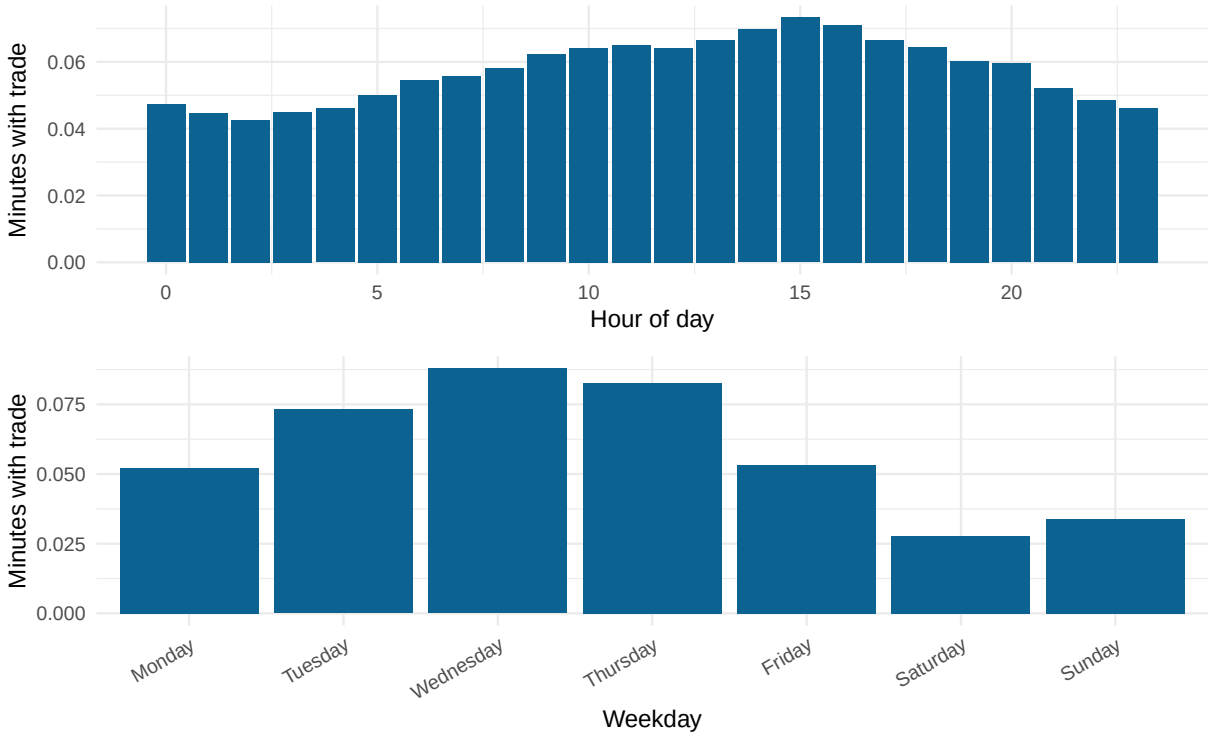


Figure E.3: Polymarket Trading Activity

Notes: This figure shows the fraction of minutes with at least one price update by hour (top panel) and weekday (bottom panel) both defined with respect to New York time.

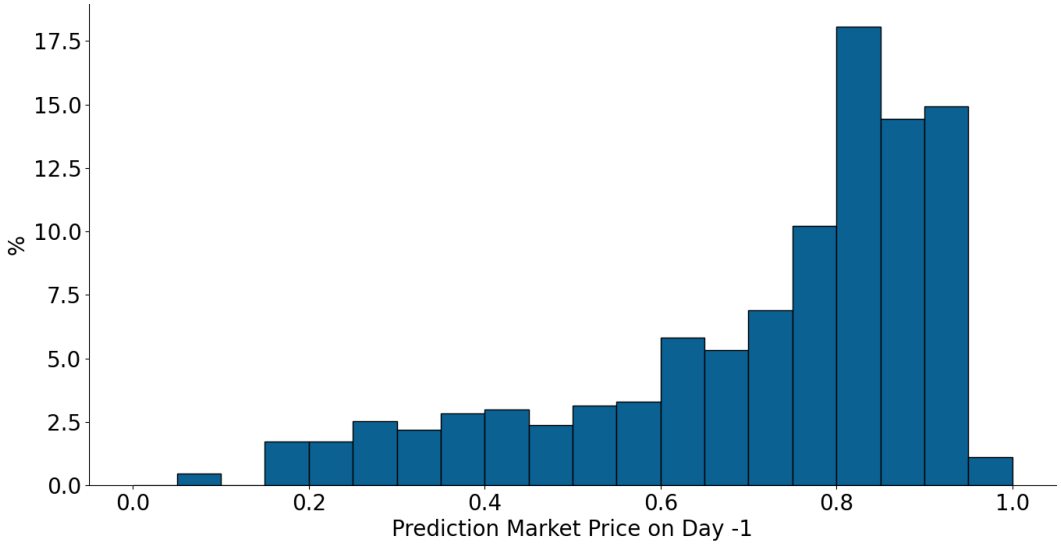


Figure E.4: Distribution of Prediction Market Prices

Notes: This figure shows the distribution of prediction market prices. The x-axis represents the prediction market prices one day prior to the earnings announcement date. The y-axis represents percentages of markets. The prediction market sample spans 15 September 2025 to 9 March 2026.

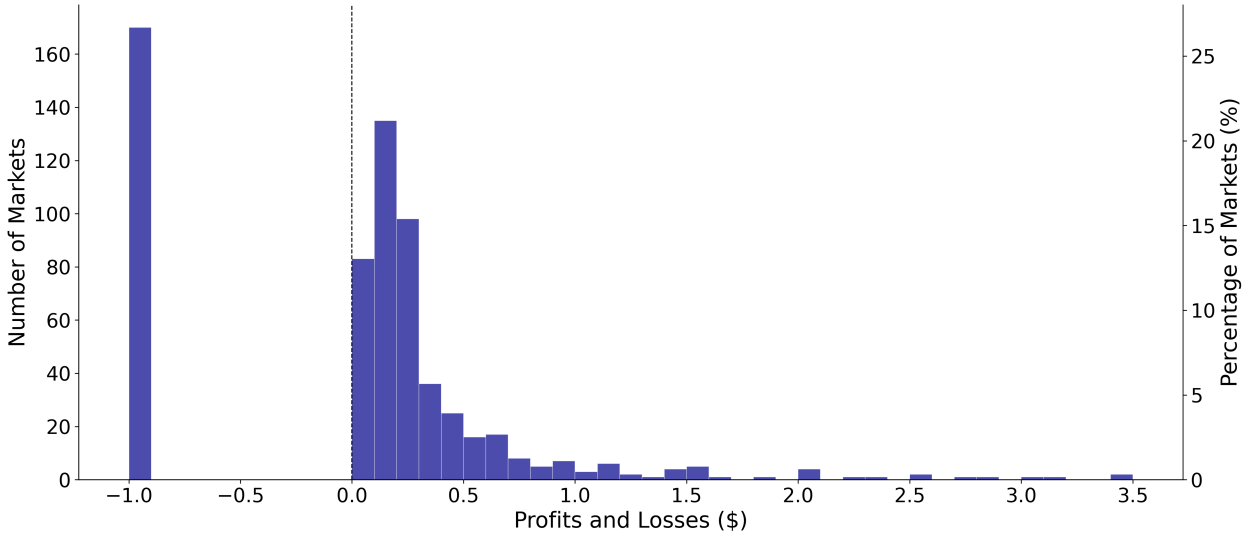


Figure E.5: Distribution of Profits and Losses

Notes: This figure presents the distribution of profits and losses (PnL) of buying 1 dollar of the Yes contract in each earnings market one day before earnings announcements. PnL is expressed in dollars. Each bin has a width of $0.10. The dashed vertical line denotes zero. The prediction market sample spans 15 September 2025 to 9 March 2026.

Table E.2: Descriptive Statistics of Account-Level Earnings Markets Trading Activity

	Number of Trades	Number of Markets	Trading Volume ($)	Total PnL ($)
Obs	28,291	28,291	28,291	28,291
Mean	14.33	4.78	610.35	−0.00
Std	141.71	17.28	12,289.32	514.56
Min	1	1	0.00	−18,795.50
1%	1	1	1.00	−505.05
5%	1	1	1.00	−91.28
10%	1	1	1.30	−29.99
20%	1	1	2.82	−8.19
30%	1	1	5.00	−2.71
40%	2	1	9.12	−1.00
Median	2	1	15.00	0.00
60%	3	2	28.70	0.22
70%	4	2	58.82	0.88
80%	7	4	128.80	3.19
90%	16	8	366.81	13.29
95%	38	17	1,100.88	46.46
99%	179	64	8,114.96	412.55
Max	10,463	597	1,277,267.50	30,206.10

Notes: This table reports descriptive statistics at the account level for all accounts that traded in earnings markets. `Number of Trades` is the count of unique transactions executed in earnings markets by the account. `Number of Markets` is the count of distinct earnings markets in which the account participated. `Trading Volume` is the total dollar value of trades in earnings markets. `Total PnL` is the realized profit and loss aggregated across all earnings markets. The sample spans September 15, 2025 to March 9, 2026.

Table E.3: Earnings Announcement Returns and Earnings Surprises

	Earnings Announcement Day Return $r_{i,t}$					
	Without r_t^m			With r_t^m		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	-0.05	-0.10	-0.12	0.16	0.10	0.08
t-statistic	[-0.13]	[-0.24]	[-0.30]	[0.39]	[0.25]	[0.20]
r_t^m	–	–	–	1.50	1.39	1.39
t-statistic	–	–	–	[4.04]	[3.76]	[3.78]
$\text{SUE}^{\text{analyst}}$	1.60	–	1.14	1.56	–	1.13
t-statistic	[4.97]	–	[3.39]	[4.90]	–	[3.41]
SUR^{pred}	–	0.06	0.05	–	0.05	0.04
t-statistic	–	[5.51]	[4.13]	–	[5.24]	[3.88]
Observations	635	635	635	635	635	635
R^2	0.038	0.046	0.063	0.062	0.067	0.084

Notes: This table reports regressions of earnings announcement returns on sell-side analyst surprises and prediction market surprises. $\text{SUE}^{\text{analyst}}$ is the difference between the actual and consensus EPS, scaled by the stock price one day before the earnings announcement. SUR^{pred} is the difference between the realized outcome dummy and the prediction market price one day before the earnings announcement. r_t^m is the value-weighted market return on the earnings announcement day. Columns (1)–(3) exclude the market return; columns (4)–(6) include it as a control. *t*-statistics are reported in brackets. The prediction market sample spans 15 September 2025 to 9 March 2026.

Table E.4: Prediction Accuracy by Auditor

	Brier Score (Prediction Market Price)	
	(1)	(2)
Intercept	0.16	0.33
t-statistic	[14.75]	[10.91]
Aud1 Dummy	−0.05	−0.05
t-statistic	[−2.38]	[−2.73]
Historical Beat Rate		−0.22
t-statistic		[−6.00]
Observations	608	608
R^2	0.007	0.049

Notes: This table reports results of regressing the contract-specific Brier Score of the prediction market price, that is, the squared difference between the outcome and the prediction market price one day before the announcement, $\left(p_{i,t} - \mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}}\right)^2$, on an Aud1 dummy variable and the firm's historical beat rate. `Aud1 Dummy` equals one if the firm's auditor is Aud1 (the auditor of firms where people have reported suspicious activity in the corresponding prediction market), and zero otherwise. `Historical Beat Rate` is the firm-quarter-specific historical beat rate computed over the preceding 20 quarters. The prediction market sample spans 15 September 2025 to 9 March 2026.

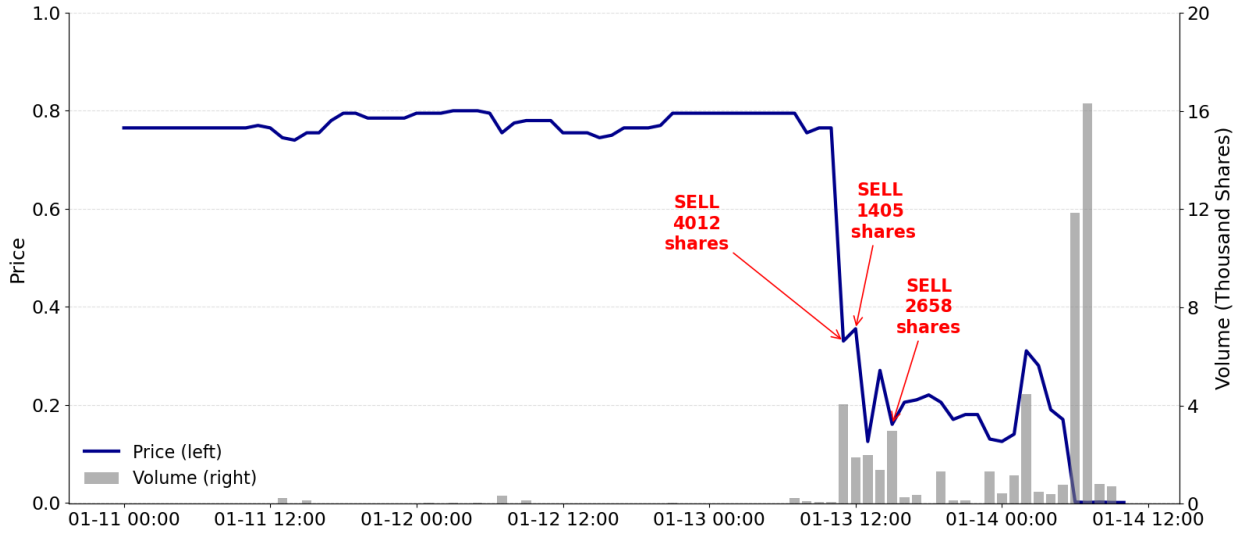


Figure E.6: Wells Fargo 2025Q4 Earnings Prediction Market

Notes: This figure displays the prediction market price, volume, and trading activities of accounts `0x4038f719b14e1fbacc31be0116605ec6efe07d42` and `0x1693a139d630da9763c5019e3ab25fe9d65e974e` for Wells Fargo's 2025Q4 earnings market. The dark blue line represents the prediction market price. The gray bar represents market trading volume in thousand shares. Annotations include trading activities of `0x4038f719b14e1fbacc31be0116605ec6efe07d42` and `0x1693a139d630da9763c5019e3ab25fe9d65e974e`.

Table E.5: Earnings Market Activities of Two WFC Winners by Auditor

Panel A: Earnings Market Activities of `0x4038f719b14e1fbacc31be0116605ec6efe07d42`

	PnL		Volume		Number of
Auditor	Amount ($)	Share (%)	Amount ($)	Share (%)	Markets
Aud1	2,055.05	86.26	12,265.71	88.13	6
Aud4	−8.35	−0.35	437.30	3.14	4
Aud2	−5.00	−0.21	5.00	0.04	1
Aud3	340.80	14.30	1,209.60	8.69	2

Panel B: Earnings Market Activities of `0x1693a139d630da9763c5019e3ab25fe9d65e974e`

	PnL		Volume		Number of
Auditor	Amount ($)	Share (%)	Amount ($)	Share (%)	Markets
Aud1	6,372.66	83.03	36,804.21	56.08	18
Aud4	1,063.03	13.85	12,567.15	19.15	9
Aud2	412.31	5.37	9,513.48	14.50	7
Aud3	−173.12	−2.26	6,741.16	10.27	8

Notes: This table reports earnings market activities of two WFC market winners by the firm's external auditor. Panel A represents account `0x4038f719b14e1fbacc31be0116605ec6efe07d42` and Panel B represents account `0x1693a139d630da9763c5019e3ab25fe9d65e974e`. PnL (volume) amount is the total profit and loss (trading volume) in dollars. PnL (volume) share is the fraction of total pnL (trading volume) accounted for by each auditor group. Number of Markets is the number of earnings prediction markets the account traded in. The prediction market sample spans 15 September 2025 to 9 March 2026.