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: March 7, 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 and introduce a methodology to convert them into implied expectations about earnings. Relative to analyst forecasts, we find that market-implied expectations are (i) more accurate; (ii) incrementally informative for earnings announcement returns; (iii) significantly less biased, though they exhibit short-term overreaction in contrast to the underreaction typical of analysts; and (iv) lead in price discovery. 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. 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.

In this paper, we construct a new measure of earnings expectations using prices from newly introduced prediction markets on company earnings announcements. 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 November 15, 2025 and serves as an initial dataset that we will continuously 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, typically 6 to 10 days prior to the announcement—and \$0 otherwise.

[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 (the platform we study) accounts for the vast majority of prediction market volume.

In these winner-take-all markets, the price can be interpreted as the market-implied probability of an earnings beat, under the assumption that risk premia are negligible. We argue that the assumption is realistic, given that the time until contract resolution is short and because a specific company's earnings beat or miss is largely firm-specific, making it unlikely that investors require a large risk premium.[2] Further, we propose a methodology to convert market-implied beat probabilities into implied expectations about earnings. Our methodology assumes that the distribution of standardized earnings surprises has the same shape for all firm-quarters, but a firm-quarter-specific median. We estimate the shape from the distribution of realized earnings surprises and use it, together with the beat probability, to estimate the implied earnings expectation.

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 market-implied beat probability correctly classifies 68% of the earnings announcements seven days before the earnings announcements and 77% on the day before. By contrast, the analyst forecasts issued after the prediction market was opened but before it closed accurately classify 62%. Prediction markets expected a high number of the contracts, 81%, to resolve as "beat" (defined as cases where the price one day before the earnings announcement was above 50%), whereas analysts only expected this to happen for 56% of the contracts. Another way to see this is to note that the average prediction market price one day before the earnings announcement was 0.71. Prediction markets were right to expect a high number of beats, as the realized beat probability in-sample was 74%.

Second, 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

[2] To illustrate why the assumption is not restrictive, suppose an investor estimates a 60% beat probability and thus an expected cash-flow of \$0.6. Given that most contracts are open for a maximum of ten days, the investor is willing to pay $p = 0.6/(1+k)^{10/365}$, where k is the investor's annual discount rate. Even if the investor perceives the payoff as risky and sets the discount rate at, say, $k = 0.2$, the price is 0.597. In a more realistic case where the investor perceives the payoff as safe and discounts it at the risk-free rate (around 4% at the time of writing), the price is 0.599. In both scenarios, the price is close to the investor's estimated probability of 0.6.

announcement return on the earnings surprise relative to analyst and prediction-market expectations, respectively, the associated R^2 is 0.036 and 0.040. When adding both surprises as explanatory variables, they both enter significantly, and the R^2 is 0.058—an increase of 61% 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 matches and even surpasses that of analysts is surprising.

Third, 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 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). However, this overreaction dissipates over longer windows, and 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.

Finally, 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.

Given the nascent nature of this prediction market, our current sample is inherently limited. We intend to continuously 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 use 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 is primarily from sell-side analysts employed by banks, and have many of the incentive-related issues mentioned in the introduction. 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 Interpreting Prediction Market Prices

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[3] 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 November 15, 2025. 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?", and participants can buy either "YES" and "NO," and a trade happens when a YES and NO side is 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 8.5 days before the earnings announcement, so think of the strike price as being the analyst consensus 8.5 days before the earnings announcement day).[4]

Polymarket interprets contract prices as the probability of the event occurring.[5] For example, they interpret a price of 0.70 as indicating a 70% probability that the firm beats the analyst consensus. But this interpretation is not quite right, as it ignores the effects of time discounting and risk premia.

To see this, let t be the time a contract is opened and $t+K$ the time it is resolved (i.e., the company earnings announcement). Further, let X_{t+K} be the realized earnings per share and $\overline{F}_t$ the analyst consensus. Then the contract payoff can be defined as the indicator, $\mathbb{1}_{\{X_{t+K}>\overline{F}_t\}}$, that resolves to 1 if $X_{t+K} > \overline{F}_t$ and 0 otherwise. Assuming the law of one price, which implies the existence of a stochastic discount factor, we can define the price of the

[3]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.

[4]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 specified for each contract ex ante.

[5]See for example, this Polymarket article about how prices are calculated, which has a section titled "Prices = Probabilities."

contract according to the following Euler equation

$$p_t = \mathbb{E}_t\Big[M_{t,t+K}\mathbb{1}_{\{X_{t+K}>\overline{F}_t\}}\Big] \tag{1}$$

$$= \frac{\mathbb{E}_t\Big[\mathbb{1}_{\{X_{t+K}>\overline{F}_t\}}\Big]}{R_{ft}^{(K)}} + \mathrm{Cov}_t\Big(M_{t,t+K}, \mathbb{1}_{\{X_{t+K}>\overline{F}_t\}}\Big) \tag{2}$$

where $M_{t,t+K}$ is the stochastic discount factor, $R_{ft}^{(K)}$ is the gross riskfree rate, and $\mathbb{E}_t[\mathbb{1}_{\{X_{t+K}>\overline{F}_t\}}]$ is the physical market probability that earnings beat the consensus forecast. From (2), we see that Polymarkets price interpretation is equivalent to assuming that the gross risk-free rate is one and that the covariance term (the risk adjustment) is zero.

While the assumptions are unlikely to hold exactly, we argue that they are probably a good approximation to the truth. In particular, the contract duration is short, implying that the risk-free rate earned is likely small, and the contract outcomes represent diversifiable firm-specific risk, implying that the risk adjustment is likely also small. Appendix A 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 assume that these assumptions hold, and interpret the prediction market price as the prediction market-implied probability of an earnings beat.

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[6]. Each request returns a minute-level time series of historical prices for that specific earnings market.

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 December 2025, I/B/E/S data are available through September 18, 2025. Consequently, our historical I/B/E/S data (spanning 1983 to September 18, 2025) only covers the first three days of the Polymarket sample period (15 September–15 November 2025).

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, on average,

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

prediction markets open 8 days prior to the earnings announcement. As shown in Column 2, the average trading volume per contract is approximately $23,000. Column 3 reports a mean pre-announcement contract price of 0.70, implying a market-implied probability of an earnings beat of 70%. This is slightly lower than the realized beat frequency of 74% reported in Column 4. The interquartile range for these probabilities is 0.61 to 0.85 (see also Figure C.3 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 $100 billion, compared to $15 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.

Table C.1 in the Appendix details the industry composition and lists the top three firms by trading activity. Figure C.1 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's 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 C.2 in the Appendix shows that the trading activity is highest around market close and on Wednesdays, whereas there is less activity during weekends.

The percentage of firms that beat analyst forecasts and the average implied probability from prediction markets are both well above the 50% that would prevail if analyst expectations were unbiased. This result is well known and is often attributed to analysts' incentives to please company management by "walking down their forecasts" to make it easier for companies to report beating earnings expectations (Richardson et al., 2004). However, the share of firms beating consensus forecasts has not always been this elevated. Figure 1 plots the yearly proportion of firms exceeding 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. For comparison, the red dot with error bands in the figure represents the distribution implied by the prediction markets. At the inception of our sample, the prediction market probabilities

Table 1: Summary Statistics

	Prediction Market Duration (days)	Trading Volume ($)	Prediction Market Price	Actual Beat	Stock Market Capitalization (Bn $)		B/M Ratio	
					Prediction	Analysts	Prediction	Analysts
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean	8.48	23,110	0.70	0.74	101.21	15.80	0.44	0.70
Std	4.59	45,886	0.19	0.44	348.68	116.34	0.40	0.86
Min	1.00	1,336	0.19	–	0.03	0.00	0.01	0.00
25%	6.00	8,104	0.61	–	3.67	0.25	0.14	0.24
Median	7.00	14,144	0.78	–	18.14	1.27	0.31	0.52
75%	10.00	25,355	0.85	–	68.42	5.64	0.59	0.89
Max	35.00	726,860	0.97	–	3,785	3,785	2.33	16.51
Obs	299	299	294	299	287	3,933	258	3,378

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 15 November 2025.

are of similar magnitude to the realized beat rate, though slightly lower.

3 Informational Content

In this section, we assess the informational content of the market-based expectations by evaluating their accuracy, how they correlate with earnings announcement returns, and testing for various biases.

3.1 Accuracy

Table 2 reports the predictive accuracy of the prediction markets. We evaluate accuracy using a confusion matrix, which compares the binary outcome implied by the market price (beat vs. no beat) against the realized earnings result.

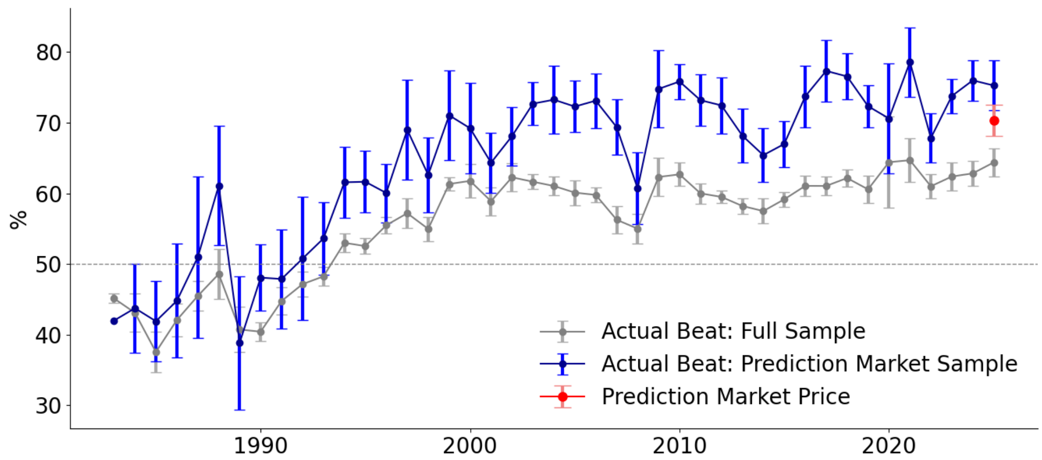

Figure 1: Percentage of Actual Beat and Prediction Market Price

Notes: This figure displays the historical actual beat percentage and prediction market price of beating analyst forecasts. The dark blue (gray) line represents the percentage of actual beat for the restricted sample of firms listed on the prediction market (full sample covered by I/B/E/S). The red dot represents the average of the prediction market prices one day before the earnings announcement. The error bands represent the 95% confidence intervals. The prediction market sample spans 15 September, 2025 to 15 November, 2025.

Table 2: Prediction Accuracy Across Earnings Windows

	Prediction Markets								Analysts	
	7 Days Before		5 Days Before		3 Days Before		1 Day Before			
Actual	Miss	Beat	Miss	Beat	Miss	Beat	Miss	Beat	Miss	Beat
Miss	4%	22%	9%	17%	11%	16%	11%	15%	16%	10%
Beat	11%	63%	9%	65%	8%	65%	7%	66%	28%	46%
Accuracy	68%		74%		76%		77%		62%	
Observations	139		258		290		294		171	

Notes: This table reports the confusion matrix comparing the prediction markets and sell-side analysts earnings prediction against actual outcomes. The prediction market predictions are shown for 7, 5, 3, 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. The prediction market sample spans 15 September, 2025 to 15 November, 2025.

We report accuracy statistics based on market prices observed seven, five, three, and one day prior to the earnings announcement. To align with the close of the U.S. stock market, we sample prediction market prices at 4:00 PM Eastern Time. Seven days before the announcement, the overall accuracy is approximately 68%. This predictive performance improves monotonically as the announcement approaches, rising to 74%, 76%, and 77% at the five-, three-, and one-day horizons, respectively.

For comparison, we also benchmark this performance against analyst forecasts issued immediately prior to the earnings announcement. Specifically, we consider all analyst forecasts issued within the short window preceding the earnings announcement but subsequent to the establishment of the consensus estimate used as the prediction market's strike price. We find that analysts achieve an overall accuracy of 62%, which is lower than the accuracy observed in the prediction markets at any of the horizons we tested. Notably, analysts predict more beats than misses, a finding that may initially seem counterintuitive given that the contract strike price is derived from the analyst consensus. However, since the strike price is fixed at contract origination, this asymmetry indicates that analysts, on average, revised their expectations upward between the contract's inception and the earnings release.

3.2 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:

$$r_{i,t} = a_1 + b_1 \text{SUE}_{i,t}^{\text{analyst}} + \epsilon_{i,t}, \tag{3}$$

$$r_{i,t} = a_2 + b_2 \text{SUR}_{i,t}^{\text{pred}} + \epsilon_{i,t}, \tag{4}$$

$$r_{i,t} = a_3 + b_3 \text{SUE}_{i,t}^{\text{analyst}} + b_4 \text{SUR}_{i,t}^{\text{pred}} + \epsilon_{i,t}, \tag{5}$$

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.[7]

Table 3 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 0.036, 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 61% (relative to analysts only) and 45% (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 4 regresses stock returns on contemporaneous changes in prediction market probabilities 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 5 basis point increase in stock returns, with a t-statistic of 2.91. Columns 2 to 5 consider narrower win-

[7] 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 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 3: Regression of Stock Returns on Analyst SUE and Prediction Market Surprise

	Earnings Announcement Return $r_{i,t}$		
	(1)	(2)	(3)
Intercept	-0.77	-0.90	-0.91
t-statistic	[-1.27]	[-1.47]	[-1.50]
$\text{SUE}^{\text{analyst}}$	1.35	–	1.00
t-statistic	[3.31]	–	[2.34]
SUR^{pred}	–	0.05	0.04
t-statistic	–	[3.48]	[2.56]
Observations	294	294	294
R^2	0.036	0.040	0.058

Notes: This table reports regressions of earnings announcement returns on sell-side analyst surprises and prediction market surprises. Analyst surprises are computed as the difference between the actual and the consensus EPS, scaled by the stock price one day before the earnings announcement. Prediction market surprises are computed as the difference between the realized outcome dummy and the market price one day before the earnings announcement. The prediction market sample spans 15 September 2025 to 15 November 2025.

dows 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.

3.3 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.

3.3.1 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 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

Table 4: Contemporaneous Stock Returns and Prediction Market Price Changes

	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	-1.35	-1.86	-1.47	-1.08	-0.88
t-statistic	[-1.78]	[-2.64]	[-2.22]	[-1.66]	[-1.46]
$\Delta p_{i,[t_1,t_2]}$	0.05	0.07	0.06	0.06	0.05
t-statistic	[2.91]	[4.39]	[3.92]	[3.57]	[3.54]
Observations	212	260	285	294	295
R^2	0.039	0.070	0.051	0.042	0.041

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 15 November 2025.

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).[8] 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 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:

$$e_{t|t+K} = a + u_t, \tag{6}$$

[8]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."

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

Panel A of Table 5 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 0.03 seems implausibly large. To see this point, note that the average contract price is approximately 0.70 and the duration is approximately 8 days. If this price reflects investors' cash-flow expectation of 0.73 plus a risk adjustment of 0.03, then it would imply a discount rate of $0.73/0.7 - 1 = 4.3\%$, which corresponds to an annualized rate of $(1 + 0.043)^{365/8} - 1 = 582.7\%$.[9]

3.3.2 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

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

Table 5: 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	0.03	0.03	0.03
t-statistic	[1.48]	[1.22]	[1.31]
Observations	294	290	258
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	0.11	0.05	0.10
t-statistic	[0.78]	[0.06]	[0.50]
$p_{i,t+K-1}$	0.89	–	–
t-statistic	[0.78]		
$p_{i,t+K-3}$	–	0.97	–
t-statistic		[0.06]	
$p_{i,t+K-5}$	–	–	0.90
t-statistic			[0.50]
Observations	294	290	258
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	0.03	0.03	0.03
t-statistic	[1.25]	[1.34]	[1.11]
$\Delta p_{i,[t+K-3,t+K-1]}$	-0.45	–	–
t-statistic	[-1.96]		
$\Delta p_{i,[t+K-5,t+K-1]}$	–	-0.31	–
t-statistic		[-1.70]	
$\Delta p_{i,[t+K-5,t+K-3]}$	–	–	-0.22
t-statistic			[-1.01]
Observations	289	255	256

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 Sep., 2025 to 15 Nov., 2025.

beating earnings are indeed more likely to do so. Formally, we estimate the regression:

$$\mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}} = \alpha + \beta p_{i,t+K-1} + u_{i,t+K}, \tag{7}$$

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 C of Table 5 reports our estimates of α and β from equation (7). 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.

3.3.3 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

$$\mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}} - p_{i,t+K-1} = \alpha + \beta \Delta p_{i,[t+K-3,t+K-1]} + u_{i,t+K}, \tag{8}$$

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 B of Table 5 reports the estimated version of equation (8). We find evidence of overreaction: the coefficient on the three-day price revision, $\Delta p_{i,[t+K-3,t+K-1]}$, is negative and statistically significant ($\beta = -0.45$, t-statistic $= -1.96$). This result 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).

However, when we extend the revision window to five days (using $\Delta p_{i,[t+K-5,t+K-1]}$), the estimated coefficient is no longer statistically significant. Similarly, regressing the forecast error $\mathbb{1}_{\{X^i_{t+K} > \overline{F}^i_t\}} - p_{i,t+K-3}$ on the lagged price change $\Delta p_{i,[t+K-5,t+K-3]}$ yields no significant coefficient. These patterns suggest that market-implied expectations tend to overshoot at very short horizons, but this relationship dissipates over longer intervals. We acknowledge, however, that the statistical power of these tests is constrained by our current sample size and that some of the apparent overreaction could reflect market microstructure noise. As the market matures and more data becomes available, we expect to estimate these coefficients with greater precision.

3.4 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 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. 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 B.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

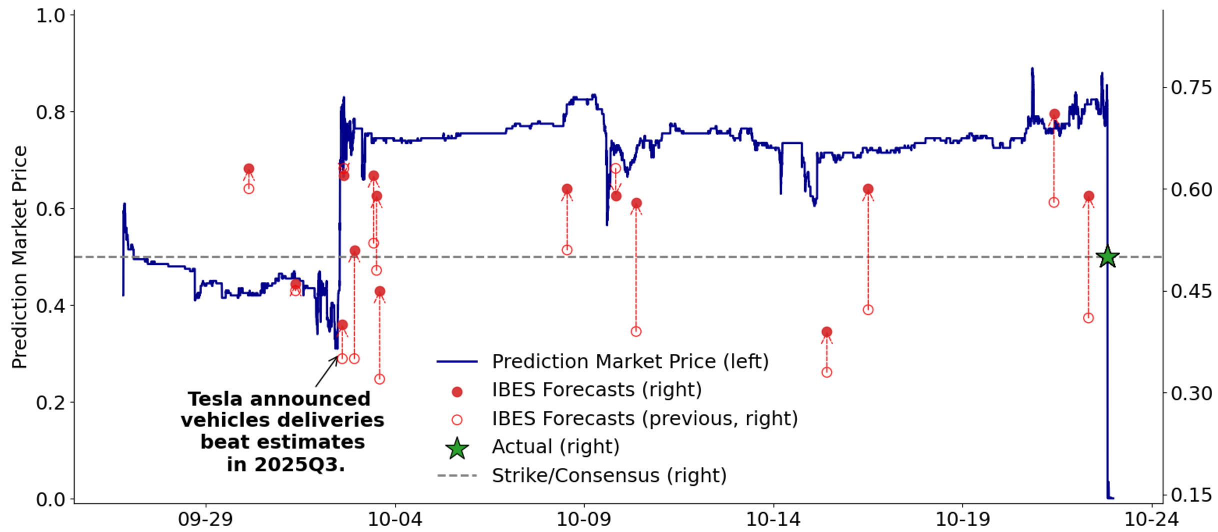

Figure 2: Case study: Tesla's 2025Q3 earnings

Notes: This figure compares the prediction market price and analyst forecasts for Tesla's 2025Q3 earnings. The darkblue 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.

prices contained more timely information than analysts.[10]

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 regression:

$$\Delta\overline{F}_{i,t} = a + b\Delta p_{i,[t-h,t]} + \epsilon_{i,t} \tag{9}$$

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

[10] Appendix B.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 15 November 2025.

the 5% level for all five horizons considered. Appendix Table B.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.

4 Translating beat probabilities to level expectations

Having established that prediction market probabilities differ from analyst expectations, and that this difference contains information, we now turn to the difficult task of translating the prediction market-implied beat probability into a prediction market-implied earnings forecast. 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:

$$\frac{X_{i,t+K} - \overline{F}_{i,t}}{p_{i,t}^{\text{stock}}} = \mu_{i,t} + \epsilon_{i,t+K}, \tag{10}$$

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, $\mathbb{E}[\epsilon_{i,t+K}] = 0$.[11]

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

$$p_{i,t} = \text{Prob}_t\left(X_{i,t+K} > \overline{F}_{i,t}\right), \tag{11}$$

and using the assumed structure, we have:

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

[11] 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.

Therefore, the shift parameter implied by the market price is

$$\mu_{i,t} = -G^{-1}(1 - p_{i,t}) \tag{12}$$

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

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

$$X_{i,t+K} = \overline{F}_{i,t} + p_{i,t}^{\text{stock}}\mu_{i,t} + p_{i,t}^{\text{stock}}\epsilon_{i,t+K}. \tag{13}$$

Taking conditional expectations given the prediction market price, plugging in the expression of $\mu_{i,t}$ from (12), and noting that $\mathbb{E}[\epsilon] = 0$, yields the prediction-market-implied earnings forecast:

$$\overline{F}_{i,t}^{M} = \overline{F}_{i,t} - p_{i,t}^{\text{stock}}G^{-1}(1 - p_{i,t}). \tag{14}$$

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 C.4 shows the distribution. We further subtract the average earnings surprise to be consistent with the normalization we assumed in our derivation.[12]

Results. Table 7 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 0.731 for analysts and 0.734 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—

[12] 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 C.4.

that is, the percentage of cases in which realized earnings exceed forecasts—is 75% for analyst expectations but only 46% 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.

Table 7: Precision and Bias of Prediction Market Expectations

	Precision		Bias	
Expectation	R^2	RMSE	Beat	SUE
Analyst	0.731	0.0130	75%	0.13
t-statistic				[2.48]
Prediction Market	0.734	0.0129	46%	0.05
t-statistic				[0.82]

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 15 November 2025.

5 Conclusion

We introduce a high-frequency measure of earnings expectations derived from financial prediction markets and develop a methodology to recover implied EPS forecasts. Our analysis demonstrates that prediction markets are efficient aggregators of information: they significantly outperform sell-side analysts in predictive accuracy and possess significant incremental explanatory power for earnings announcement returns. Furthermore, while analysts systematically underestimate earnings and exhibit underreaction, prediction market expectations are unbiased in levels and largely consistent with rational benchmarks, though we do observe short-term overreaction. Finally, we show that prediction markets lead price discovery for earnings-related information, incorporating new information before analyst forecast revisions.

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.

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.

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., 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 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:

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

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 riskfree 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. $\mathrm{Cov}_t\left(M_{t,t+K}, \mathbb{1}_{\{X_{t+K} > \overline{F}_t\}}\right) = 0.$

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.[13] 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 $\mathbb{E}_t\left[\mathrm{CF}_{t+K}\right] = 0.6$. If the contract opens ten days before the earnings announcement, the investor is willing to pay

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

[13] 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.

B Price Discovery in Financial Prediction Markets

B.1 Prediction Markets Lead Analyst Forecasts: Robustness

Table B.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 15 November 2025.

B.2 Case Study: Tesla

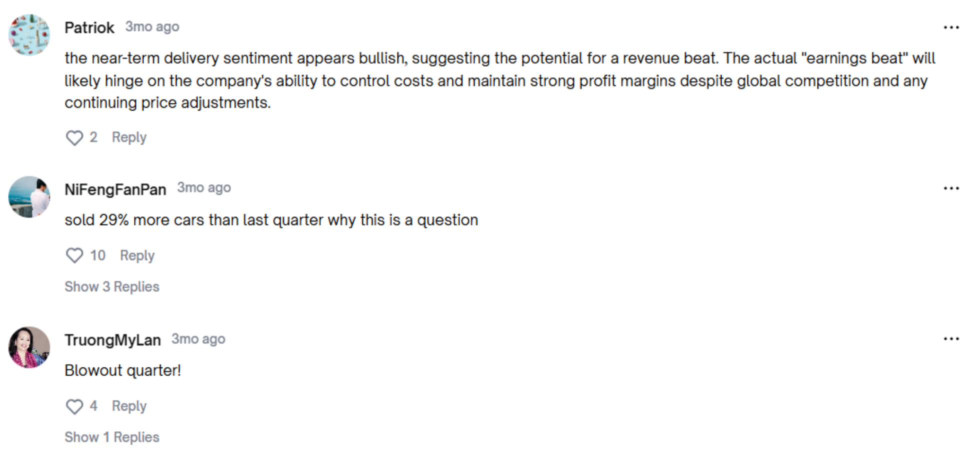

Figure B.1: Tesla 2025Q3 Earnings Prediction Market Comments

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

B.3 Case Study: JP Morgan

Figure B.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, JP Morgan announced on their 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 B.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 $ 3500 in total.

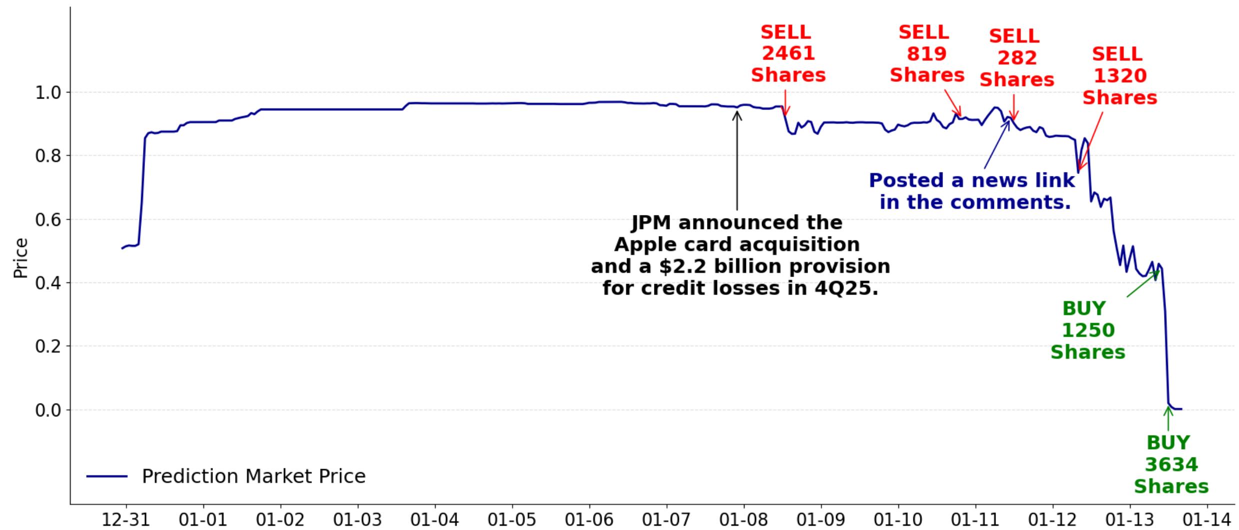

Figure B.2: JP Morgan 2025Q4 Earnings Prediction Market Comments

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

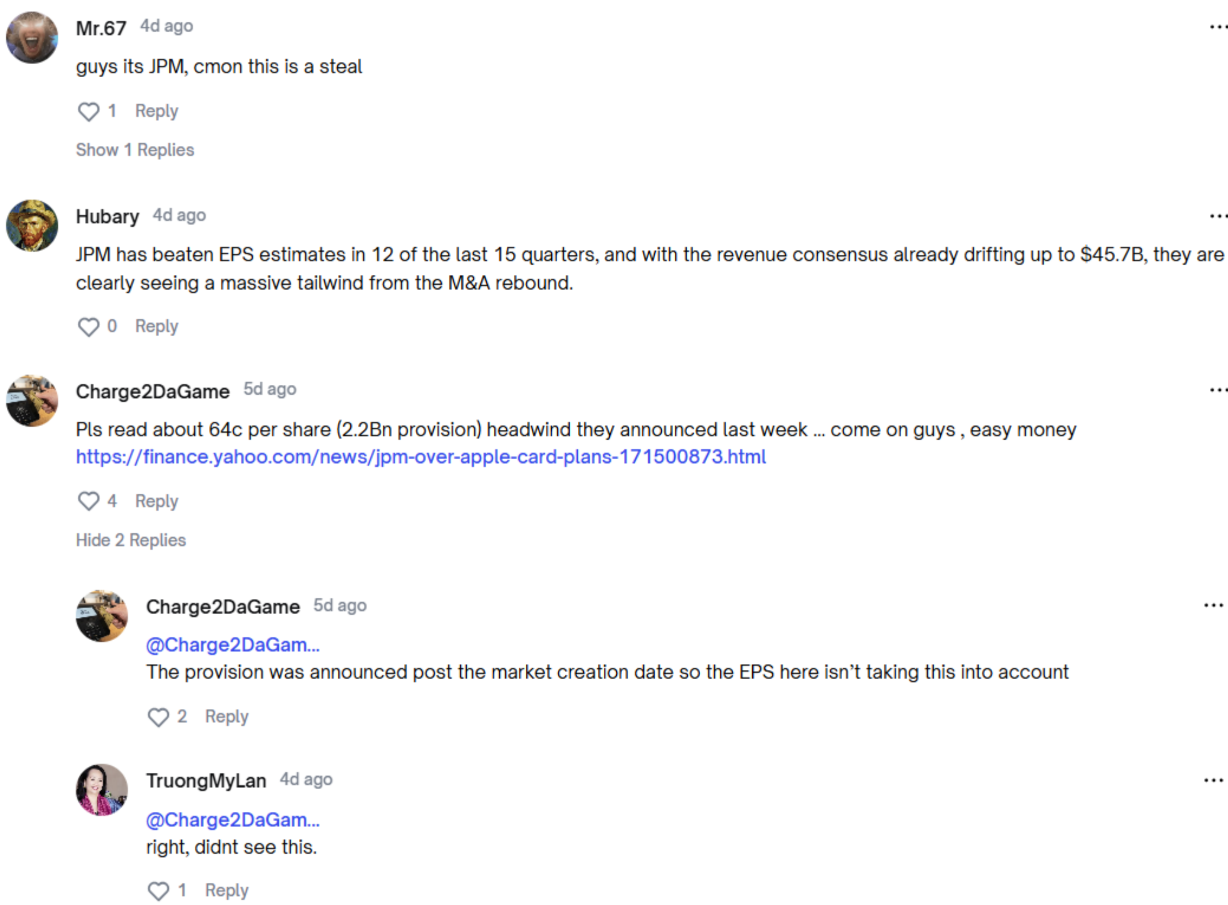

Figure B.3: JPM 20260113 Comments

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

C Additional Figures and Tables

Table C.1: Industry Distribution of Firms

	Prediction		Analysts		
NAICS-2d Industry	Obs.	%	Obs.	%	Top 3 Firms by Trading Volume
31–33: Manufacturing	87	29.10	1,665	38.14	Tesla, Nike, Apple
52: Finance and Insurance	62	20.74	636	14.57	Robinhood Markets, Pathward Financial, Sofi Technologies
51: Information	56	18.73	440	10.08	Alphabet, Palantir Technologies, Netflix
72: Accommodation and Food Services	15	5.02	60	1.37	Mcdonald's, Starbucks, Airbnb
48–49: Transportation and Warehousing	15	5.02	121	2.77	Uber Technologies, Delta Air Lines, Virgin Galactic Holdings
44–45: Retail Trade	9	3.01	125	2.86	Carmax, Amazon, Bed Bath & Beyond
54: Professional, Scientific, and Technical Services	6	2.01	109	2.50	Accenture, Omnicom Group, Gambling.Com Group
42: Wholesale Trade	6	2.01	84	1.92	Sysco, Richardson Electronics, Fastenal
62: Health Care and Social Assistance	5	1.67	64	1.47	Quest Diagnostics, Teladoc Health, Humana
71: Arts, Entertainment, and Recreation	5	1.67	28	0.64	Draftkings, TKO Group Holdings, Flutter Entertainment

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 15 November 2025.

Figure C.4 shows the distribution of scaled earnings surprises estimated from all firm-quarters in our prediction market sample, which is the distribution we use for G. The distribution has heavier tails than a normal distribution, as shown by overlaying the normal distribution

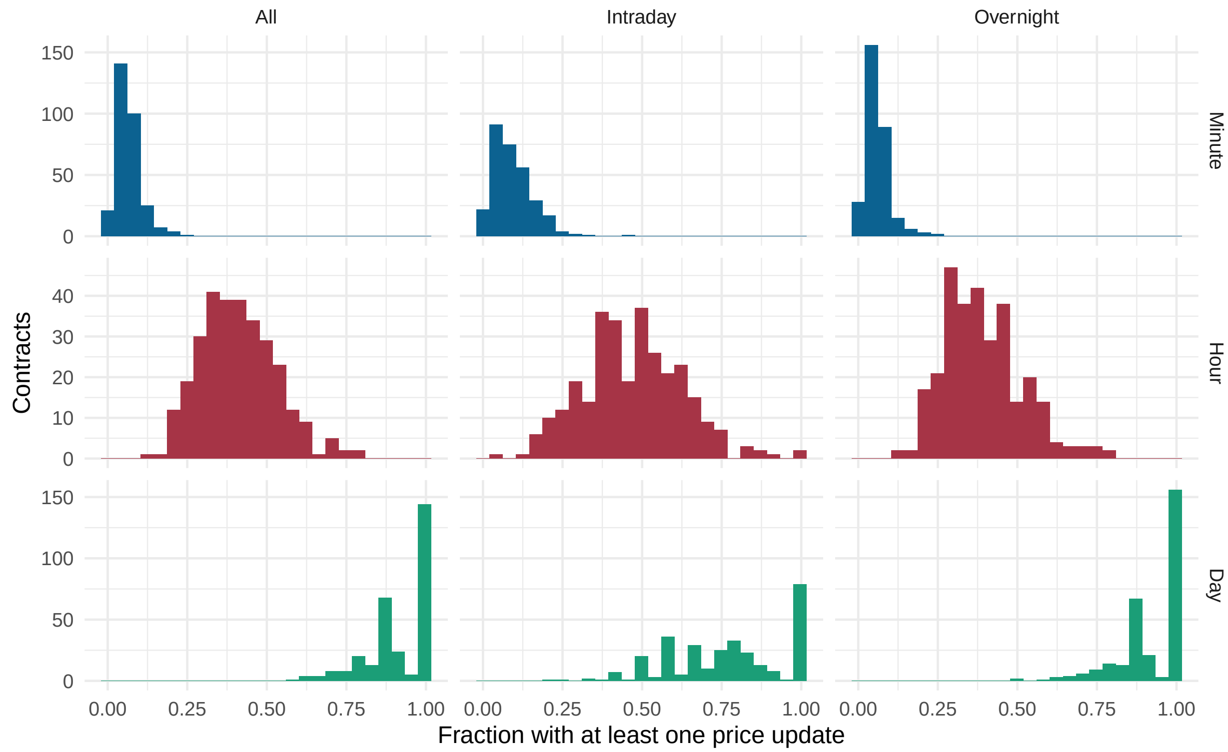

Figure C.1: Polymarket Trading Activity

Notes: This figure shows the fraction of time units with a 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.

with the sample mean and standard deviation.

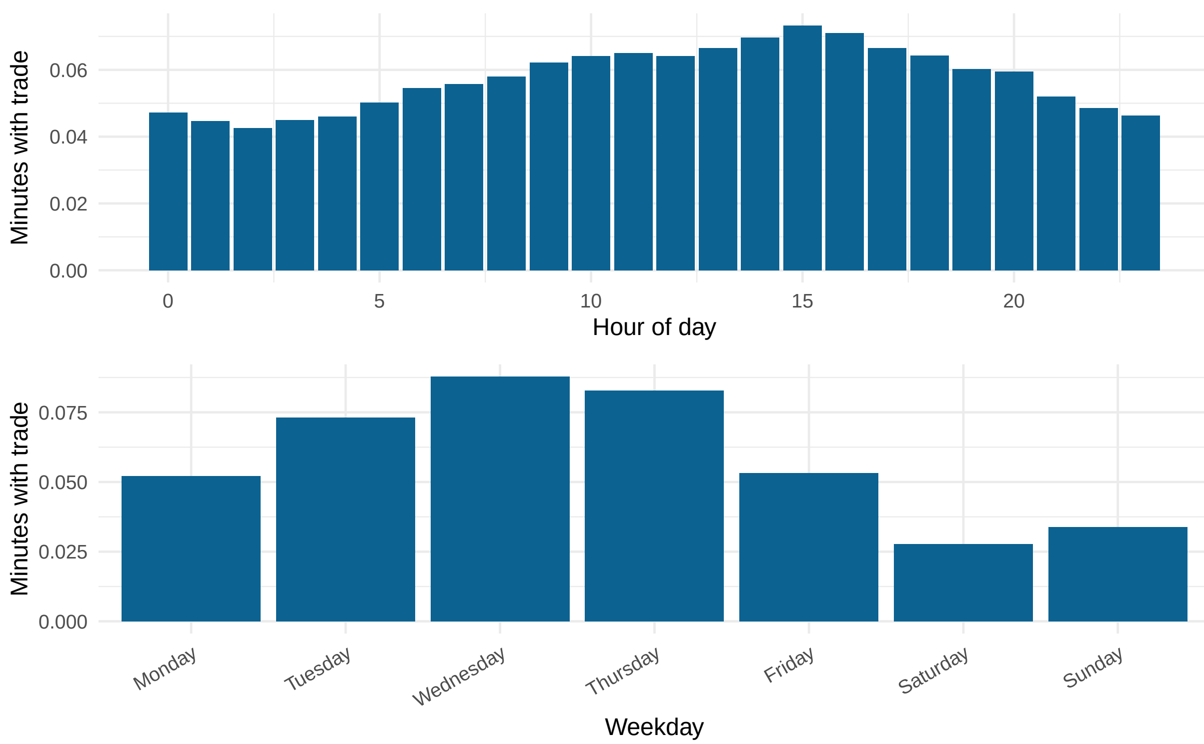

Figure C.2: Polymarket Trading Activity

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

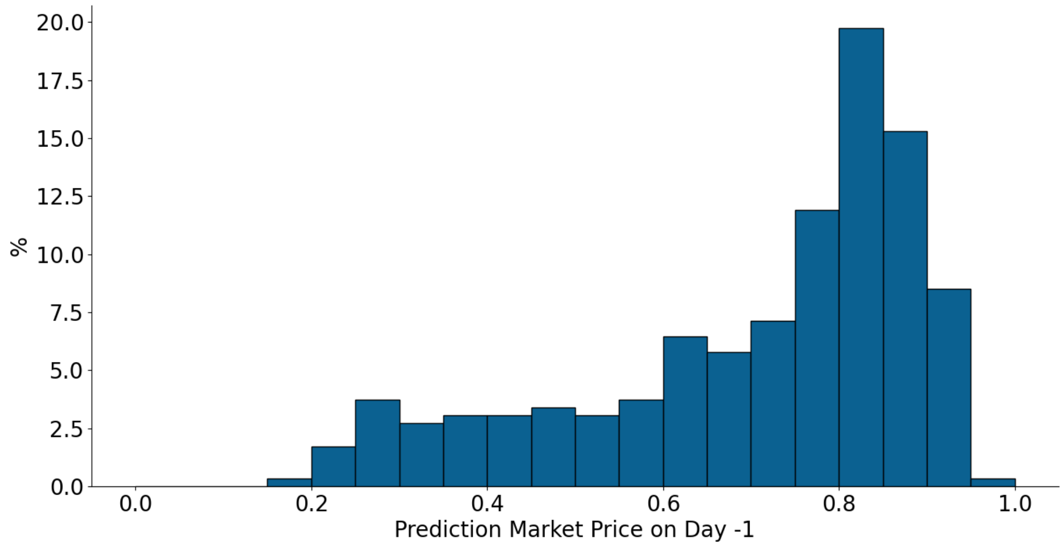

Figure C.3: Distribution of Prediction Market Prices

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 15 November 2025.

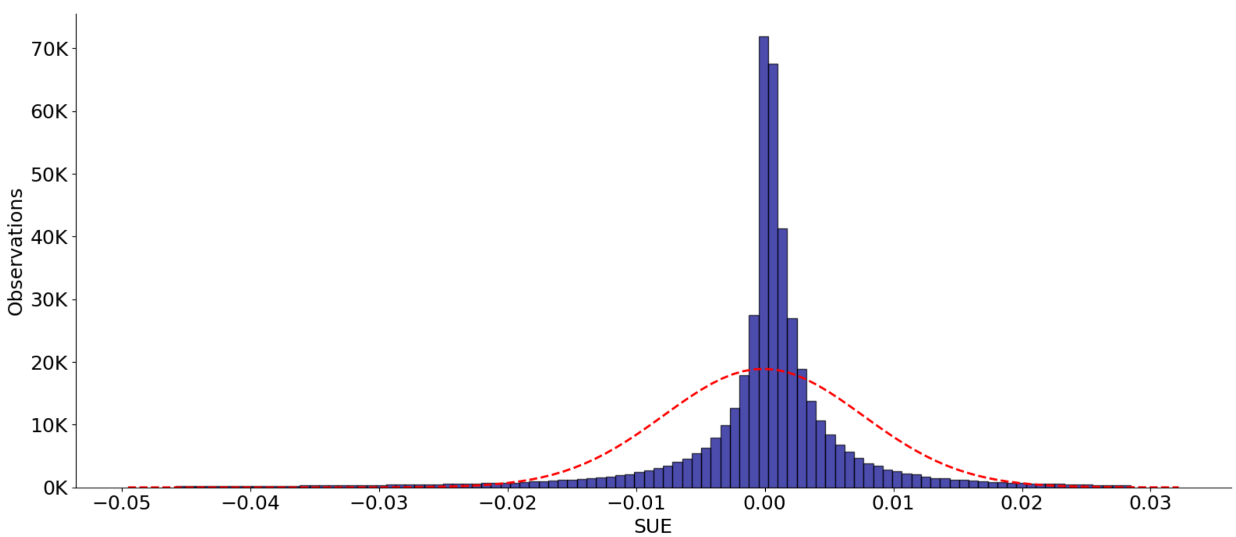

Figure C.4: Distribution of Analyst SUEs

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.