Underreaction to Known Changes in Fundamental Value: Evidence from a Laboratory Experiment and Ex-Dividend Days

Tania Diaz-Bazan*
Yale School of Management

September 7, 2026

Abstract

A long literature has documented that on ex-dividend days stock prices fall by less than the dividend, even though this represents an anticipated and publicly known change in the value of the stock; on average, across 432,582 U.S. events since 1963, the price drop is only three quarters per dollar of dividend. The standard explanations are taxes and market microstructure. I show that the same underreaction arises in a laboratory market where neither applies. In the Yale Stock Trading Game, a rich market simulation environment, a public announcement reveals that the fundamental value of a stock has changed by an exact dollar amount; across 232 announcements in 58 trading sessions, prices move by only 0.177 per announced dollar at the first trade and 0.278 after ten trades, even though participants clearly notice the announcement and trading activity increases when it occurs. I show that this underreaction pattern can be explained within the framework of cognitive uncertainty of Enke and Graeber (2023): investors who are uncertain about the fundamental value of a stock stay close to its pre-event price (the cognitive default) and adjust only part of the way toward the new value. This behavioral explanation is consistent with the fact that we see the underreaction in the lab setting, where standard explanations are ruled out. It also makes an additional testable prediction: that we should see more underreaction where fundamental value is more uncertain. I verify this prediction empirically both in the game and in the market response on ex-dividend days.

1 Introduction

A fundamental prediction of the efficient market hypothesis is that prices react correctly to news about the fundamental value of the asset. As is well known, evaluating whether this is the case requires a model of the correct reaction (Fama, 1991). A widely-studied exception to this is the case of dividend payments: on ex-day, the dividend distribution is completely known ex ante, and therefore in a frictionless efficient world (in which Modigliani-Miller holds), the price of the stock paying the dividend should drop by exactly the amount of the dividend.

That prices indeed do *not* fall on ex-day by the entire amount of dividends on average is a puzzle first documented more than 70 years ago (Campbell and Beranek, 1955; Readett, 1956; Durand and May, 1960), and replicated numerous times since then. In the sample I build from CRSP to replicate this result, that includes 432,582 ex-dividend events from 1963 to 2025, the price falls by

*Yale School of Management. tania.diaz-bazan@yale.edu. I am grateful to Roger Ibbotson for sharing the Stock Trading Game data.

about three quarters of the dividend on average (−0.767), and by nine tenths after 2000 (−0.896), both significantly different from −1.

The two standard explanations in the literature are taxes and market microstructure. Tax-clientele models beginning with Elton and Gruber (1970) explain the puzzle based on the differential tax treatment of dividends and capital gains: in particular, they predict a drop smaller than the dividend when the marginal investor is taxed more heavily on dividends than on capital gains. Microstructure explanations point to tick-size discreteness, bid–ask bounce, and transaction costs (Bali and Hite, 1998; Frank and Jagannathan, 1998). While both likely play a role in explaining the puzzle, the empirical evidence on both mechanisms is mixed.[1] A related literature (Kalay, 1982; Lakonishok and Vermaelen, 1986; Michaely and Vila, 1995) studies the effect on the equilibrium price drop of the presence of arbitrageurs (short-term traders who buy the stock cum-dividend and sell it ex-dividend). In these theories the shortfall still originates in taxes, but the equilibrium price drop is bounded by the costs and risk of this trade, and, therefore, when such traders are marginal it reflects their costs and not the tax rate of any long-term holder.

In this paper, I study the price reaction to a publicly known change in the fundamental value of a stock in a laboratory setting: the Yale Stock Trading Game. This is a rich market simulation environment, in which the standard explanations for the ex-day puzzle do not apply: taxes are absent by design, and the microstructure explanations can be ruled out directly. I show that in this setting prices severely underreact to the publicly known change in fundamental value, just like they do on ex-dividend days, suggesting that the underreaction may be due to behavioral factors. I propose a behavioral explanation for this underreaction by adapting the cognitive uncertainty framework of Enke and Graeber (2023): I take the pre-event price as the *cognitive default* and assume that greater fundamental uncertainty increases cognitive noise, strengthening investors' tendency to anchor to that price. The mechanism relates closely to recent research showing that investors lack conviction about what the price level of stocks should be even when given information about fundamentals, and that this can explain observed patterns in asset prices (Hartzmark and Sussman, 2025); and to recent evidence that people who are uncertain about their own judgment compress their responses to new information (Enke and Graeber, 2023) and underreact to strong, precise signals in particular (Augenblick et al., 2025). Under these assumptions, the model predicts that underreaction is larger where fundamental value is more uncertain. I examine this additional prediction in the Stock Trading Game and in the stock market, on the sample of ex-dividend events.

The Stock Trading Game is a laboratory market experiment designed and maintained by Roger Ibbotson at the Yale School of Management. My sample covers 58 sessions played between 2020 and 2024 by undergraduate and graduate students, alumni, and executive-education participants. The game was built to be a realistic market, in which participants trade 4 stocks for a total of 25 minutes, and features public and private information, an order book, leverage, and short-sale constraints. It

[1] Evidence supporting a role for taxes is in Poterba (1986), Green and Rydqvist (1999), Bell and Jenkinson (2002), Elton et al. (2005), Zhang et al. (2008), and Kreidl (2020); qualifications to tax-based interpretations are discussed in Boyd and Jagannathan (1994), Michaely (1991), Chetty et al. (2005), Eades et al. (1984), and Grinblatt et al. (1984). Evidence supporting microstructure explanations is in Bali and Hite (1998) and Frank and Jagannathan (1998); challenges to these explanations are discussed in Graham et al. (2003).

was not designed to study the reaction to known changes in fundamental value like the one that occurs on an ex-dividend day, but it turns out to be especially well suited to that purpose, for three reasons. First, once per stock in every game, a change in the fundamental value of that stock is publicly announced to all participants as an exact dollar amount, similarly to the distribution of the dividend on an ex-dividend day. Second, this setting has no taxes and no trading commissions, and other microstructure effects that have been studied in the literature can be ruled out. Third, the data are very rich: every order, every transaction, and every piece of information available to participants is recorded with a precise timestamp.

I document that across all events in which a change in fundamental value of a stock is announced in the Stock Trading Game (there are 232 of these events across the 58 sessions), prices move by 0.177 per dollar of change at the first post-announcement trade (SE 0.033), and by 0.278 after ten trades, thus displaying severe underreaction to the actual change in fundamentals. Taxes cannot explain it, because there are none. Nor can the standard microstructure explanations for the ex-dividend puzzle: first, every announced change in fundamental value is a whole number of dollars (thus a multiple of the tick size, \$1), and the shortfall is the same for announcements of \$20 or more, which rules out the rounding mechanism of Bali and Hite (1998). The quote midpoint, which is free of bid–ask bounce, moves even less than transaction prices, which rules out the microstructure mechanism of Frank and Jagannathan (1998). And there are no commissions or trading fees at all.

The shortfall observed in the lab is also not due to inattention. While limited attention is a leading behavioral account of underreaction to other events like earnings news (DellaVigna and Pollet, 2009; Hirshleifer et al., 2009), the announcement here is scheduled, displayed to every participant very prominently, and one of only three public events per stock in the entire game. In fact, I do see participants visibly act on it. Immediately after the announcement of the change in value of a stock, order submissions for the announced stock rise sharply for that stock, and the number of distinct active traders rises by about half, while activity in the other stocks in the game falls at that time. Therefore, participants clearly notice the announcement and trade on it; what they do not do is move the price by the amount the change would imply.

After documenting the underreaction and ruling out both the standard explanations and limited attention, I investigate a potential behavioral explanation for this phenomenon. I adapt the cognitive uncertainty framework of Enke and Graeber (2023) to explain this underreaction. In their framework, a decision maker who can form only a noisy estimate of her best decision does not act on it in full: she starts from a *cognitive default*, the choice she would make without deliberating, and adjusts only part of the way toward her estimate, the less so the noisier the estimate. In my application, the decision is whether to trade a stock and at what price, and I take the prevailing pre-announcement price as the cognitive default: the participant (who has her own set of public and private signals) understands that the price does not necessarily reflect the true value of the stock, but trading at the current price is the natural choice for someone with no view of her own: that price is the cognitive default. A participant who is uncertain about the stock's value understands that the announced change shifts the fundamental value, and her own estimate of it, by the full amount; but the default,

the visible pre-announcement price, has not moved. In this application, her offer to buy or sell the stock is then a weighted average of her estimate and the pre-announcement price, with more weight on the price the noisier her estimate. Averaged across participants, prices end up biased toward the pre-announcement price, which is observed as underreaction to the known change in fundamental value.

The model has one additional testable prediction: that underreaction is larger where participants have more uncertainty about the fundamental value. I study this prediction empirically in the Stock Trading Game and in the stock market. In the Stock Trading Game, I compare announcement events across stocks where the uncertainty about the fundamental value is plausibly different. Because uncertainty about the fundamental value cannot be observed directly, I rely on two proxies. I first study early announcements within each game and compare them to later announcements, and show that the very first announcement (which happens just five minutes into the game, so where the participants had the least information about the fundamental value) has significantly larger underreaction than the later announcements (whose pattern is however nonmonotonic and noisily estimated). Second, I study announcements that are preceded by different amount of private signals purchased by participants, and find significantly less underreaction when participants have acquired more private information.

Finally, I move back to the field evidence to specifically investigate the role of uncertainty about the fundamental value as a potential driver of underreaction (controlling for the other factors that have been studied in the literature). The ex-dividend sample is far larger than the lab sample and offers better proxies for uncertainty about fundamental value. I use two: the stock's realized volatility over the year before the event, the cross-sectional dispersion across analysts of forecasts of the firm's long-term growth (LTG), which does not depend on the stock price (but only on its long-term fundamental value). Holding the effects (documented in the literature) ofsize, turnover, and dividend yield constant, I show that the ex-day shortfall (underreaction) is indeed larger for stocks whose value is more uncertain: the ex-day slope rises from about -0.95 in the bottom quintile of volatility to about -0.73 in the top quintile, and it is 0.121 to 0.187 less negative in the top quintile of LTG dispersion than in the bottom quintile, depending on the minimum number of analysts I require to include the event ($t = 2.46$ and $t = 3.24$).

I also examine whether the pattern reflects limits to arbitrage (Shleifer and Vishny, 1997; Pontiff, 2006): the estimated volatility gradient is somewhat larger among small stocks than among large ones (0.349 against 0.252 among low-turnover stocks), but a strong gradient remains even among large and liquid stocks, which raises the question of why arbitrageurs do not correct the mistake at least for those stocks. The answer is that volatility is itself a risk that discourages arbitrage. Michaely and Vila (1995) make this point specifically for the ex-dividend day, where the risk of the position limits how much any trader is willing to hold, and Pontiff (2006) makes it for arbitrage in general. Here the trade that would remove the gradient requires concentrated positions in exactly the stocks with the highest volatility. Overall, my evidence is consistent with the underreaction on

the ex-day being driven at least in part by behavioral factors related to cognitive uncertainty, with arbitrageurs facing significant costs and risks in correcting the mistake.

The paper makes two main contributions to the literature. First, it provides novel evidence from a laboratory experiment, which provides a setting of public and deterministic change in fundamental value that is similar to the distribution of dividends on ex-day, but where the standard explanations for that phenomenon can be ruled out. Second, it shows that, in line with a theory of underreaction due to fundamental uncertainty, the underreaction is stronger when investors have more fundamental uncertainty, providing evidence both from the lab and from the stock market, using different proxies for uncertainty.

In addition to the various papers mentioned above on the ex-dividend day puzzle, this paper relates to a few other strands of the literature. First, it relates to behavioral accounts of how investors treat dividends and price levels, most notably Hartzmark and Solomon (2013, 2019). It also relates to a literature on the psychology of updating under uncertainty: anchoring and adjustment (Tversky and Kahneman, 1974), cognitive uncertainty (Enke and Graeber, 2023; Woodford, 2020), and underinference from strong signals (Augenblick et al., 2025). The paper also relates to models of underreaction based on conservatism, overconfidence, or slow diffusion (Barberis et al., 1998; Daniel et al., 1998; Hong and Stein, 1999). I study whether uncertainty about the level of fundamental value impedes price adjustment even when the change in value is public and known exactly; dividend neglect alone does not specify this uncertainty mechanism. Related evidence links information uncertainty to return continuation (Zhang, 2006; Jiang et al., 2005) and valuation difficulty to investors' behavioral biases (Kumar, 2009). The paper also relates to the literature that uses cross-analyst forecast dispersion as a proxy for differences of opinion and fundamental uncertainty (Miller, 1977; Diether et al., 2002; Johnson, 2004; Barron et al., 1998), and dispersion in long-term growth forecasts in particular (Anderson et al., 2005; Hong and Sraer, 2016). Finally, it relates to experimental asset markets, where prices famously fail to track a known, declining fundamental value (Smith et al., 1988; Palan, 2013), in part because of confusion about the fundamental-value process (Lei et al., 2001; Kirchler et al., 2012).

The rest of the paper proceeds as follows. Section 2 builds the ex-dividend sample and replicates the puzzle, both unconditionally and as a function of dividend yield, size, and turnover. Section 3 describes the Stock Trading Game and documents the underreaction to the change in fundamental value in that setting. Section 4 proposes a mechanism for the underreaction, built on the cognitive uncertainty framework of Enke and Graeber (2023), and derives the prediction that underreaction is larger where fundamental value is more uncertain. Section 4.1 examines this prediction in the game, and Section 4.2 shows that the underreaction in price movements on ex-day in the stock market lines up with proxies for fundamental uncertainty, as predicted by the theory. Section 5 concludes.

2 The Ex-Dividend Puzzle

I begin by replicating the key results of the ex-dividend puzzle literature in an updated sample of U.S. stocks. I estimate the slope of the regression of the ex-day price change on the dividend that was distributed on ex-day, both scaled by the cum-day price. Following the literature, I estimate this relationship both unconditionally and as a function of dividend yield, size, and turnover (all of which play a role in the potential explanations of the puzzle proposed by the existing literature).

Sample. I use CRSP daily data for ordinary common stocks on NYSE, AMEX, and Nasdaq, and ordinary cash dividends declared and paid in U.S. dollars. Following the literature, I require nominal cum-day prices to be at least \$5 and nominal dividends at least \$0.01, and exclude the lowest one percent of dividends as a fraction of cum-day prices (that is, the case where dividends were a trivial fraction of the price; in practice, the cutoff dividend is 0.067% of the price). The resulting panel contains 440,173 events from 1963 to 2025; the market adjustment described below requires 60 prior trading days, which leaves 432,582 events in the main regressions (Appendix Table A.1). Appendix A reports the filters and timing conventions in detail. Table A.2 in that appendix reports summary statistics for the variables used in the analysis: the median dividend payment as fraction of the price is 0.7 percent, and the median ex-day price change is −0.45 percent.

Specification. Following the literature, the dependent variable in my main specification is the market-adjusted relative price change on ex-day,

$$y_{it}^{\text{adj}} = \frac{P_{it}^{\text{ex}} - P_{it}^{\text{cum}}}{P_{it}^{\text{cum}}} - \widehat{\alpha}_{it} - \widehat{\beta}_{it} r_{M,t}, \tag{1}$$

where the market-model coefficients are estimated over the 60 trading days strictly preceding the ex-day (Brown and Warner, 1985). This price change on ex-day is then regressed on the dividend paid as a fraction of the cum-day price: $x_{it} = D_{it}/P_{it}^{\text{cum}}$:

$$y_{it}^{\text{adj}} = \alpha + \beta x_{it} + \varepsilon_{it} \tag{2}$$

A slope of −1 in this regression means that the price falls by the full dividend. Standard errors are clustered by ex-date, which allows for dependence across all firms sharing the same ex-day.

Unconditional slope. The first column of Table 1 reports the estimates of regression (2) in the full sample, which gives a slope of $\hat{\beta} = -0.767$, strongly statistically different from −1; the second column shows a slope of −0.896 after 2000 (still significantly different from −1). These results were obtained using close-to-close price changes. Measuring the change from the cum-day close to the ex-day open instead gives a slope of −0.776 (column 3, which starts in 1992, when opening prices become available); these results show that the underreaction (relative to the slope of −1) is not created by news arriving during the ex-day. Finally, column 4 shows the close-to-close slope

Table 1: Ex-dividend price adjustment: complete baseline regressions

	Full sample	Post-2000	Close-to-open	Unadjusted
Constant	0.00041***	0.00055***	−0.00018*	0.00139***
	(0.00008)	(0.00013)	(0.00011)	(0.00013)
Dividend yield	−0.767***	−0.896***	−0.776***	−0.801***
	(0.009)	(0.018)	(0.014)	(0.010)
Years	1963–2025	2000–2025	1992–2025	1963–2025
R^2	0.039	0.043	0.059	0.037
Events	432,582	153,920	205,776	440,057
Ex-dates	15,663	6,481	8,364	15,665

Note: Column (1) shows the results of regression (2) on the full sample (1963-2025), using close-to-close returns. Column (2) repeats the analysis after 2000. Column (3) uses close-to-open returns, which are available from 1992, and subtracts half of the full-day market adjustment. Column (4) uses the full sample and does not adjust for the market. Standard errors cluster by ex-date; ***, **, * denote significance at 1%, 5%, and 10% confidence levels. The benchmark of a fully offsetting price decline is a slope of −1.

without any market adjustment. All the slopes in the table are strongly statistically different from −1, confirming the existence of the ex-dividend puzzle in my sample.

Yield, size, and trading conditions. As discussed in the previous section, the existing literature has shown that the extent of underreaction of prices on ex-day is related to a number of stock characteristics. The most important ones are the dividend yield, the firm size, and trading activity. Figure 1 replicates these results. Specifically, I construct quintiles for each of these three variables,[2] and run a joint regression of the price change y_{it}^{adj} on the dividend paid relative to the price (x_{it}) interacted with each quintile dummy. The figure plots the estimated average slope in each quintile. The regression coefficients are reported in Appendix Table C.1. The yield and size gradients are large, negative and significant: the slope moves from −0.59 in the lowest yield quintile to −1.09 in the highest, and from −0.65 among the smallest firms to −1.10 among the largest. Turnover adds little once size and yield are included.

[2]All quintiles are relative to other stocks in the ex-day sample. Quintiles of size (market cap) are relative to the other stocks in the same ex-day; quintiles of dividend yield (x_{it}) and turnover are relative to other stocks in the same calendar year.

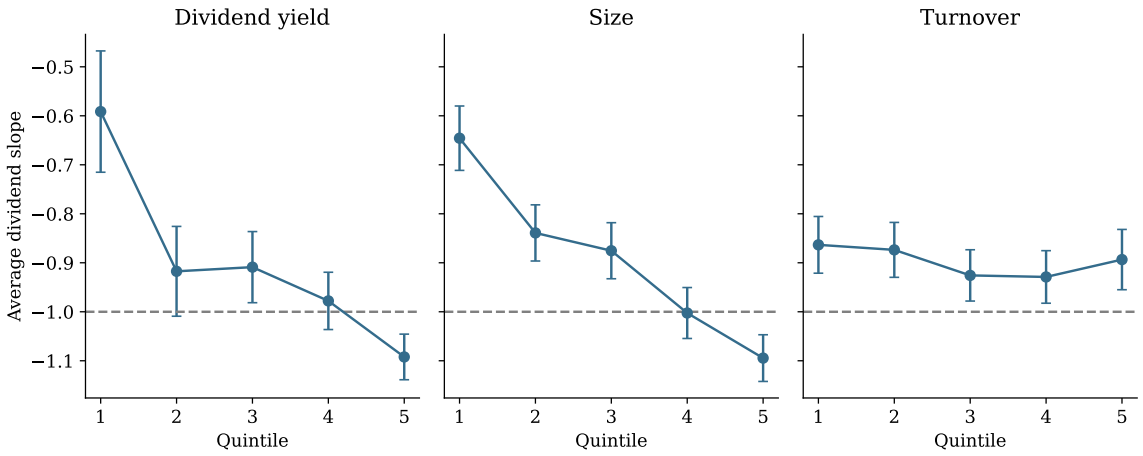


Figure 1: Dividend-yield, size, and turnover profiles. Each panel averages the fitted dividend slope over the other controls in the same joint regression. Bars are 95% ex-date-clustered confidence intervals.

These results are consistent with the findings in the existing literature. The yield and size gradients are both what costly arbitrage predicts: a short-term trader's profit is proportional to the dividend and the cost to the price, so the deviation that survives is larger when the yield is low and when trading costs are high (Kalay, 1982; Shleifer and Vishny, 1997; Pontiff, 2006). Tax clienteles give a second reason for the yield gradient (Elton and Gruber, 1970).

3 The Stock Trading Game

This section describes the Yale Stock Trading Game, a laboratory simulation of a stock market. A key feature of the game is that during the 25-minute trading sessions, there is a public announcement of a change in the fundamental value of a known amount for each of the stocks traded. This allows me to study the price reaction to known-value changes in the fundamental value in a clean setting, where many of the channels that confound the ex-day puzzle are absent by design: there are no taxes, no transaction costs, no clientele effects.

In this clean environment I can construct a close analog to the ex-dividend event, and observe whether prices respond to it correctly. As I discuss below, they do not: across 58 game sessions, and 232 events, prices move, on average, by only about 18 cents for every dollar of change in fundamental value. The lab evidence thus produces an underreaction to a known change in fundamental value that mirrors the ex-day puzzle in direction and is, if anything, larger in magnitude.

3.1 Description of the game

The data analyzed in this section are drawn from sessions of a stock trading game played by undergraduate and graduate students, alumni, and executive-education participants at Yale. The game was originally designed, and is currently maintained, by Roger Ibbotson, who kindly shared

the historical data with me. It is a rich classroom simulation of a stock market environment, designed to give students a feel for how a market with public and private information, an order book, leverage, and short-selling actually functions. Although its primary purpose is pedagogical, the design generates a dataset that is unusually well-suited to studying how prices respond to a clean, publicly announced change in fundamental value.

Each session of the game involves (with a few exceptions) at least 20 players (on average 42) and lasts 25 minutes. Four stocks are traded simultaneously, identified by colours: Green (G), Orange (Q), Red (R), and Blue (B). At the start of each session, ten playing cards are secretly drawn from standard decks for each of the four stocks. Each card has a numerical score in the range -20 (black ace) to $+40$ (red ace).[3] The sum of these ten scores, truncated below at zero and expressed in dollars, defines the stock's *initial fundamental value* (IFV). Players *do not* observe the IFV. Because the unconditional mean per card is \$4, the unconditional expected fundamental value of each stock (before truncation) is \$40.

Each player begins the session with five shares of each of the four stocks plus \$200 in cash. During the 25-minute trading session, players can submit limit or market orders against a publicly observable order book. Prices are quoted in whole dollars.[4] There is no monetary reward; players are ranked publicly by final wealth, and that is the sole incentive. The game has no dividend taxes and no transaction commissions. It does have an order book, financing and short-sale limits, and a cost of private information.

Three kinds of information arrive during a session. First, every stock is the subject of two public *Security Analysis Reports* (SARs), one at the start of the game and one at minute 12.5. A SAR draws five of the ten cards at random, sums them, adds \$20 (the expected value of the other five), truncates at zero, and displays the result. It is a noisy public signal about the fundamental value. Second, once per stock, at a fixed time (minute 5 for G, 10 for Q, 15 for R, and 20 for B), an *EPS event* replaces one of the ten cards with a fresh draw and publicly displays the change in card value: if a -10 card is replaced by a $+6$ card, everyone sees "+16." The underlying card sum changes by exactly the displayed amount. This is the game's analogue of an ex-dividend distribution: public, scheduled, and exact in dollars, though unlike a dividend it can be positive or negative. Third, players can purchase *peeks*. A peek costs \$1 and shows three of the ten cards, drawn with replacement, to the buyer alone (so they are a private signal). Peeks can be repeated without limit, and the data record which cards each player saw and when.

The sample I analyze consists of 58 sessions of the game, played between 2020 and 2024. With four stocks per session, the 58 sessions yield $58 \times 4 = 232$ EPS events in total.

[3] All cards from 2 to 10 are worth their face value, Jack, Queen, and King are worth 10, and Aces are worth 20. The score is multiplied by -1 for black cards and by $+2$ for red cards. The unconditional expected score per card is therefore \$4.

[4] Every executed price in the data is an integer, so one dollar is the effective price grid. Appendix B.1 shows that rounding to this grid cannot account for the response to large announcements. The player guide is at https://stg.som.yale.edu/how-to-play.

3.2 Data

For each game, the dataset contains the ten final cards that determine the truncated fundamental value at the end of the session, the timestamp of every order submitted to the book, the price and size of every transaction that clears, and a complete log of every public signal (SAR and EPS event) and every private peek with the requesting player's identifier. All data are linked by stock id (G, Q, R, B) and by anonymized player ids that are unique within (but not across) games. For the analyses that follow, I use only completed-transaction prices; player-level peek behavior is summarized at the stock-game level (the total number of peeks generated against each stock before each EPS event), and I do not exploit the player-level structure of the data otherwise, leaving this to future research.

Table 2 reports summary statistics at four levels of aggregation. Panel A summarizes the universe of completed transactions: the average transaction price is \$40, exactly the unconditional non-truncated expected stock value, and the substantial dispersion in prices reflects both the wide range of realized fundamental values and the within-session price drift around news arrivals. Order sizes are small: more than 90% of transactions are for a single share. Panel B reports stock-game-level statistics ($N = 232$): the average initial fundamental value is \$43.8, somewhat above \$40 because of the truncation at zero. The two SARs and the final fundamental value all have means in the same range. The EPS value change averages near zero and ranges from -50 to $+60$ in my sample. The average stock-game has 174 trades and 127 peeks, with substantial dispersion. Panel C aggregates to the game level ($N = 58$): on average about 42 players participate per session, and the equal-weighted market return (computed as the return to holding the starting portfolio of each player, valued at the ex-ante price of \$1000) across the four stocks averages about 5%. Panel D reports player-level statistics across 2,420 player-game observations.

3.3 The benchmark price reaction

The focus of my analysis of the Stock Trading Game is the price reaction to the "EPS event". I therefore start with a discussion of the theoretically correct reaction to the EPS event, and then turn to the data. Mechanically, the EPS event swaps one of the ten cards that determine the stock's fundamental value with a new draw, and *publicly displays the change in card value*. Absent any truncation, this means that the change in the (non-truncated) fundamental value of the stock is exactly equal to the announced EPS value change, by construction. A rational, fully informed player therefore understands that the fundamental value has moved by the announced EPS value change, and so – to the extent that the price reflects the fundamental value – should expect the stock price to move by exactly the EPS value change, irrespective of the prior distribution she might have had (just like an ex-dividend event).

The truncation at zero introduces a small wrinkle. Because both the initial and final fundamental values are truncated below at zero, the EPS event has no actual effect on the (truncated) fundamental value when the sum of card scores is sufficiently negative both before and after the swap. From a player's point of view, that means there is some probability — depending on her belief about how close the non-truncated fundamental value is to zero — that the EPS event has zero effect on the

Table 2: Summary statistics: the Yale Stock Trading Game

Variable	Mean	SD	Min	P10	Median	P90	Max	*N*
Panel A: Completed transactions								
Price	40.0	15.5	1.0	20.0	40.0	59.0	200.0	40,446
Shares per transaction	1.1	0.4	1.0	1.0	1.0	1.0	10.0	40,446
Panel B: Stock–game observations								
Initial payoff	43.8	36.9	0.0	0.0	41.0	94.8	150.0	232
First SAR	42.7	30.4	0.0	1.0	40.5	85.9	124.0	232
Second SAR	39.9	27.6	0.0	3.0	40.0	73.9	108.0	232
EPS value change	-1.5	21.2	-50.0	-28.0	-2.0	28.0	60.0	232
Final payoff	42.4	35.9	0.0	0.0	40.0	86.0	168.0	232
Last price	39.8	18.1	1.0	16.0	40.0	62.0	125.0	232
Transactions	174.3	105.8	6.0	54.1	159.0	336.9	531.0	232
Private peeks	127.2	87.6	0.0	35.1	114.0	265.0	502.0	232
Panel C: Games								
Players	41.7	32.3	5.0	19.7	33.5	66.0	196.0	58
Market return (percent)	4.8	36.0	-66.0	-34.9	-2.0	48.7	109.5	58
Median return (percent)	3.7	37.3	-79.5	-36.6	0.9	47.0	106.3	58
Panel D: Player–game observations								
Return (percent)	0.7	54.6	-298.8	-59.7	-4.7	69.9	346.4	2,420
Excess return (percentage points)	-1.7	41.4	-284.3	-43.0	0.0	40.0	308.4	2,420
Private peeks	12.2	18.5	0.0	0.0	6.0	31.0	218.0	2,420
Transactions	33.4	30.5	0.0	4.0	27.0	66.1	332.0	2,420

Note: Statistics use completed transactions, stock–game observations, sessions, and player–game observations as indicated. Player excess returns subtract the recorded market return for the same game.

actual, truncated, fundamental value. In equilibrium, this should bias the price reaction toward zero, but only by a small amount, since the truncation rarely binds in the games in my sample. Conditional on observing a public signal (a SAR) that reduced the truncation probability sufficiently (that is, a sufficiently positive SAR), the truncation issue should disappear entirely.

Formally, let V_{gs}^0 be the pre-replacement untruncated card sum of stock s in game g and EPS_{gs} the EPS value change. The card sum changes by exactly EPS_{gs}, and the payoff changes by

$$\Delta V_{gs} = \max\{0, V_{gs}^0 + \text{EPS}_{gs}\} - \max\{0, V_{gs}^0\}. \tag{3}$$

To quantify the effect of truncation on the benchmark price reaction, in column (1) of Table 3 I report the average benchmark price reaction in my sample, accounting for the truncation: I regress ΔV_{gs} defined as in equation (3) on EPS_{gs} across all the EPS events in my sample. The slope is 0.883 in the full sample and 0.985 when the first SAR exceeds 40, where a negative card sum is very unlikely. Truncation thus accounts for about a tenth of the EPS value change in the full sample and for almost none of it when the first SAR is high (much less than the shortfall in price reaction I document below). Figure 2 shows the two relations in a graph.

Table 3: Effect of truncation on the benchmark price reaction

Sample	(1) Payoff slope	(2) SE	(3) Events
All	0.883	0.024	232
First SAR > 30	0.967	0.014	145
First SAR > 40	0.985	0.007	116
First SAR > 50	0.996	0.003	90

Note: Column (1) regresses the change in the (unobservable to participants) truncated fundamental value on the announced EPS value change, with standard errors clustered by game in column (2). Column (3) reports the number of stock-games in each sample. High-SAR restrictions reduce the realized effect of truncation.

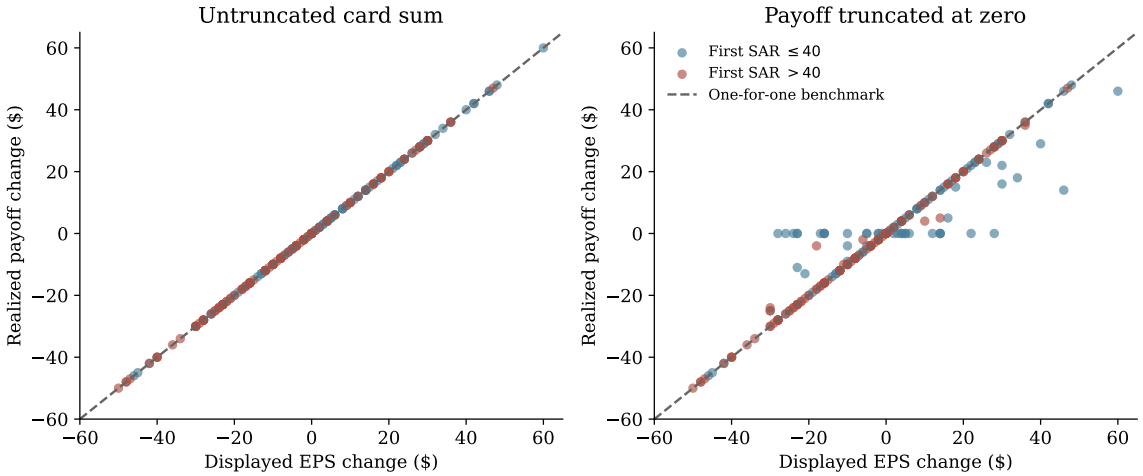


Figure 2: The EPS value change and the realized payoff change. Left: the untruncated card-sum change, an identity. Right: the change after truncating the card sums at zero. Colours distinguish whether the first SAR exceeds 40.

Finally, note that this benchmark treats the EPS announcement as a shift in the payoff and assumes away any inference participants might draw from it about the value of the card removed. I regard such inference as infeasible for participants in real time, given the fast pace of the game and the complexity of the calculations required (see some examples in Appendix B.3). In fact, to the extent participants understand that they should be doing these calculations, this would represent exactly the kind of cognitive effort that could give rise to the bias that the behavioral model (section 4) predicts. Appendix B.3 nevertheless reports, by simulation, what a trader who learned optimally from the announcement, drawing that full inference, would do. The observed price responses documented below once again fall far short of the full-learning benchmark that incorporates this inference, so nothing below depends on this assumption.

3.4 Price reaction to the public change in value

To study price reaction to the EPS event, I estimate the regression

$$\Delta P_{gs}^{(k)} = \alpha + \beta \, \mathrm{EPS}_{gs} + \varepsilon_{gs}, \tag{4}$$

where $\Delta P^{(k)}$ is the difference between the price at the kth transaction timestamp after the announcement and the price at the last timestamp before it, with standard errors clustered by game. Panel A of Table 4 shows the results for different values of k.

The average price reaction to the EPS event is $\hat{\beta} = 0.177$ (SE 0.033) at the first transaction after the event. The slope rises as the window widens, to 0.278 at the tenth timestamp and to 0.259 thirty seconds after the announcement, and windows that average two or five prices on each side give similar results (last two columns). Every estimate rejects the one-for-one benchmark at any

conventional level. Restricting to events whose first SAR exceeds 30, 40, or 50, where truncation is immaterial, gives similar results (Panel B).

Table 4: Price adjustment across transaction windows and public-signal restrictions

	1 trade after	2 trades after	5 trades after	10 trades after	30 seconds after	2 before/after means	5 before/after means
Panel A: Post-event windows							
Constant	-0.529	-1.565^{**}	-2.373^{***}	-1.766^{**}	-1.265^{*}	-0.951^{*}	-1.577^{***}
	(0.591)	(0.740)	(0.783)	(0.788)	(0.729)	(0.540)	(0.451)
EPS value change	0.177^{***}	0.208^{***}	0.253^{***}	0.278^{***}	0.259^{***}	0.179^{***}	0.217^{***}
	(0.033)	(0.044)	(0.054)	(0.048)	(0.056)	(0.033)	(0.032)
Zero EPS	12	12	12	11	12	12	12
Zero price change	56	46	20	12	50	13	3
p: slope = 1	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
R^2	0.153	0.137	0.168	0.186	0.184	0.171	0.280
Events	231	231	229	223	232	231	226
Games	58	58	58	57	58	58	58
Panel B: First-SAR restrictions, first post-event price							
Sample	All	SAR > 30	SAR > 40	SAR > 50			
Constant	-0.529	0.283	0.228	0.073			
	(0.591)	(0.740)	(0.718)	(0.993)			
EPS value change	0.177^{***}	0.192^{***}	0.202^{***}	0.196^{***}			
	(0.033)	(0.042)	(0.049)	(0.053)			
p: slope = 1	< 0.001	< 0.001	< 0.001	< 0.001			
R^2	0.153	0.174	0.176	0.176			
Events	231	145	116	90			
Games	58	57	55	50			

Note: Panel A uses the last pre-event price and the indicated post-event transaction timestamp, except the 30-second column and the final two columns. The 30-second column uses the last transaction price at or before 30 seconds after the announcement. The final two columns average the stated number of timestamps on each side. Panel B conditions on the first SAR. All columns include an intercept. Standard errors in parentheses are clustered by game. ***, **, and * denote two-sided significance against zero at the 1%, 5%, and 10% levels. The row p: slope = 1 reports the game-clustered p-value of the test that the announcement coefficient equals the one-for-one benchmark.

A striking feature of the data is how often the first price after the announcement is exactly the price before it: 56 of 231 events (24.2%), against only 12 announcements of zero change in value. Figure 3 shows the event-level price changes against the announced value change. The fitted response is far flatter than the one-for-one line across the whole range of announcements, positive and negative.

3.5 Additional analysis of the price reaction

Unlike a dividend, an announcement can raise or lower the payoff. Columns (1) and (2) of Table 5 split the sample by sign: the slope is 0.121 for negative and 0.200 for positive announcements, both far below one, and the difference between them is imprecise ($p = 0.394$). The shortfall is not a feature of one direction of the change.

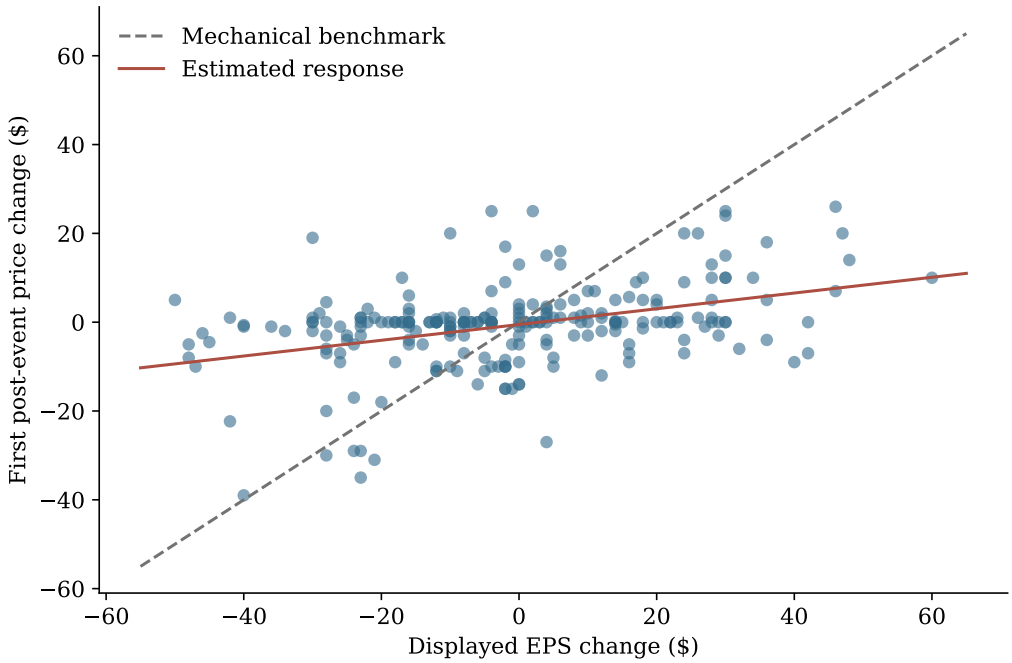


Figure 3: First transaction-price adjustment to the EPS event. The fitted line is the first column of Table 4; the dashed line is the one-for-one benchmark.

Table 5: Responses by announcement sign and for newly initiated orders

	(1) Negative EPS	(2) Positive EPS	(3) Newly initiated orders
Constant	-1.588	-0.672	-1.219
	(1.456)	(1.589)	(0.820)
EPS value change	0.121	0.200***	0.160***
	(0.073)	(0.073)	(0.046)
p: slope = 1	< 0.001	< 0.001	< 0.001
R^2	0.024	0.092	0.088
Events	123	96	229
Games	54	47	58

Note: Columns (1) and (2) split the non-zero announcements by sign. The p-value for equality of the two slopes reported in the text comes from a pooled regression of these events on the EPS value change, a positive-announcement indicator, and their interaction. Column (3) uses all announcements and replaces the first post-event price with the earliest price of transactions whose orders were both submitted and completed strictly after the announcement, so no record initiated before the announcement enters that price; the sample is the stock-games with at least one such transaction. Standard errors in parentheses are clustered by game. ***, **, and * denote two-sided significance against zero at the 1%, 5%, and 10% levels. The row p: slope = 1 reports the game-clustered p-value of the test that the announcement coefficient equals the one-for-one benchmark.

One Potential concern with the analysis above is that the price reaction may be due to orders that were submitted before the announcements, and new orders then simply being executed against orders that were already in the book. To test this, in column (3) I perform the same regression after discarding every transaction record initiated before the announcement, and forming the first post-event price only from orders submitted after it. The slope is 0.160 (SE 0.046), a similar level of underreaction. Adding game intercepts, or restricting to sessions with at least twenty players (Appendix Table B.1), leaves the slope between 0.17 and 0.20.

Finally, I study how the order book (rather than the trades) responds to the announcement. I rebuild the book from every order record and track the best bid, the best ask, and their midpoint from one second before the announcement to 60 seconds after it (Table 6). On the 165 events with a two-sided book at every horizon, the midpoint moves by 0.081 per dollar after fifteen seconds and 0.221 (SE 0.038) after 60 seconds, with the bid at 0.207 and the ask at 0.236. Quotes move less than transaction prices at every horizon, and both sides of the book move alike. Trades are therefore not executing against stale limit orders while the book has already adjusted: the standing quotes underreact at least as much as the trades do.

Table 6: Order-book quotes and transaction prices after the announcement

Seconds	Midpoint	Best bid	Best ask	Transaction price	Events
15	0.081 (0.020)	0.103 (0.030)	0.060 (0.021)	0.209 (0.049)	165
30	0.117 (0.022)	0.108 (0.022)	0.127 (0.031)	0.228 (0.060)	165
60	0.221 (0.038)	0.207 (0.037)	0.236 (0.045)	0.354 (0.052)	165

Note: Slopes on the EPS value change of the change in the order-book midpoint, best bid, and best ask from one second before the announcement to the indicated horizon, with game-clustered standard errors in parentheses. The book is rebuilt from every order record: an order is live from its initiation time until its recorded finish time, and orders still open at the end of the game are treated as live throughout. The sample is fixed across horizons at the 165 events in 55 games with a two-sided book at every horizon; 203 of 232 events have a two-sided book one second before the announcement. The transaction-price column uses the last transaction price at or before each horizon, on the same events.

3.6 Standard explanations for the underreaction

Several standard explanations for the ex-dividend puzzle do not apply here, as discussed above (e.g., taxes, trading fees). Other explanations are still possible. One of them, behavioral in nature but different from the one I propose in this paper, is that participants simply do not notice the announcement. That is quite unlikely in this context, because the announcement is scheduled, displayed to every participant, and one of only three public events for the stock.

That explanation can also be directly ruled out by studying trading activity around the announcement. For each of the 232 events I compare order submissions and market participation (number of active traders) in the two minutes before and after the announcement, and compare it to the participation in the other three stocks in the game at the same moment (which did not get an EPS event at that time).

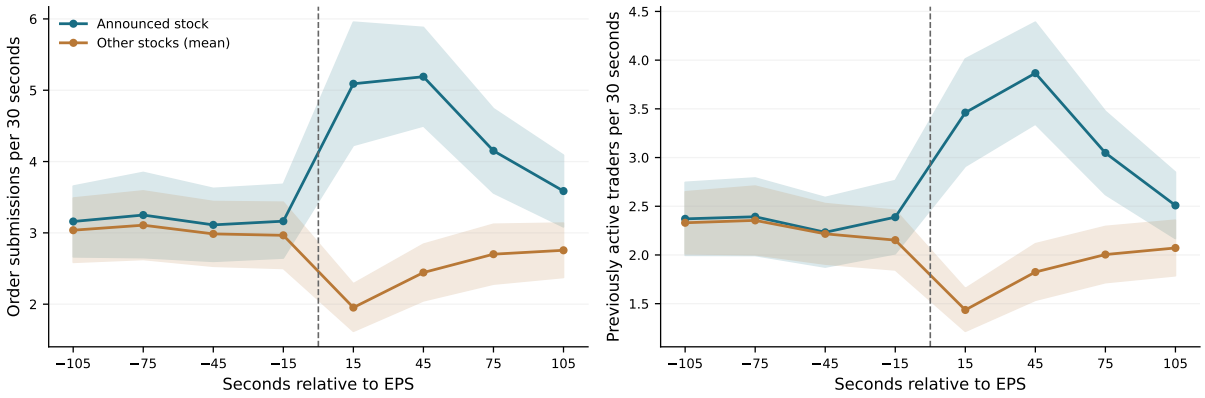


Figure 4: Trading activity around the EPS announcement. Each point is the mean across the 232 events of the number of order submissions (left) or of distinct previously active traders submitting an order (right) in the announced stock during a 30-second bin, where previously active traders are those who had submitted an order in the game more than two minutes before the announcement. The orange line is the same quantity averaged over the three other stocks in the same game. Shading denotes pointwise 95% game-clustered intervals.

Figure 4 reports activity in 30-second bins from two minutes before to two minutes after the EPS event. The left panel counts order submissions in the stock, and the right panel counts the number of distinct participants submitting them, restricted to participants who had already been active in the game more than two minutes before the announcement so that the count is not inflated by players drawn in by the announcement itself. Both are means across the 232 events of per-event counts in each 30-second bin. In both panels, the blue line refers to the announced stock and the orange line to the average of the other three stocks in the same game. The figure shows that clearly participants see the announcement and act on it within seconds. So lack of notice is not the explanation for the underreaction.

Microstructure explanations – which have been used to explain the ex-dividend puzzle in the stock market – could potentially explain the underreaction in the lab experiment as well. One of them is the rounding effect due to tick size. In this case, the announced change is always an exact multiple of the tick size ($1), so there is no rounding at all. In addition, Appendix B.1 shows that the same underreaction is present also for announcements that are large in dollar terms: for announcements of $20 or more the first-price slope is 0.194. The other main microstructure explanation is related to the bid-ask bounce, which is ruled out by the order book analysis above.

Yet another possibility is that financing and/or position limits are the reason for the underreaction: that is, perhaps participants are unable to trade enough shares to fully react to the announcement. To study that, in Appendix B.2 I reconstruct every player's cash and holdings from the order flow, and show that the underreaction is the same in events where at least five active players held enough cash or shares to trade at least five shares against the announcement (0.179 at the first trade). These results are robust to the thresholds of liquidity required.

3.7 Summary of the Lab Results

The Yale Stock Trading Game provides a clean public announcement whose value is objective and known to all participants, and it is in many ways similar to the ex-dividend case in the stock market. Because several of the key features of the market that have been used to explain the ex-dividend puzzle do not apply here, or can be directly ruled out, the experiment can help exclude some specific theories for the ex-dividend puzzle (to the extent, of course, that the results of the experiment can be applied to the real stock market setting), and can point towards explanations for the underreaction in the stock market.

Summing up, the Stock Trading Game analysis shows that underreaction to a known change in fundamental value in the lab experiment is not due to 1. taxes; 2. transaction costs; 3. lack of attention; 4. discrete price grid; 5. financing and/or position limits. In the next section, I propose an explanation for the underreaction based on uncertainty about the fundamental value of the stock, and I explore its predictions both in the game and in the stock market on ex-dividend days.

4 Cognitive Uncertainty and Underreaction to a Known Change in Value

The previous section shows that prices in the Stock Trading Game underreact to a *known* change in value, and that taxes, trading costs, microstructure concerns, and inattention cannot explain it. This section proposes a mechanism to explain the underreaction built on the cognitive uncertainty framework of Enke and Graeber (2023). In their model, a decision maker who can form only a noisy estimate of her best decision does not act on it in full; she instead starts from a *cognitive default* – the decision she would make without deliberating – and adjusts only part of the way toward her actual estimate; the noisier her estimate, the more she anchors to the cognitive default.

Applying this general framework to the specific setting of the Stock Trading Game requires specifying the decision and the cognitive default. I model the decision of a participant as whether to trade a stock, and at what price. The cognitive default is her willingness to pay before any deliberation about value; a natural way to model it is the latest price at which the stock is trading: not because she takes it to be the true value of the stock, but because trading at the prevailing price is a natural choice for someone who has no view of her own (which is what the cognitive default is meant to capture). Consider then a participant who is uncertain about a stock's value, and must piece together the public signals, her private signals, and the announced change in value (the "EPS event") into an estimate of what the stock is worth. She understands that the known change shifts the fundamental, and her estimate, by the full amount. But the cognitive default (the visible pre-announcement price) has not moved. According to the model, then, her offer to buy or sell the stock reflects a weighted average of her best estimate and the pre-announcement price, with more weight on the price the noisier her estimate. Averaged across participants, this behavior yields a post-announcement price that is biased toward the pre-announcement price, generating observable underreaction to the change in fundamental value.

The remainder of this section briefly presents the formal derivations of this prediction. Consider a stock with unknown payoff V, and call I the investor's information. Write $v = \mathbb{E}[V \mid I]$ for the valuation an investor holds and $U = \mathrm{Var}(V \mid I)$ for her uncertainty about the payoff, which I call fundamental uncertainty. An event changes the payoff from V to $V + \delta$, where δ is public and known exactly and carries no information about V.

Rational updating. If all investors are rational, their valuation moves by:

$$\mathbb{E}[V + \delta \mid I] - \mathbb{E}[V \mid I] = \delta \tag{5}$$

independently of the level of uncertainty U. If rational updating describes the updating of all traders that participate in the market, then the equilibrium price should simply move by δ.

Cognitive uncertainty. Next, I apply the framework of Enke and Graeber (2023), modeling the cognitive default as the current (pre-announcement) price as described above (I refer to the default as d). At the time the change in value for a stock is announced, market participants focus on that stock and try to determine whether they should trade it and at what price. To do so, they consider the public information, the private signals, and, of course, the newly announced δ. Call the participant's own estimate of the fundamental value of the stock $\tilde{v}$:

$$\tilde{v} = (v + \delta) + \eta$$

where v is the pre-announcement fundamental value and η is the noise in her estimate, which has mean zero and positive variance. Note that the participant understands that the fundamental value has shifted by δ; it is just that her own estimate is noisy, and she is aware of the noise.

Applying equation (1) of Enke and Graeber (2023) with the pre-announcement price as the default, her post-announcement valuation of the stock is a weighted average of her estimate and the default:

$$p^* = (1 - \lambda)\, d + \lambda\, \tilde{v}, \tag{6}$$

where λ increases with the precision of her estimate. An investor with a precise estimate acts on nearly all of it; an investor with more uncertainty stays closer to the default. Averaged across investors, whose errors are independent and wash out, the aggregate valuation is $\bar{p}^* = (1 - \lambda)d + \lambda\,(v + \delta)$.

Just before the announcement the same rule gives $(1 - \lambda)d + \lambda\, v$: the investor's estimate was centered on v, and anchored to the last observable price. The announcement moves every estimate by δ but leaves the default where it is, so the aggregate valuation changes by

$$\Delta\bar{p}^* = \lambda\, \delta. \tag{7}$$

Note that no one fails to add δ to their estimate in this model. What determines the partial price adjustment is how much of their own estimates the investors act on.

An additional prediction. In the model of Enke and Graeber (2023), observed underreaction to the change in value (determined by λ) depends on the precision of the investors' estimate of fundamental value. While that is not the same as the amount of uncertainty about the fundamental value U, the two are related, because, everything else equal, higher uncertainty about the fundamental value will lead to higher variance in someone's estimates of it. Under this assumption, I can then derive an additional testable prediction of this model: that the displayed underreaction should be more severe when fundamental uncertainty is higher. This is a prediction that the next two subsections explore, in the Stock Trading Game and in the stock market for the ex-dividend case, where much larger samples allow sharper tests.

4.1 Uncertainty and underreaction in the Stock Trading Game

To test this prediction in the Stock Trading Game, I look for proxies for participants' uncertainty about the valuation, and test whether events in which participants are more uncertain display more underreaction.

Even with a small number of events (232 across all games), the evidence from the Stock Trading Game shows that uncertainty about the fundamental value is associated with underreaction. I propose two pieces of evidence. The first exploits the timing of the announcements of stocks within the game; the second exploits the additional information I observe about the participants' private information purchases ("peeks") before the announcement ("EPS event").

Timing across stocks. The first piece of evidence exploits the differential timing of the announcement across stocks. In particular, G stocks always have the "EPS event" early in the game, exactly 5 minutes in. By this time, participants will have accumulated little information about the value of the stock, so one would expect the response to be weakest for G. The other stocks then have their own "EPS events" spaced out every 5 minutes after that, and so one would expect the response to become stronger for those. Table 7 computes the slope separately for the four types of stocks, pooling across the 58 games. The table shows that indeed the slope for stock G is by far the lowest, at 0.075, compared to 0.31, 0.16, and 0.20 for the other three stocks, whose EPS event arrives later in the game. The difference between the slope for G and the pooled slope of the other three stocks is 0.154, statistically significant at the 1% level. While one would expect the response to be monotonic with time, in the data the slope of the reaction for the later stocks fluctuates non-monotonically, though it stays consistently above the G slope (and the differences between the slopes of the other three stocks are not statistically significant).

Private information. In the Stock Trading Game, participants can purchase private information about the stock's value at any time ("peeks"). Peeks measure private information acquisition directly; presumably, events before which participants have purchased more peeks should be associated with less underreaction (because participants have better information about the stock's value). To test whether this is the case, I estimate a pooled regression of the price reaction for stock s in game g on

Table 7: Underreaction across the four stocks

Stock	Minute	Slope	SE
G	5	0.075	0.046
Q	10	0.309	0.058
R	15	0.163	0.094
B	20	0.204	0.046

Note: Separate regressions of first-post-event price changes on EPS. The pooled non-G minus G slope difference is 0.154 (SE 0.057, clustered by game).

the corresponding EPS value change, interacting the slope of that reaction with z_{gs}, the log of one plus the cumulative number of peeks purchased on that stock in that game before the announcement (standardized to mean zero and unit variance across events). Specifically, I estimate:

$$\Delta P_{gs} = \alpha_g + \mu_s + \beta \,\mathrm{EPS}_{gs} + \gamma z_{gs} + \theta \,\mathrm{EPS}_{gs} z_{gs} + \varepsilon_{gs}, \tag{8}$$

where ΔP_{gs} is the change from the last pre-announcement transaction price to the price after announcement, and EPS_{gs} is the EPS value change. The regression also includes game fixed effects and stock fixed effects. Standard errors are clustered by game.

Table 8 reports the estimates on the 223 events in 57 games for which all four horizons are observed. The interaction term is essentially zero when computing the change from the first trade, but is positive after 5 trades, and positive and statistically significant after 10 trades and 30 seconds after the announcement. The magnitudes are economically significant. For example, after 10 trades, a one-standard-deviation increase in log peek purchases is associated with an additional \$0.091 response per announced dollar (about a third of the average response). As predicted by the model, more precision of the estimate of the stock's value is associated with a stronger response to the announcement. Appendix Table B.5 reports all the coefficients of the regression (for the specification that uses the tenth trade after announcement).

These results show that despite the limited amount of data in the 232 events available in the Stock Trading Game, there is significant evidence that indeed underreaction is higher when participants have more uncertainty about the stock's value, in line with the predictions of the model.

4.2 Uncertainty and underreaction on ex-dividend days

Next, I explore the prediction of the model that the extent of underreaction of prices on ex-day is related to fundamental uncertainty just like in the Stock Trading Game. Compared with the Stock Trading Game, the ex-day data provide a much larger sample and richer proxies for uncertainty about the fundamental value. In my analysis, I consider two alternative measures of uncertainty: past return volatility and the dispersion of analysts' long-term growth forecasts. The key hypothesis is that ex-day adjustment is weaker when investors have a harder time pinning down the fundamental value.

Table 8: Prior private information and adjustment to the EPS event

	1 trade after	5 trades after	10 trades after	30 seconds after
EPS ($\hat{\beta}$)	0.198***	0.286***	0.281***	0.282***
	(0.049)	(0.068)	(0.065)	(0.070)
EPS × prior information ($\hat{\theta}$)	−0.030	0.072	0.091**	0.124**
	(0.038)	(0.064)	(0.043)	(0.054)
p-value, interaction	0.438	0.265	0.040	0.027
Events	223	223	223	223
Games	57	57	57	57
R^2	0.399	0.422	0.427	0.415

Note: The table reports $\hat{\beta}$ and $\hat{\theta}$ in equation (8). Prior information is $\log(1 + \text{peeks})$, counting purchases on the announced stock strictly before its EPS event and standardized over all 232 events. Price changes are relative to the last transaction price before the announcement. Game-clustered standard errors appear in parentheses. ***, **, and * denote two-sided significance at 1%, 5%, and 10%.

Past realized volatility is computed as the mean squared daily return over the 252 trading days before the event. *Long-term-growth (LTG) dispersion* is the cross-sectional dispersion of long-term growth forecasts across the analysts covering the firm: the cross-sectional standard deviation of their I/B/E/S forecasts, in growth-rate points. It measures how much analysts disagree about the firm's long-term growth. I use the latest snapshot strictly before the ex-date and at most 62 calendar days old, include zero dispersion when it occurs, but require a minimum number of analysts. Note that when studying LTG dispersion, I always control flexibly (with dummy variables) for the number of analysts in the sample, so that the dispersion is not mechanically due to a different number of analysts covering the firms.

The two measures are both meant to capture the uncertainty that investors have on the ex-day about the fundamental value of the firm. To study how these measures are related to the extent of underreaction of prices on ex-day, I extend the regression of Section 2 (equation 2) in which the slope of the price reaction on ex-day is allowed to depend on quintiles of size, turnover, and dividend yield (Figure 1 and Appendix Table C.1). I either add to this regression interaction of the dividend paid with five quintiles of past realized volatility, or the interaction with five quintiles of LTG dispersion (as well as interactions with the number of analysts as control).[5]

Figure 5 plots the average estimated ex-day price drop across the LTG dispersion and volatility quintiles. For LTG dispersion, I plot the results conditioning on having at least 3 and at least 8 analysts, respectively.

LTG dispersion. The left graph of the figure shows that, conditional on how the ex-day slope varies with size, turnover, dividend yield, and analyst count, this slope significantly varies with fundamental uncertainty as measured by the LTG dispersion (while standard errors are large, the

[5]Quintiles for both uncertainty measures are computed within year; I do not jointly include both measures in the same regression, as they are highly correlated in the panel and I consider them alternative proxies for uncertainty.

differences across uncertainty quintiles are statistically significant, as discussed next). The ex-day slope is 0.121 less negative in the top quintile of LTG dispersion than in the bottom quintile at the three-analyst threshold ($t = 2.46$), and 0.187 at the eight-analyst threshold ($t = 3.24$). Of course, raising the threshold changes the precision with which dispersion is estimated but also changes the population of firms: this naturally explains the different level across thresholds. That said, the analysis shows a meaningful relation between the extent of underreaction of prices on ex-day and the uncertainty that investors have as measured by the LTG dispersion.

Volatility. Even stronger results hold for the case of past volatility. As the right graph of Figure 5 shows, there is a clear relation between volatility and the ex-day dividend drop, with the smallest drop in absolute value concentrated among the most volatile stocks.

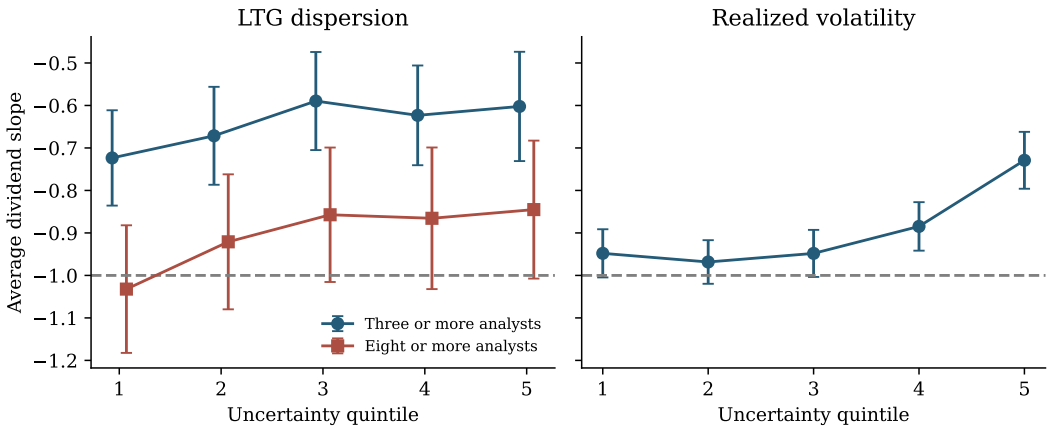


Figure 5: Average dividend slopes by uncertainty quintile. Each series averages over the other controls in its own regression sample. The left panel shows LTG dispersion at the three-analyst (circles) and eight-analyst (squares) thresholds; the right panel shows past volatility. Bars are 95% ex-date-clustered intervals from the full covariance matrix; the dashed line marks the frictionless slope of −1.

4.2.1 Interpretation of the results and limits to arbitrage

These results show that, holding size, turnover, and dividend yield constant, as well as analyst coverage for the LTG dispersion measure, the ex-day discount is larger for stocks whose value is more uncertain, whether uncertainty is measured by LTG dispersion or by realized volatility. This is what the model of cognitive uncertainty predicts: investors who hold only a loose sense of what the stock is worth stay closer to the price before ex-day, and underreact more.

If these patterns are, at least in part, due to behavioral forces, one would expect rational arbitrageurs to try and trade against them and, when possible, arbitrage them away. For this pattern to survive in equilibrium, it must then be sufficiently costly for arbitrageurs to not be worth or possible to correct it fully (Shleifer and Vishny, 1997; Pontiff, 2006). To study this, I consider two standard proxies for the cost of arbitrage, size and turnover, and ask whether the uncertainty gradient

is weaker where arbitrage is cheaper (note that I focus this analysis on the volatility patterns, becuase LTG dispersion can be computed only for a group of large, liquid stocks that receive sufficient analyst coverage). Table 9 splits the sample into four groups of stocks: small or large firms (the two bottom or the two top size quintiles of the main specification), interacted with low or high turnover (the two bottom or the two top size quintiles of the main specification). Within each group I sort stocks into within-year quintiles of past volatility and estimate the ex-day slope of the stocks by uncertainty quintile, and report the average slope in the lowest-uncertainty quintile (column 1) and in the highest-uncertainty quintile (column 2), together with their difference (column 3). Two results are notable. First, within every group, the prices of high-uncertainty stocks react substantially less to the dividend than those of low-uncertainty stocks. Second, underreaction is indeed more severe in small firms with low turnover, and within that group, especially so for stocks with high uncertainty.[6] Overall, the results in this section are consistent with a behavioral explanation of the uncertainty-underreaction gradient (as measured by volatility and LTG dispersion), limited to some extent by the action of rational arbitrageurs.

Table 9: Uncertainty gradient within size and turnover groups

	(1) Low uncertainty	(2) High uncertainty	(3) Difference (2)−(1)
Small firms, low turnover	-0.793	-0.444	0.349
	(0.032)	(0.056)	(0.065)
Small firms, high turnover	-1.044	-0.729	0.315
	(0.043)	(0.102)	(0.111)
Large firms, low turnover	-1.168	-0.916	0.252
	(0.023)	(0.050)	(0.055)
Large firms, high turnover	-1.114	-0.846	0.268
	(0.022)	(0.050)	(0.055)

Note: Each row reports, for the indicated group of stocks, the ex-day dividend slope of the stocks in the lowest (column 1) and highest (column 2) within-year quintile of 252-day realized volatility, from a separate regression for each group of the market-adjusted ex-day return on the dividend yield interacted with the volatility-quintile indicators. Volatility quintiles are formed within the group. Column (3) is the difference between the two slopes. Standard errors, clustered by ex-date, are in parentheses. Number of events in each group: 80,155; 39,358; 42,669; 82,149.

5 Conclusion

In financial markets, the fundamental values of assets are generally not observable. In specific cases, however, changes in the underlying value are known and observed exactly. A dividend removes a known amount of cash from the firm; analogously, an EPS event in the Yale Stock Trading Game

[6]For large stocks, the underreaction more muted, yet there is still some volatility gradient, suggesting that other costs of arbitrage – among which volatility itself, that may pose risks to arbitrageurs as modeled in Michaely and Vila (1995) – are playing a role in determining the equilibrium underreaction.

adds or removes a known amount from a stock's payoff. Empirically, in both settings, prices move by much less than the change in underlying value. In the game, prices move by 0.177 per dollar of EPS value change at the first trade and by 0.278 after ten trades, with no taxes and no commissions, and even though participants notice the event and trade on it. On ex-dividend days, the price falls by about three quarters of the dividend.

Whereas several explanations have been proposed (with mixed empirical support) for the ex-dividend day puzzle, none of these can explain the pattern in the Stock Trading Game. I propose an alternative behavioral explanation that can explain underreaction in both settings, based on cognitive uncertainty: investors who are unsure what a stock is worth stay close to its pre-event price and adjust only part of the way toward the new value. Consistent with this explanation, I show that the underreaction is larger where fundamental value is more uncertain, and this pattern is present in both settings.

Of course, the fact that a model based on cognitive uncertainty can fit qualitatively the patterns does not exclude other explanations for the ex-dividend day puzzle, and is likely to be just one of the many drivers of the underreaction. The Stock Trading Game, with its rich structure and data (including on individual traders' information acquisition behavior) offers the opportunity to further explore this question, and to identify the specific mechanisms that drive the underreaction in different settings.

References

Anderson, E. W., Ghysels, E., and Juergens, J. L. (2005). Do heterogeneous beliefs matter for asset pricing? *Review of Financial Studies*, 18(3):875–924. Uses cross-analyst LTG-forecast dispersion (in addition to FY1 EPS dispersion).

Augenblick, N., Lazarus, E., and Thaler, M. (2025). Overinference from weak signals and underinference from strong signals. *Quarterly Journal of Economics*, 140(1):335–401.

Bali, R. and Hite, G. L. (1998). Ex-dividend day stock price behavior: Discreteness or tax-induced clienteles? *Journal of Financial Economics*, 47(2):127–159.

Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343.

Barron, O. E., Kim, O., Lim, S. C., and Stevens, D. E. (1998). Using analysts' forecasts to measure properties of analysts' information environment. *The Accounting Review*, 73(4):421–433.

Bell, L. and Jenkinson, T. (2002). New evidence of the impact of dividend taxation and on the identity of the marginal investor. *Journal of Finance*, 57(3):1321–1346.

Boyd, J. H. and Jagannathan, R. (1994). Ex-dividend price behavior of common stocks. *Review of Financial Studies*, 7(4):711–741.

Brown, S. J. and Warner, J. B. (1985). Using daily stock returns: The case of event studies. *Journal of Financial Economics*, 14(1):3–31.

Campbell, J. A. and Beranek, W. (1955). Stock price behavior on ex-dividend dates. *Journal of Finance*, 10(4):425–429.

Chetty, R., Rosenberg, J., and Saez, E. (2005). The effects of taxes on market responses to dividend announcements and payments: What can we learn from the 2003 dividend tax cut? Working Paper 11452, National Bureau of Economic Research.

Daniel, K., Hirshleifer, D., and Subrahmanyam, A. (1998). Investor psychology and security market under- and overreactions. *Journal of Finance*, 53(6):1839–1885.

DellaVigna, S. and Pollet, J. M. (2009). Investor inattention and Friday earnings announcements. *Journal of Finance*, 64(2):709–749.

Diether, K. B., Malloy, C. J., and Scherbina, A. (2002). Differences of opinion and the cross section of stock returns. *Journal of Finance*, 57(5):2113–2141. Uses cross-analyst SD of next-fiscal-year (FY1) EPS forecasts as the disagreement measure.

Durand, D. and May, A. M. (1960). The ex-dividend behavior of American Telephone and Telegraph stocks. *Journal of Finance*, 15(1):19–31.

Eades, K. M., Hess, P. J., and Kim, E. H. (1984). On interpreting security returns during the ex-dividend period. *Journal of Financial Economics*, 13(1):3–34.

Elton, E. J. and Gruber, M. J. (1970). Marginal stockholder tax rates and the clientele effect. *Review of Economics and Statistics*, 52(1):68–74.

Elton, E. J., Gruber, M. J., and Blake, C. R. (2005). Marginal stockholder tax effects and ex-dividend-day price behavior: Evidence from taxable versus non-taxable closed-end funds. *Review of Economics and Statistics*, 87(3):579–586.

Enke, B. and Graeber, T. (2023). Cognitive uncertainty. *Quarterly Journal of Economics*, 138(4):2021–2067.

Fama, E. F. (1991). Efficient capital markets: II. *Journal of Finance*, 46(5):1575–1617.

Frank, M. and Jagannathan, R. (1998). Why do stock prices drop by less than the value of the dividend? evidence from a country without taxes. *Journal of Financial Economics*, 47(2):161–188.

Graham, J. R., Michaely, R., and Roberts, M. R. (2003). Do price discreteness and transactions costs affect stock returns? comparing ex-dividend pricing before and after decimalization. *Journal of Finance*, 58(6):2611–2636.

Green, R. C. and Rydqvist, K. (1999). Ex-day behavior with dividend preference and limitations to short-term arbitrage: The case of Swedish lottery bonds. *Journal of Financial Economics*, 53(2):145–187.

Grinblatt, M. S., Masulis, R. W., and Titman, S. (1984). The valuation effects of stock splits and stock dividends. *Journal of Financial Economics*, 13(4):461–490.

Hartzmark, S. M. and Solomon, D. H. (2013). The dividend month premium. *Journal of Financial Economics*, 109(3):640–660.

Hartzmark, S. M. and Solomon, D. H. (2019). The dividend disconnect. *Journal of Finance*, 74(5):2153–2199.

Hartzmark, S. M. and Sussman, A. B. (2025). Price agnostic demand. Working paper, Chicago Booth Research Paper No. 25-12; SSRN 5186859.

Hirshleifer, D., Lim, S. S., and Teoh, S. H. (2009). Driven to distraction: Extraneous events and underreaction to earnings news. *Journal of Finance*, 64(5):2289–2325.

Hong, H. and Sraer, D. A. (2016). Speculative betas. *Journal of Finance*, 71(5):2095–2144. Uses cross-analyst LTG-forecast dispersion as the canonical disagreement measure in the speculative-overvaluation cross-section.

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

Jiang, G., Lee, C. M. C., and Zhang, Y. (2005). Information uncertainty and expected returns. *Review of Accounting Studies*, 10(2–3):185–221.

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

Kalay, A. (1982). The ex-dividend day behavior of stock prices: A re-examination of the clientele effect. *Journal of Finance*, 37(4):1059–1070.

Kirchler, M., Huber, J., and Stöckl, T. (2012). Thar she bursts: Reducing confusion reduces bubbles. *American Economic Review*, 102(2):865–883.

Kreidl, F. (2020). Stock-market behavior on ex-dates: New insights from German stocks with tax-free dividend. *International Journal of Financial Studies*, 8(3):58.

Kumar, A. (2009). Hard-to-value stocks, behavioral biases, and informed trading. *Journal of Financial and Quantitative Analysis*, 44(6):1375–1401.

Lakonishok, J. and Vermaelen, T. (1986). Tax-induced trading around ex-dividend days. *Journal of Financial Economics*, 16(3):287–319.

Lei, V., Noussair, C. N., and Plott, C. R. (2001). Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality. *Econometrica*, 69(4):831–859.

Michaely, R. (1991). Ex-dividend day stock price behavior: The case of the 1986 Tax Reform Act. *Journal of Finance*, 46(3):845–859.

Michaely, R. and Vila, J.-L. (1995). Investors' heterogeneity, prices, and volume around the ex-dividend day. *Journal of Financial and Quantitative Analysis*, 30(2):171–198.

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

Palan, S. (2013). A review of bubbles and crashes in experimental asset markets. *Journal of Economic Surveys*, 27(3):570–588.

Pontiff, J. (2006). Costly arbitrage and the myth of idiosyncratic risk. *Journal of Accounting and Economics*, 42(1–2):35–52.

Poterba, J. M. (1986). The market valuation of cash dividends: The Citizens Utilities case reconsidered. *Journal of Financial Economics*, 15(3):395–405.

Readett, Jr., P. B. (1956). The price behavior of stocks on their ex-dividend dates. S.M. thesis, Massachusetts Institute of Technology, School of Industrial Management.

Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.

Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56(5):1119–1151.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.

Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, 12:579–601.

Zhang, X. F. (2006). Information uncertainty and stock returns. *Journal of Finance*, 61(1):105–137.

Zhang, Y., Farrell, K. A., and Brown, T. A. (2008). Ex-dividend day price and volume: The case of 2003 dividend tax cut. *National Tax Journal*, 61(1):105–127.

A Data construction

Prices and dividends. My main sample is the CIZ CRSP file from WRDS. I select ordinary common stocks on NYSE, AMEX, and Nasdaq. I keep dividend distributions that are classified as ordinary dividend income, declared and paid in U.S. dollars, at monthly, quarterly, semiannual, or annual frequency; where more than one qualifying distribution is recorded on a firm-date, amounts are summed. I filter out stocks worth less than \$5 and dividends below \$0.01, and also remove the events below the first percentile of dividend to cum-dividend price. Size is one-day lagged nominal market capitalization; turnover is the prior 21-day mean of daily volume divided by shares outstanding. Return volatility uses 252 lagged rolling means of squared returns; the realized volatility series excludes daily returns outside ±100%. Table A.1 traces the event count through these screens and the availability restrictions of each regression, and Table A.2 summarizes the variables before the regression-specific availability restrictions.

Table A.1: Construction and coverage of the market sample

Sample definition	Events
Extracted qualifying events, 1963–2025	457,299
After nominal price/dividend screens	444,619
After lowest one-percent yield exclusion	440,173
With valid close-to-close market adjustment	432,582
With yield/size/turnover controls	398,371
With RV and baseline controls	379,300
LTG with at least three analysts	104,813
LTG with at least eight analysts	37,262

Note: The first three rows are sequential screens. Subsequent rows are analysis-specific availability restrictions, not a set of sequential sample restrictions. Market-model estimation requires 60 preceding daily observations; RV requires 252 price-return observations.

Table A.2: Market variables and analyst availability

Variable	*N*	Mean	SD	P10	Median	P90
Dividend to cum-dividend price (percent)	440,173	0.824	0.572	0.246	0.708	1.542
price change (percent)	440,057	-0.520	2.371	-2.963	-0.450	1.852
Adjusted price change (percent)	432,582	-0.589	2.226	-2.890	-0.550	1.623
Nominal cum-day price	440,173	33.807	49.445	9.971	24.350	60.250
Daily RV ($\times 10^4$)	408,561	5.221	5.307	1.358	3.746	10.451
LTG dispersion (growth points)	137,884	3.310	4.618	0.710	2.380	6.260
LTG analyst count	187,520	4.611	4.345	1.000	3.000	11.000

Note: Summary statistics on market variables.

Forecasts. I/B/E/S files were linked to CRSP using link scores at most two and date-valid links. I take the most recent earnings forecasts snapshot strictly preceding the ex-date and no more than 62 calendar days old. LTG is in growth-rate points.

B Additional laboratory evidence

Table B.1 reports the first-trade slope after restricting to sessions with at least twenty players and after adding game intercepts.

Table B.1: Participation and game-level controls

Specification	Slope	SE	Events	Games
All games	0.177	0.033	231	58
At least 20 players	0.174	0.038	208	52
Game fixed effects	0.195	0.050	231	58

Note: The dependent variable is the first-post-event price change. The participation restriction uses the recorded total player count. The final row adds game intercepts. Standard errors are clustered by game in every row.

B.1 Large announcements and price granularity

In this section, I examine announcements of at least \$10, \$20, and \$40 in absolute value. Table B.2 reports first-price regressions in dollars, with an intercept and game-clustered inference.

Table B.2: Price adjustment to large announcements

Minimum |EPS|	First price [95% CI]	Events	Games
10	0.176 [0.108, 0.244]	152	56
20	0.194 [0.118, 0.270]	89	50
40	0.171 [0.037, 0.305]	19	17

Note: OLS of dollar price change on dollar EPS with an intercept. Cutoffs of 10, 20 and 40 dollars were fixed before estimating these exploratory tests; 20 dollars is the primary cutoff. Prices are averaged within completion timestamps; the last timestamp at or before EPS is the baseline. The first response is strictly later. Intervals cluster by game. The primary test against a unit slope has a p-value below 0.001.

B.2 Reconstructed trading capacity

I reconstruct every player's cash and holdings from the initial endowment, all completed transactions, and peek purchases. For positive announcements, I consider a previously active player as having five-share capacity if her pre-event cash covers five shares at the pre-event price plus the announcement; for negative announcements she must own at least five shares. An event is defined to have ample capacity if at least five such players exist. Table B.3 reports the price response in these events and under alternative capacity definitions.

B.3 A full-learning benchmark for the announcement

The main text measures the response against a one-for-one benchmark of Section 4. The announcement, however, reports only the difference between the new card and the card it replaces, and the replaced card is one of the ten that make up the pre-replacement sum. A trader who used the

Table B.3: Price adjustment with measured trading capacity

Sample	First price [95% CI]	Events	Games
Five players with five-share capacity	0.179 [0.105, 0.253]	200	56
Five players with one-share capacity	0.179 [0.110, 0.247]	213	57
Ample capacity: positive announcements	0.214 [0.036, 0.393]	79	41
Ample capacity: negative announcements	0.131 [-0.018, 0.279]	121	53

Note: Nonzero announcements only. A previously active player has five-share capacity after a positive announcement if strictly pre-event cash covers five shares at the pre-event price plus EPS; after a negative announcement the player must own at least five shares of the stock. Ample capacity requires at least five qualifying players. OLS includes an intercept; intervals cluster by game. The slope in the primary sample is significantly below one.

announcement to update her estimate of the replaced card would rationally revise her estimate of the pre-replacement value as well: a large positive announcement makes it likely that a low card was removed, so the pre-replacement sum was probably lower than she thought. The full-learning price change is $E[V_{\text{post}} \mid I, A] - E[V_{\text{pre}} \mid I]$, where A is the announcement and I the public information available before it, and its regression slope on A is below one whenever the ten cards are not fully known.

Table B.4 reports that slope in 20,000,000 simulated stock-games that reproduce the game's card process. With the first SAR as the only public information, the situation of stocks G and Q at their announcements, the full-learning slope is 0.748 for the untruncated sum and 0.630 for the payoff truncated at zero. With both SARs, the situation of R and B, it is 0.833 and 0.694. When the ten cards are known the slope is one, and the truncated value 0.841 is the simulated counterpart of the realized payoff slope in Table 3. The benchmark rises with the information available before the announcement because better information about the card sum leaves less to infer about the replaced card; private peeks move it further toward one.

Table B.4: Full-learning benchmark slopes for the displayed announcement

Public information before the announcement	Untruncated payoff	Payoff truncated at zero
First SAR only (G and Q at their announcement)	0.748	0.630
Both SARs (R and B at their announcement)	0.833	0.694
All ten cards known	1.000	0.841

Note: Each cell is the regression slope of the full-learning price change $E[V_{\text{post}} \mid I, A] - E[V_{\text{pre}} \mid I]$ on the displayed announcement A in 20,000,000 simulated stock-games that reproduce the game's card process: ten cards drawn with replacement from a standard deck, each SAR the sum of five randomly chosen cards plus \$20 truncated at zero, and one uniformly chosen card replaced by a fresh draw. Conditional expectations are cell means on one half of the simulations and are evaluated on the other half, so cell noise cannot inflate the slope. The final row is the accounting identity when the ten cards are known; its truncated entry is the simulated counterpart of the realized payoff slope in Table 3. The benchmark rises with the information available before the announcement because better information about the card sum leaves less to infer about the replaced card.

The observed responses of 0.177 at the first trade and 0.278 after ten trades are between a quarter and a half of the full-learning benchmark, so the shortfall documented in Section 3 does not depend

on whether participants draw this inference; I regard the inference as beyond what participants can do in real time,[7] and the main text maintains the one-for-one benchmark.

Table B.5: Private information and the response after ten trades: full specification

Coefficient	Estimate	Clustered SE	t	p
EPS	0.281	0.065	4.339	< 0.001
Prior information (z_{gs})	4.617	3.208	1.439	0.156
EPS $\times$ prior information	0.091	0.043	2.104	0.040
Stock FE	✓			
Game FE	✓			
Events	223			
Games	57			
R^2	0.427			

Note: Coefficient estimates for the ten-trade column of Table 8, using equation (8). The dependent variable is the price at the tenth distinct post-announcement transaction timestamp minus the last price at or before the announcement; simultaneous transaction prices are averaged. Prior information is the standardized logarithm of one plus cumulative pre-event peeks on the announced stock, using the mean and standard deviation of all 232 events, so the EPS coefficient is the slope at $z_{gs} = 0$. The regression includes game and stock fixed effects, whose coefficients are not reported. Standard errors cluster by game with the conventional finite-sample correction.

C Additional market results

Table C.1 reports the regression behind Figure 1: the ex-day price change on the dividend yield, interacted with quintiles of dividend yield, size, and turnover (Section 2).

[7] Two examples convey what the inference requires, even before any SAR is used. Suppose the screen shows +30. The trader must list the ordered pairs (removed card, new card) in the deck of footnote 3 whose difference is exactly 30: a −20 replaced by a +10, a −10 by a +20, or a +10 by a +40. The deck holds 2, 8, and 2 cards of the three removed values and 2, 8, and 2 of the three new values, so these cases account for $2 \times 2 = 4$, $8 \times 8 = 64$, and $2 \times 2 = 4$ of the $52 \times 52 = 2{,}704$ equally likely pairs. The removed card was therefore worth −20, −10, or +10 with probabilities 4/72, 64/72, and 4/72, and its expected value is $(-80 - 640 + 40)/72 = -9.44$ rather than the unconditional +4: the pre-replacement sum was 13.44 lower than she believed, and the correct revision of the value is $30 - 13.44 = 16.56$, or 0.55 per dollar announced. A +6 is harder, not easier. Ten removed values are consistent with it (−10, −9, −8, −2, +4, +6, +8, +10, +12, +14), accounting for 16, 4, 4, 4, 4, 4, 4, 4, 4, and 16 of the 64 consistent pairs, so the expected removed card is $148/64 = 2.31$, the pre-replacement sum is revised down by 1.69, and the correct update is 4.31, or 0.72 per dollar. Conditioning on a SAR adds to each calculation the probability that the removed card was among the five cards the SAR summed, the distribution of the ten cards given the SAR total, and the truncation at zero, which is why Table B.4 is computed by simulation.

Table C.1: Ex-day price changes and dividend yield, size, and turnover: joint regression

	Main effect	Interaction with x_{it}
Dividend yield (x_{it})	-0.306^{***}	
	(0.070)	
Dividend-yield quintile		
Quintile 2	0.00181^{***}	-0.326^{***}
	(0.00029)	(0.074)
Quintile 3	0.00262^{***}	-0.318^{***}
	(0.00031)	(0.069)
Quintile 4	0.00376^{***}	-0.386^{***}
	(0.00034)	(0.067)
Quintile 5	0.00494^{***}	-0.501^{***}
	(0.00040)	(0.067)
Size quintile		
Quintile 2	0.00089^{***}	-0.193^{***}
	(0.00031)	(0.034)
Quintile 3	0.00072^{**}	-0.230^{***}
	(0.00031)	(0.034)
Quintile 4	0.00137^{***}	-0.357^{***}
	(0.00029)	(0.032)
Quintile 5	0.00131^{***}	-0.449^{***}
	(0.00028)	(0.031)
Turnover quintile		
Quintile 2	-0.00047^{*}	-0.010
	(0.00027)	(0.029)
Quintile 3	0.00000	-0.062^{**}
	(0.00025)	(0.027)
Quintile 4	0.00015	-0.066^{**}
	(0.00026)	(0.028)
Quintile 5	-0.00011	-0.030
	(0.00028)	(0.034)
Constant	-0.00185^{***}	
	(0.00033)	
R^2	0.051	
Events	398,371	
Ex-dates	15,657	

Note: OLS of the market-adjusted ex-day price change y_{it}^{adj} on the dividend yield x_{it}, indicators for quintiles 2–5 of dividend yield, size, and turnover, and the interaction of each indicator with x_{it}. The lowest quintile of each characteristic is the omitted category. The first column reports the coefficient on each variable and the second the coefficient on its interaction with x_{it}: the coefficient on x_{it} is the slope for events in the lowest quintile of all three characteristics, and each interaction is the change in that slope. Figure 1 plots the implied slope in each quintile, averaged over the sample distribution of the other two characteristics. Size quintiles are formed within ex-date, and dividend-yield and turnover quintiles within calendar year, among sample events. Standard errors, in parentheses, cluster by ex-date; ***, **, * denote significance at 1%, 5%, and 10%.