Mispricing and Information Acquisition

Haisong Lin* **Lihong Zhang†**
Xiaoquan (Michael) Zhang‡

October 16, 2026

Abstract

Standard theory holds that mispricing attracts informed capital until prices correct. We show the reverse can occur. When an asset is already mispriced, acquiring information and exploiting the mispricing become substitutes, so large mispricing deters acquisition even when information is free. In a strategic trading model with competitive informed traders, acquiring information raises both the opportunity cost of exploiting the mispricing and the acquirer's price impact. The value of information is therefore nonmonotonic in mispricing: moderate underpricing stimulates information production, while large mispricing crowds it out. In a dynamic extension, deterred acquisition slows price correction, so large anomalies can become partly self-sustaining: inefficiency that shields itself from the capital that would erase it.

Keywords: Information Acquisition, Mispricing, Price-driven Trading

JEL Codes: D82, D83, G12, G14

*School of Economics and Management, Tsinghua University. Email: linhs20@mails.tsinghua.edu.cn
†School of Economics and Management, Tsinghua University. Email: zhanglh2@sem.tsinghua.edu.cn
‡CUHK Business School, Chinese University of Hong Kong. Email: zhang@cuhk.edu.hk

1 Introduction

Informed trading is the force that pushes prices toward fundamental value, and the standard presumption is that the force scales with the deviation: the further the price lies from value, the larger the profit from becoming informed, and the faster the correction (Grossman and Stiglitz, 1980; Kyle, 1985; Verrecchia, 1982). This paper shows that the presumption fails once the deviation is already in the price. A trader who can become informed at no cost prefers not to when the price is far from value, and the preference is stronger the larger the deviation. The reason is that acquiring information and exploiting a known mispricing are substitutes: acquisition makes the trader's order flow more informative, raises the trader's price impact, and thereby lowers the return to trading against the mispricing.

We develop the argument in a one-period strategic trading model with two risk-neutral informed traders, a competitive market maker, and noise traders, building on Kyle (1985) and Foster and Viswanathan (1994). The price carries a proportional wedge θ relative to the market maker's conditional expectation of the liquidation value, with θ_0 the exogenous wedge at the outset and θ_1 the wedge after trading. Both informed traders observe a common signal about the liquidation value. One of them, B, can acquire at no cost an additional signal that reveals the liquidation value, and the acquisition decision is observed by the market before trading. A trader who has not acquired information still profits by trading against the wedge, buying when the asset is underpriced and selling when it is overpriced; we call this price-driven trading, and it is absent from efficient-market models, in which private information is the only source of profit. We ask how B's decision depends on the size of θ_0.

The main result (Proposition 4.2) is that B acquires information if and only if θ_0 lies in a bounded interval $[\underline{\theta}, \bar{\theta}]$ around the efficient price $\theta_0 = 1$. The interval is determined by two primitives: the variance of the acquirable signal relative to the common signal, q, and the mispricing rent κ, the squared prior mean of value relative to the variance of the common signal. Two forces generate the bound. The first, which we call the exploitation effect, is an opportunity cost: acquisition raises price impact for every trader, which lowers the profit from price-driven trading, and the forgone profit grows without bound in the distance of θ_0 from one. The second, which we call the price impact effect, is the proportional increase in price impact caused by acquisition. This increase is decreasing in θ_0 (Lemma 4.1): underpricing

makes the market deep before acquisition, so acquisition raises price impact by a large proportion, whereas overpricing makes the market shallow and the proportional increase small. Under overpricing the exploitation effect alone closes the acquisition region; under underpricing the two forces reinforce each other. The interval widens in q, narrows in κ, collapses to the efficient price as the common signal vanishes (Propositions 4.3 and 4.4), and does not depend on the volume of noise trading.

The value of information is nonmonotone in the wedge. Although large mispricing deters acquisition, the gain from acquiring is maximized at an interior level of underpricing $\theta^* < 1$, not at the efficient price (Proposition 4.5). Underpricing deepens the market in absolute terms, so an informed trader disguises larger positions and earns more on the acquired signal, and for moderate underpricing this gain exceeds the exploitation cost. Section 6 applies the result to initial public offerings, where the offer price is a deliberate choice: the model predicts an inverse-U relationship between underpricing and pre-market information production, with the information-maximizing discount increasing in q and decreasing in κ.

Section 5 extends the model to N trading rounds. For two rounds we solve the equilibrium analytically (Supplemental Appendix C.5); for more rounds we compute it numerically. In the one-period model the better-informed trader always earns more. Over several rounds the less-informed trader trades the mispricing more aggressively than the better-informed trader in every round but the last, a property we prove for two rounds and observe numerically for four, and when θ_0 is far from one this aggression is enough for the less-informed trader to earn more in total. Acquisition then fails even for a trader who cares only about outperforming the rival. The dynamics also close a loop: deterred acquisition slows the correction of the wedge, and slow correction prolongs the environment in which exploitation dominates.

The mechanism bears on several questions in the economics of information. It explains why falling costs of information need not raise price informativeness (Farboodi and Veldkamp, 2020): the binding constraint is the opportunity cost of forgoing exploitation, not the cost of acquisition. It implies that heavy participation by sophisticated investors in a mispriced asset need not reflect fundamental research, and that liquidity improvements can raise the informativeness of prices about information already held while lowering it about information not yet acquired, which reconciles the finding that algorithmic trading improves short-run price efficiency while reducing the production of costly fundamental information (Weller, 2018). It also offers

an account of the slow decay of return predictability after an anomaly is documented (McLean and Pontiff, 2016): the traders best positioned to eliminate a large anomaly prefer to exploit it.[1]

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 characterizes the equilibrium of each subgame and introduces price-driven trading. Section 4 analyzes the acquisition decision. Section 5 extends the model to a dynamic setting, which we solve numerically. Section 6 applies the model to IPO underpricing and derives testable predictions. Section 7 presents extensions. Section 8 concludes. Proofs are in the Supplemental Appendix.

Related Literature. Our paper speaks to four literatures. The first studies endogenous information acquisition in financial markets (Grossman and Stiglitz, 1980; Verrecchia, 1982; Hébert and La'O, 2023), where the acquisition decision turns on a quality–cost trade-off and the prevailing price is correct. Mispricing is assumed away, and with it the very margin we study. We show that once prices can deviate from value, a second trade-off appears: acquiring information competes with exploiting the mispricing already in the price, because acquisition raises a trader's price impact and erodes the exploitation return. This exploration–exploitation channel inverts the canonical comparative static, deterring acquisition rather than encouraging it, and operates even when information is costless.

A second literature establishes that more information can be privately or socially harmful, and two papers are closest to our deterrence result. Nezafat and Schroder (2023) show that the value of private information turns negative for risk-averse traders when noise-trading volume is high; our deterrence requires neither risk aversion nor heavy noise trading, and is driven by mispricing alone. Chaigneau et al. (2018) show that better information can lower welfare in a contracting setting through an incentive channel; our channel is instead a market one, in which information raises price impact and thereby destroys the return to exploiting mispricing.

A third literature asks why mispricing persists (Stambaugh et al., 2015; Jacobs, 2016). Existing explanations locate persistence in forces external to information production: limits to arbitrage (Gromb and Vayanos, 2010), behavioral biases (Barberis et al., 1998), or feedback effects (Edmans et al., 2015). We add an internal one.

[1]Quantitative investors now account for roughly 35% of U.S. stock market ownership, 60% of institutional equity assets under management, and 60% of trading volume (*The Economist*, 2019, as reported in Dyer et al., 2025).

Large mispricing deters the information production that would correct it, so the deviation protects itself from the capital that would close it. Persistence then reflects not only risk and constraints but the endogenous diversion of informed effort away from fundamental learning.

A fourth literature studies strategic trading in the tradition of Kyle (1985). These models assume efficient prices, so the only motive to trade is private information. Embedding mispricing in the competitive structure of Foster and Viswanathan (1994) adds a second motive, price-driven trading, through which even an uninformed trader profits by correcting the distortion. This motive also separates our dynamics from the multi-period analyses of Holden and Subrahmanyam (1992) and Back et al. (2000): the race to exploit mispricing, absent when prices are efficient, is what lets a less-informed trader outperform a better-informed rival. On the empirical side, Anderson (2022) finds that underpricing induced by mutual-fund flows raises the consumption of accounting information, consistent with our prediction that moderate underpricing stimulates information production.

2 The Model

2.1 Model Setting

We incorporate information acquisition and asset mispricing into a Kyle (1985) model with competition among informed traders. The economy has three dates, $t = 0, 1, 2$, and the market operates only on date 1. Figure 1 describes the timeline. There is a single risky asset with a date-2 liquidation value v, which decomposes as

$$v = \bar{v} + f + e, \tag{1}$$

where $\bar{v}$ is the prior mean, $f \sim N(0, \Omega_0)$ is common information revealed to informed traders before trading, and $e \sim N(0, \Lambda_0)$ is extra information that is unobservable without information acquisition. Without loss of generality, we assume $\bar{v} \geq 0$, and that f and e are independent. Let $\Sigma_0 = \Omega_0 + \Lambda_0$ denote the total variance of the liquidation value.

There are three groups of market participants: two risk-neutral informed traders, B and L, who are endowed with common information; one risk-neutral market maker; and noise traders. B and L differ in their ability to acquire information, in that only B

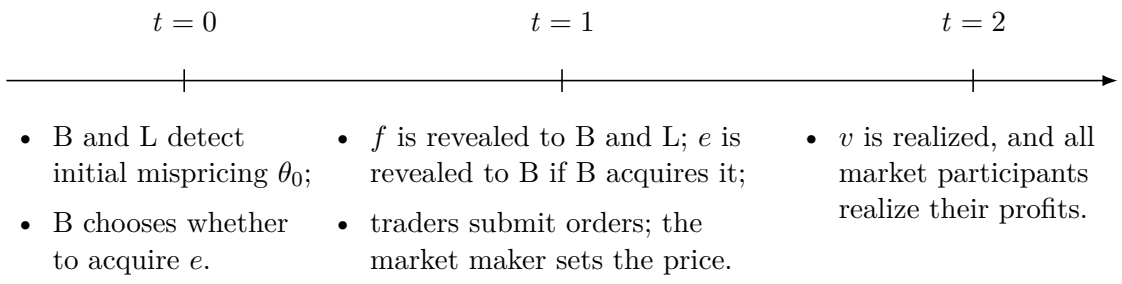

Figure 1: Timeline.

can acquire extra information at date 0 to become better informed. We assume that information acquisition is costless.[2] If information acquisition is chosen, the extra information is revealed to B before trading, together with the common information. With the extra information, B directly observes the liquidation value.[3] B chooses to acquire information only if B's unconditional expected profits (abbreviated as ex ante profits) increase after acquiring the extra information.

Common information f is shared by all investors (for example, through earnings calls or collaborative field research), whereas extra information e is privately accessible only to B (for example, from a unique firm relationship).[4]

ASSUMPTION 2.1 ***(Observable Acquisition)*** *B's acquisition decision is made at date* 0*, before any private signal is observed, and is observed by L and the market maker before trading at date* 1*.*

Assumption 2.1 is the assumption on which the deterrence results rest, and it deserves a precise statement of its role. Information is costless, so if B's decision were unobservable, then holding the pricing rule and L's strategy fixed, acquiring e would raise B's expected profit by the value of trading on e alone, which is strictly positive (Corollary 3.1). B would then always acquire. Deterrence arises only because the market anticipates that B is better informed and prices order flow accordingly. The assumption is natural when the identity of informed players is public, as when

[2]Information acquisition costs do not affect the core mechanism of the exploration–exploitation trade-off. We consider costly information acquisition in the smooth information acquisition extension, and the baseline results remain valid.

[3]The assumption that B can directly observe the liquidation value greatly simplifies the inference problem when the model is extended to a dynamic setting. We maintain consistency between the one-shot and dynamic settings.

[4]We restrict L's ability to acquire information to ensure that e, once acquired, is incremental to the economy, rather than simply shifting informational competition with a counterparty.

a fund's research mandate or an analyst's coverage decision is disclosed, or when algorithmic traders are known to operate in a security (Weller, 2018). Allowing B to acquire after observing f would add a signaling dimension to the acquisition decision without altering the mechanism.

Trading takes place at date 1. As in Kyle (1985), B and L do not observe the date-1 price when submitting orders, so they submit market orders x_B and x_L to maximize expected trading profits conditional on their private information:

$$\max_{x_i} \mathbb{E}\left[\pi_i \mid \mathcal{F}^i\right], \quad \text{where} \quad \pi_i = x_i(v - p_1), \quad i \in \{B, L\}. \tag{2}$$

Here p_1 is the date-1 equilibrium price set by the market maker, and π_B and π_L denote the trading profits of B and L. L's information set is always $\mathcal{F}^L = \{f\}$, whereas B's information set is $\mathcal{F}^B = \{f, e\}$ with information acquisition and $\mathcal{F}^B = \{f\}$ without information acquisition.

2.2 Mispricing

In addition to the orders submitted by the informed traders, noise traders submit random market orders $z \sim N(0, \sigma_z^2)$, independent of all other random variables. The market maker observes the total order flow $y = x_B + x_L + z$. We measure mispricing as a proportional wedge between the price and the market maker's rational Bayesian conditional expectation. Specifically, the date-1 price satisfies

$$p_1 = \theta_1 \mathbb{E}\left[v \mid \mathcal{F}^M = \{y\}\right], \quad \theta_1 > 0. \tag{3}$$

The pricing rule that generates this wedge is the primitive of the model. The market maker anchors to the prevailing price and adjusts linearly in total order flow, $p_1 = p_0 + \lambda y$, with the price impact λ and the wedge θ_1 jointly determined by the requirement that the resulting price satisfy Equation (3). Two features distinguish this rule from efficient pricing. First, the conditional expectation in Equation (3) is computed correctly, so the wedge does not stem from errors in inference. Second, the price responds to total order flow rather than to its surprise component, so the predictable order flow generated by traders who exploit the wedge moves the price. The date-1 wedge θ_1 is therefore not chosen by anyone: it is the value for which an anchored linear rule is consistent with a correct conditional expectation,

and it is through this consistency requirement that exploitation corrects mispricing (Section 3.1).

We regard the rule as a reduced form for an anchored pricing mechanism rather than as the outcome of an optimizing market maker. Explicit frictions deliver wedges of this kind: noise trader sentiment (de Long et al., 1990), limits to arbitrage (Gromb and Vayanos, 2010), behavioral overreaction and underreaction (Barberis et al., 1998), and stochastic mispricing errors (Brennan and Wang, 2010). The parameter θ abstracts from the source of the wedge, is a sufficient statistic for its size, and is empirically measurable using composite mispricing indices (Stambaugh and Yuan, 2017). Section 8 discusses what a full microfoundation would add.

Given an exogenous initial mispricing parameter $\theta_0 > 0$ and B's choice of information acquisition, θ_1 is determined in equilibrium. At date 0, both B and L observe θ_0 and the initial price $p_0 = \theta_0 \bar{v}$, the price that prevails when no private information is available.

When $\theta_0 = 1$, the market is already efficient, and the model collapses to the standard strategic trading framework of Kyle (1985). When $\theta_0 \neq 1$, the actual price does not correctly reflect all available public information. We refer to the case where $\theta_0 > 1$ as initial overpricing and the case where $\theta_0 < 1$ as initial underpricing, as the mispricing either amplifies or attenuates the forecast of the liquidation value based on all available information.

3 Equilibrium Results

We analyze two distinct subgames, contingent on whether B acquires extra information. When B does not acquire information, B and L share an identical information set derived solely from common information; we call this the symmetric subgame. When B acquires extra information, B holds an informational advantage over L; we call this the asymmetric subgame. We mark parameters in the symmetric subgame with a hat.

We restrict attention to linear equilibria, defined as follows.

DEFINITION 3.1 ***(Linear Equilibrium)*** *Fix $\theta_0 > 0$ and B's acquisition decision. A linear equilibrium consists of order strategies x_B and x_L, each linear in the trader's information advantages and in p_0; a pricing rule $p_1 = p_0 + \lambda y$ with $\lambda > 0$;*

and a wedge $\theta_1 > 0$, such that (i) given the other trader's strategy and the pricing rule, each informed trader's order maximizes Equation (2)*; and (ii) at the equilibrium order flow, the price satisfies Equation* (3)*.*

Condition (i) is the standard best-response requirement. Condition (ii) replaces the zero-profit condition of an efficient market: the market maker's anchored linear rule must coincide with a scaled conditional expectation, and the scale θ_1 is determined jointly with λ rather than chosen. Throughout, equilibrium refers to a linear equilibrium.

3.1 Symmetric Subgame: Without Information Acquisition

In the symmetric subgame, we first discuss the structure of the optimal order-submission strategy. By ex ante homogeneity of B and L, we omit subscripts and write each informed trader's order as $\hat{x}$. As in Kyle (1985), to conceal private information, informed traders trade on their information advantage from common information relative to the market maker, which we denote s_L because it is also L's advantage in the asymmetric subgame:

$$s_L = \mathbb{E}\left[v \mid f\right] - \bar{v} = f. \tag{4}$$

We call the strategy of trading on an information advantage *information-driven trading*; it exploits the gap between a trader's private valuation and the market maker's forecast. When the price also deviates from market efficiency, that is, from the perspective of public information, traders adopt *price-driven trading* alongside information-driven trading, and through it the initial mispricing is endogenously corrected. The unique linear equilibrium in the symmetric subgame is characterized in the following proposition.

PROPOSITION 3.1 ***(Symmetric Subgame)*** *When B does not acquire extra information, there exists a unique linear equilibrium of the form:*

$$\begin{cases} \hat{x} = \hat{\gamma} s_L + \hat{\xi} p_0, \\ \hat{p}_1 = p_0 + \hat{\lambda} \hat{y}. \end{cases} \tag{5}$$

The date-1 mispricing parameter is

$$\hat{\theta}_1 = \frac{\theta_0}{3} + \frac{2}{3}. \tag{6}$$

The price impact is

$$\hat{\lambda} = \frac{1}{3\sigma_z}\sqrt{2\theta_0\Omega_0}. \tag{7}$$

The trading intensities are

$$\hat{\gamma} = \frac{1}{3\hat{\lambda}} \quad \textit{and} \quad \hat{\xi} = \frac{1-\theta_0}{3\hat{\lambda}\theta_0}. \tag{8}$$

The optimal order-submission strategy combines information-driven trading and price-driven trading, with intensities $\hat{\gamma}$ and $\hat{\xi}$, respectively. Consistent with the literature, informed traders place positive weight on their information advantage from common information. The novel element is the existence of endogenous price-driven trading to exploit mispricing: in addition to information-driven trading, the informed traders buy (sell) extra units based on p_0 under initial underpricing (overpricing). To see why price-driven trading arises, consider the price increment when informed traders submit orders with price-driven trading intensity $\hat{\xi}$:

$$\hat{p}_1 - p_0 = \left(\frac{\hat{\theta}_1}{\theta_0} - 1 - 2\hat{\lambda}\hat{\xi}\right) p_0 + \hat{\lambda}\hat{y}. \tag{9}$$

In an efficient market, $\frac{\hat{\theta}_1}{\theta_0} - 1 = 0$. This implies that any price-driven trading would violate the linear structure defined in Equation (5); thus, informed traders rely exclusively on their information advantage. Conversely, in the presence of initial mispricing, $\frac{\hat{\theta}_1}{\theta_0} - 1 \neq 0$. Here, price-driven trading with intensity $\hat{\xi}$ serves to exactly neutralize this term, preserving the linear equilibrium structure and the price impact.

Price-driven trading stems from mispricing rather than private information and would arise even for an uninformed trader; it does not constitute risk-free arbitrage because conditioning on historical prices entails ex ante risk.

The equilibrium mispricing parameter $\hat{\theta}_1$ is a weighted average of θ_0 and one (with weights 1/3 and 2/3), so price-driven trading moves the market toward efficiency. Because $\hat{\theta}_1$ is bounded below by 2/3, severe underpricing is substantially corrected,

while significant overpricing can persist.

This correction mechanism is illustrated by Equation (9). To maintain that the price increment is proportional to total order flow, the condition $\frac{\hat{\theta}_1}{\theta_0} - 1 = 2\hat{\lambda}\hat{\xi}$ must hold. Under initial underpricing (overpricing), informed traders buy (sell) additional units to exploit the distortion, thereby driving $\hat{\theta}_1$ toward one. This establishes a reciprocal relationship between mispricing and price-driven trading: mispricing creates opportunities for price-driven traders to exploit, whose subsequent activity corrects the price distortion.

Next, price impact $\hat{\lambda}$ increases in θ_0. In the absence of information acquisition, extra information does not alter price impact, but initial underpricing (overpricing) attenuates (amplifies) the effect of common information on price impact. Market depth, the reciprocal of price impact, therefore decreases in θ_0: a larger θ_0 implies greater illiquidity. As $\theta_0 \to 0$, the market becomes infinitely deep and prices relatively stable. As $\theta_0 \to +\infty$, the market becomes infinitely shallow and prices relatively unstable.

3.2 Asymmetric Subgame: With Information Acquisition

We now examine the asymmetric subgame in which B acquires extra information. As in the symmetric case, informed traders engage in both information-driven and price-driven trading. Because L cannot infer the content of B's extra information in the one-shot setting, L's informational advantage relative to the market maker remains $s_L = f$. B's advantage, however, now consists of two components: (i) B's informational advantage relative to L, denoted by s_B in Equation (10); and (ii) L's advantage relative to the market maker, s_L.

$$s_B = v - \mathbb{E}[v \mid f] = e. \tag{10}$$

The unique linear equilibrium in the asymmetric subgame is characterized in the following proposition.

PROPOSITION 3.2 ***(Asymmetric Subgame)*** *When B acquires extra informa-*

tion, there exists a unique linear equilibrium of the form:

$$\begin{cases} x_B = \beta s_B + \gamma s_L + \xi p_0, \\ x_L = \psi s_L + \zeta p_0, \\ p_1 = p_0 + \lambda y. \end{cases} \tag{11}$$

*The date-*1 *mispricing parameter is*

$$\theta_1 = \frac{\theta_0}{3} + \frac{2}{3}. \tag{12}$$

The price impact is

$$\lambda = \frac{1}{6\sigma_z}\sqrt{(6\Lambda_0 + 8\Omega_0)\theta_0 + 3\Lambda_0}. \tag{13}$$

The trading intensities are

$$\beta = \frac{1}{2\lambda}, \quad \gamma = \psi = \frac{1}{3\lambda}, \quad \textit{and} \quad \xi = \zeta = \frac{1-\theta_0}{3\lambda\theta_0}. \tag{14}$$

As in the symmetric subgame, λ increases with θ_0. Due to the incremental value in forecasting, extra information enters the expression for price impact via the term $(6\theta_0 + 3)\Lambda_0$ under the square root. This makes λ strictly larger than $\hat{\lambda}$, implying a shallower market than in the symmetric case. The date-1 mispricing parameter is the same across the two subgames, so information acquisition does not affect the correction of mispricing in the one-shot model.

B trades on both extra and common information (with intensities β and γ, respectively), while L trades only on common information.[5] Because e, f, and p_0 are independent, B and L share the same intensities for common-information-driven and price-driven trading. Both intensities are lower than in the symmetric subgame because the impact of B's extra information makes the market shallower. B trades more aggressively on extra information than on common information because the latter faces competition from L.

[5]Throughout, "trading on common (extra) information" means trading on the informational advantage derived from common (extra) information.

3.3 Exploitation of Mispricing

The informed traders' ex ante profits are characterized in the following corollary.

COROLLARY 3.1 *The following properties of informed traders' ex ante profits hold:*

(i) In the symmetric subgame, each informed trader's ex ante profits are

$$\mathbb{E}\hat{\pi} = \underbrace{\frac{\Omega_0}{9\hat{\lambda}}}_{\textit{Information-driven}} + \underbrace{\frac{(1-\theta_0)^2\bar{v}^2}{9\hat{\lambda}}}_{\textit{Price-driven}}. \tag{15}$$

(ii) In the asymmetric subgame, B and L's ex ante profits are

$$\begin{aligned} \mathbb{E}\pi_B &= \underbrace{\frac{\Lambda_0}{4\lambda} + \frac{\Omega_0}{9\lambda}}_{\textit{Information-driven}} + \underbrace{\frac{(1-\theta_0)^2\bar{v}^2}{9\lambda}}_{\textit{Price-driven}}, \\ \mathbb{E}\pi_L &= \underbrace{\frac{\Omega_0}{9\lambda}}_{\textit{Information-driven}} + \underbrace{\frac{(1-\theta_0)^2\bar{v}^2}{9\lambda}}_{\textit{Price-driven}}. \end{aligned} \tag{16}$$

B's ex ante profits strictly exceed those of L.

In both subgames, ex ante profits decompose into information-driven and price-driven trading profits, so exploiting mispricing through price-driven trading is profitable. Both components are proportional to market depth $(1/\lambda)$: a deeper market lets informed traders better conceal their private information, trade more aggressively, and earn higher profits. Beyond market depth, profits from common and extra information are proportional to Ω_0 and Λ_0, which capture the value of holding such information.

Price-driven trading profits can also be written as the product of the price-driven trading intensity and $(1-\theta_0)^2\bar{v}^2/3$. A higher price-driven intensity contributes more to correcting mispricing and earns a higher market reward. Holding $\bar{v}$ constant, a larger deviation of θ_0 from one widens the exploitation opportunity and raises price-driven profits. The appearance of $\bar{v}$ in the ex ante profit function is novel: in an efficient market only the de-meaned distribution of v matters, so $\bar{v}$ is irrelevant (Kyle, 1985). Under mispricing, for a given θ_0, a larger $\bar{v}$ implies a larger absolute

deviation of the price from fundamental value and hence a greater opportunity to exploit mispricing. We therefore interpret $\bar{v}^2$ as the mispricing rent.

In the asymmetric subgame, B and L generate identical profits from common-information-driven trading and price-driven trading because they share the same trading intensities in these activities; however, B strictly outperforms L by capturing additional profits from extra information. Compared to the symmetric subgame, the profits from common-information-driven and price-driven trading are lower due to the larger price impact resulting from information acquisition.

3.4 Price Informativeness

Information acquisition raises price informativeness regardless of the initial mispricing. Following Grossman and Stiglitz (1980) and Goldstein et al. (2025), we define price informativeness as the precision of the forecast of the liquidation value conditional on the equilibrium price. In the two subgames it is

$$\begin{aligned} \hat{PI} &= \frac{1}{\mathrm{Var}\left(v \mid \hat{p}_1\right)} = \left(\Lambda_0 + \frac{\theta_0}{\theta_0+2}\Omega_0\right)^{-1}, \\ PI &= \frac{1}{\mathrm{Var}\left(v \mid p_1\right)} = \left[\frac{2\theta_0+1}{2(\theta_0+2)}\Lambda_0 + \frac{\theta_0}{\theta_0+2}\Omega_0\right]^{-1}. \end{aligned} \tag{17}$$

While the literature often equates price informativeness with market efficiency, our model underscores the distinction between the two. Price informativeness measures the extent to which the price reduces uncertainty (precision), whereas mispricing captures the deviation of the price from the fundamental value (bias). Even if the market aggregates all available information efficiently, price informativeness remains bounded if the underlying signals are noisy. Conversely, mispricing persists even when the liquidation value is fully revealed to the informed traders, because the wedge is a property of the pricing rule rather than of the information set.

In Equation (17), the weight on Ω_0 is identical across subgames, but the weight on Λ_0 is smaller under information acquisition, yielding a smaller conditional variance and thus higher price informativeness. The improvement strictly decreases in θ_0, so initial underpricing amplifies the gain in informativeness relative to the efficient-market benchmark. Underpricing enhances market liquidity and induces informed traders to trade more aggressively. As $\theta_0 \to +\infty$ the improvement vanishes; as

$\theta_0 \to 0$ the ratio $PI/\hat{PI}$ approaches its supremum of 4.

LEMMA 3.1 *The ratio of price informativeness in the two subgames,*

$$\frac{PI}{\hat{PI}} = 1 + \frac{3\Lambda_0}{2\Sigma_0\theta_0 + \Lambda_0}, \tag{18}$$

is strictly greater than one and decreases with θ_0.

4 Choice of Information Acquisition

This section analyzes how initial mispricing affects absolute information acquisition (Definition 4.1). Large initial mispricing deters information acquisition, whereas moderate initial underpricing raises the gains from it. At date 0, after observing θ_0, B compares ex ante profits in the symmetric and asymmetric subgames and acquires extra information only if doing so yields a higher payoff.

DEFINITION 4.1 ***(Absolute Information Acquisition)*** *B acquires extra information if B's ex ante profits over the entire trading period in the asymmetric subgame exceed those in the symmetric subgame.*[6]

To analyze the incentives for absolute information acquisition, we decompose the change in B's and L's ex ante profits resulting from information acquisition as follows:

$$\begin{aligned}
\mathbb{E}\pi_B - \mathbb{E}\hat{\pi} &= \underbrace{\frac{\Lambda_0}{4\lambda} + \left(\frac{1}{\lambda} - \frac{1}{\hat{\lambda}}\right)\frac{\Omega_0}{9}}_{\text{Information-driven}} + \underbrace{\left(\frac{1}{\lambda} - \frac{1}{\hat{\lambda}}\right)\frac{(1-\theta_0)^2\,\bar{v}^2}{9}}_{\text{Price-driven}}, \\
\mathbb{E}\pi_L - \mathbb{E}\hat{\pi} &= \underbrace{\left(\frac{1}{\lambda} - \frac{1}{\hat{\lambda}}\right)\frac{\Omega_0}{9}}_{\text{Information-driven}} + \underbrace{\left(\frac{1}{\lambda} - \frac{1}{\hat{\lambda}}\right)\frac{(1-\theta_0)^2\,\bar{v}^2}{9}}_{\text{Price-driven}}.
\end{aligned} \tag{19}$$

4.1 Absence of Information Acquisition

Because information acquisition increases price impact, it reduces common-information-driven and price-driven trading profits for both B and L. Since L relies exclusively on these sources, L is strictly worse off whenever B acquires information. B, however,

[6] In Section 5, we define relative information acquisition to consider B's strategic motives.

also gains from trading on private information, and in an efficient market these gains dominate: B always acquires information.

PROPOSITION 4.1 ***(No Mispricing)*** *In an efficient market:*

(i) B always engages in absolute information acquisition;

(ii) L is strictly worse off when B acquires extra information.

In an efficient market, extra information benefits the acquirer and harms the non-acquirer. Under mispricing, B instead suffers a reduction in price-driven trading profits relative to the efficient-market benchmark, and this reduction can be substantial when initial mispricing is large. The gains from trading on extra information then need not outweigh these losses, so acquiring extra information can reduce B's total ex ante profits.

Figure 2 previews the main results. Panel (b) confirms that L's profits fall when B acquires information. Panel (a) shows that B acquires information only when $\theta_0 \in [\underline{\theta}, \bar{\theta}]$ (the green shaded region), so large mispricing deters acquisition (the red shaded region).[7] Under large initial underpricing, even the change in B's information-driven trading profits from acquisition (blue line) is negative. We now characterize the two mechanisms formally.

To facilitate the analysis, we normalize the quality of extra information and the mispricing rent by the variance of common information:

$$q = \frac{\Lambda_0}{\Omega_0} \quad \text{and} \quad \kappa = \frac{\bar{v}^2}{\Omega_0}.$$

Here, q and κ quantify the value of exploring extra information and exploiting mispricing relative to common information, respectively. Since information acquisition is assumed to be costless, Equation (19) implies that the decision to acquire information is jointly determined by q, κ, and θ_0. This decision rule is characterized in the following proposition.

PROPOSITION 4.2 *Given $q > 0$ and $\kappa > 0$, the following statements hold:*

[7] A positive fixed acquisition cost shifts the net-benefit curve downward, narrowing the interval around $\theta_0 = 1$.

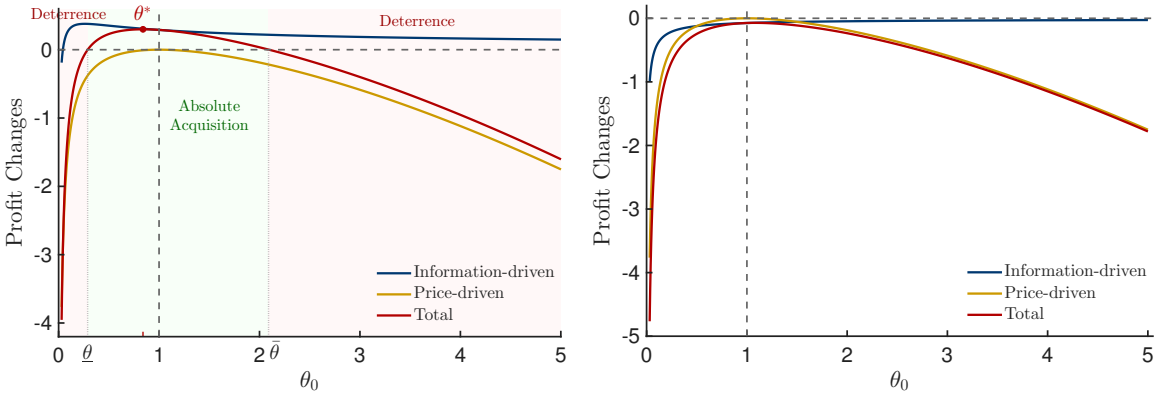

(a) Informed Trader B (b) Informed Trader L

Figure 2: This figure shows the changes in ex ante profits resulting from information acquisition for informed traders B ($\mathbb{E}\pi_B - \mathbb{E}\hat{\pi}$, Panel (a)) and L ($\mathbb{E}\pi_L - \mathbb{E}\hat{\pi}$, Panel (b)), as a function of θ_0. The total change (red) is decomposed into information-driven (blue) and price-driven (yellow) components. In Panel (a), $\underline{\theta}$ and $\bar{\theta}$ denote the zero-crossing points of the red line and θ^* its maximizer. Parameters: $\Lambda_0 = \Omega_0 = \sigma_z = 1$, $\bar{v} = 2$.

(i) There exist two thresholds, uniquely determined by q and κ, denoted by $\underline{\theta}(q,\kappa) \in (0,1)$ and $\bar{\theta}(q,\kappa) > 1$. B engages in absolute information acquisition if and only if $\theta_0 \in [\underline{\theta}(q,\kappa), \bar{\theta}(q,\kappa)]$.

(ii) L is strictly worse off when B acquires extra information.

The efficient market case is nested within this result: $\theta_0 = 1$ always falls within the interval $[\underline{\theta}(q,\kappa), \bar{\theta}(q,\kappa)]$. Introducing mispricing expands this single point into an interval, the boundaries of which are determined by the relative value of exploring extra information (q) and exploiting mispricing (κ).[8]

Relative to the standard Kyle (1985) model, mispricing introduces two distortions. First, it induces price-driven trading, a new source of profit. Second, it modifies the price impact in both subgames. We now analyze how these distortions reshape the condition for absolute information acquisition. The condition $\mathbb{E}\pi_B \geq \mathbb{E}\hat{\pi}$ is equivalent to:

$$\frac{\lambda}{\hat{\lambda}} - 1 \leq \frac{9q}{4[1+\kappa(1-\theta_0)^2]}. \tag{20}$$

[8]Noise trading is irrelevant for this threshold because it scales the price impact in both subgames proportionally (by $1/\sigma_z$), leaving the ratio $\lambda/\hat{\lambda}$ unchanged. In extensions with risk-averse informed traders, however, noise trading does influence the acquisition decision.

B chooses absolute information acquisition only when the proportional increase in price impact (LHS) is not excessive relative to the potential gains (RHS). We refer to the effect of price-driven trading on the RHS of Equation (20) as the *exploitation effect*. The term $\kappa(1-\theta_0)^2$ in the denominator reflects the opportunity cost of acquiring information: as mispricing grows (whether underpricing or overpricing), the potential profits from price-driven trading in the deep market without information acquisition increase. The bound therefore tightens, making information acquisition less attractive. This captures the trade-off between exploration (acquiring extra information) and exploitation (trading on mispricing).

We refer to the effect of mispricing on the proportional increase in price impact (the LHS of Equation (20)) as the *price impact effect*. Direct calculation of the price impact ratio yields the following lemma:

LEMMA 4.1 ***(Price Impact Effect)*** *The ratio of price impact in the two subgames,*

$$\frac{\lambda}{\hat{\lambda}} = \sqrt{1 + \frac{3q}{4} + \frac{3q}{8\theta_0}}, \tag{21}$$

is strictly greater than $\frac{\sqrt{4+3q}}{2}$ *and is strictly decreasing in* θ_0.

Unlike the exploitation effect, which is symmetric in the magnitude of mispricing, the price impact effect is asymmetric. Relative to the efficient market, initial underpricing amplifies the proportional increase in price impact, whereas overpricing dampens it. Under large initial underpricing, the market is highly liquid with minimal baseline price impact, so acquiring information causes a disproportionately large relative increase in price impact, generating a strong price impact effect that discourages acquisition.

The final decision depends on the interaction of the two effects. Under initial overpricing, they oppose each other: the exploitation effect discourages acquisition, while the price impact effect encourages it. For sufficiently large overpricing ($\theta_0 > \bar{\theta}(q,\kappa)$), the exploitation effect dominates and acquisition ceases.[9] Under initial underpricing, both effects discourage acquisition. For sufficiently large underpricing ($\theta_0 < \underline{\theta}(q,\kappa)$), the price impact effect becomes overwhelming,[10] leaving the post-acquisition market

[9] As $\theta_0 \to +\infty$, the RHS of Equation (20) approaches 0, forcing the inequality to fail.

[10] As $\theta_0 \to 0$, the LHS of Equation (20) approaches $+\infty$.

too illiquid to justify the switch. The two sides of the acquisition region differ in robustness. Deterrence under overpricing follows from the exploitation effect alone and survives every variation of the competitive structure we consider. Deterrence under underpricing relies on the divergence of the price impact ratio as $\theta_0 \to 0$, which in turn relies on B and L sharing identical common information; Section 7.1 shows that with imperfectly correlated signals the ratio stays bounded and the lower threshold exists only when the exploitation effect is strong. Table 1 summarizes how mispricing alters the incentives for information acquisition by decomposing the net effect into the exploitation channel and the price impact channel.

Table 1: Comparison of Incentives for Information Acquisition (Relative to Efficient Market)

	Underpricing ($\theta_0 < 1$)	**Overpricing** ($\theta_0 > 1$)
Exploitation Effect	**Deters** (forgone exploitation profits rise)	**Deters** (forgone exploitation profits rise)
Price Impact Effect	**Deters** (proportional rise in price impact is amplified)	**Encourages** (proportional rise in price impact is dampened)

The exploitation effect persists whenever $\bar{v} > 0$, but the price impact effect is highly sensitive to the form of competition in private information. In Section 7.1, where we allow for more general forms of competition, we find that less intense competition weakens the price impact effect for small θ_0. The absence of information acquisition under large initial underpricing therefore occurs only when the exploitation effect is strong enough to compensate for the diminished price impact effect: that is, when the potential to exploit mispricing is substantial relative to the quality of extra information.

Under perfect competition in common information (the baseline), the price impact effect is strongest: information acquisition is deterred even absent the exploitation channel. When $\theta_0 < \frac{2}{20+27q}$, acquiring information reduces B's information-driven trading profits, as Panel (a) of Figure 2 illustrates.

COROLLARY 4.1 ***(Information-driven Trading Profits)*** *Given $q > 0$, when B acquires extra information, relative to the symmetric subgame:*

(i) if $\theta_0 \geq \frac{2}{20+27q}$, B's information-driven trading profits increase; conversely, if $0 < \theta_0 < \frac{2}{20+27q}$, these profits strictly decrease;

(ii) L's information-driven trading profits strictly decrease.

4.2 Properties of the Acquisition Thresholds

The width of $[\underline{\theta}(q,\kappa), \bar{\theta}(q,\kappa)]$ determines the range of market inefficiency under which information acquisition is sustained. This interval depends solely on the relative value of exploring extra information (q) and exploiting mispricing (κ). We characterize its monotonicity and limiting behavior.

PROPOSITION 4.3 ***(Properties of $\underline{\theta}(q,\kappa)$ and $\bar{\theta}(q,\kappa)$)*** *For $q > 0$ and $\kappa > 0$, the following statements hold:*

(i) $\underline{\theta}(q,\kappa)$ is strictly decreasing in q, while $\bar{\theta}(q,\kappa)$ is strictly increasing in q.

(ii) $\lim_{q\to+\infty} \underline{\theta}(q,\kappa) = 0$ and $\lim_{q\to+\infty} \bar{\theta}(q,\kappa) = +\infty$. Furthermore, $\lim_{q\to 0^+} \underline{\theta}(q,\kappa) = \underline{\theta}(\kappa)$ and $\lim_{q\to 0^+} \bar{\theta}(q,\kappa) = \bar{\theta}(\kappa)$, where $\underline{\theta}(\kappa)$ and $\bar{\theta}(\kappa)$ are the two positive roots of the cubic equation:

$$2\kappa x^3 - 3\kappa x^2 - 10x + 1 + \kappa = 0.$$

(iii) $\underline{\theta}(q,\kappa)$ is strictly increasing in κ, while $\bar{\theta}(q,\kappa)$ is strictly decreasing in κ.

(iv) $\lim_{\kappa\to+\infty} \underline{\theta}(q,\kappa) = \lim_{\kappa\to+\infty} \bar{\theta}(q,\kappa) = 1$. Furthermore, $\lim_{\kappa\to 0^+} \underline{\theta}(q,\kappa) = \frac{2}{20+27q}$ and $\lim_{\kappa\to 0^+} \bar{\theta}(q,\kappa) = +\infty$.

Proposition 4.3 establishes the monotonicity of $\underline{\theta}$ and $\bar{\theta}$ in q and κ, as Panels (a) and (b) of Figure 3 depict. Higher information quality (higher q) makes acquisition more profitable and widens the acquisition region $[\underline{\theta}(q,\kappa), \bar{\theta}(q,\kappa)]$. A larger mispricing rent (higher κ) raises the relative attractiveness of exploiting mispricing without acquiring extra information and narrows the interval.

By monotonicity, the limiting cases define the maximal bounds of the interval $[\underline{\theta}(q,\kappa), \bar{\theta}(q,\kappa)]$. First, we consider the limits with respect to q. As $q \to +\infty$, trading

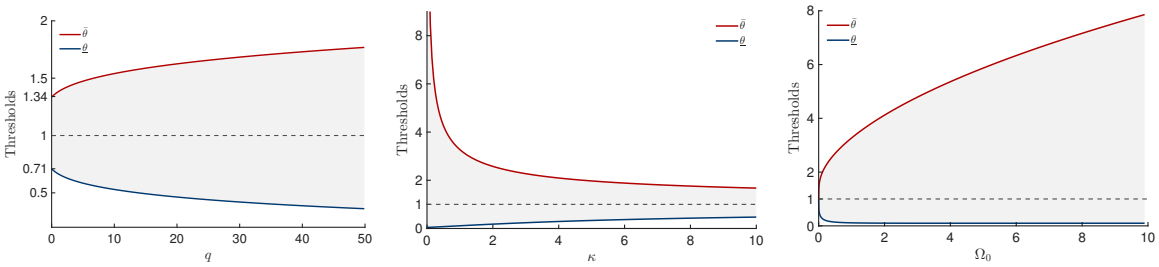

(a) Relation with q (b) Relation with κ (c) Relation with Ω_0

Figure 3: Comparative statics of the thresholds $\underline{\theta}(q,\kappa)$ (blue) and $\bar{\theta}(q,\kappa)$ (red). Panel (a) plots the thresholds against q (with $\kappa = 30$). Panel (b) plots the thresholds against κ (with $q = 1$). Panel (c) plots the thresholds against Ω_0 (with $\Lambda_0 = \bar{v} = 1$).

on extra information becomes arbitrarily valuable to B. Regardless of the magnitude of initial mispricing, the potential losses from common-information-driven and price-driven trading are fully offset by the gains from extra information. Absolute information acquisition is therefore always optimal.

The case where $q \to 0$ is more nuanced. If extra information is entirely absent ($q = 0$), the interval effectively expands to the entire real line. However, the limit as $q \to 0^+$ converges to a finite interval $[\underline{\theta}(\kappa), \bar{\theta}(\kappa)]$. This discontinuity at $q = 0$ arises because a total lack of information content eliminates both the exploitation effect and the price impact effect simultaneously. Panel (a) of Figure 3 visualizes this jump: the thresholds do not converge to 1, but rather define a distinct interval (approximately $(0.71, 1.34)$ for $\kappa = 30$) as $q \to 0$. By monotonicity, the absolute information acquisition is sustained for any $\theta_0 \in [\underline{\theta}(\kappa), \bar{\theta}(\kappa)]$, independent of q.

Second, we examine the limits with respect to κ. As $\kappa \to +\infty$, the cost of price impact becomes prohibitive. Whatever the quality of extra information, any deviation from efficiency deters B from acquiring, because B cannot sustain the unbounded losses from price-driven trading under information acquisition. Panel (b) of Figure 3 shows both thresholds converging to 1 as $\kappa \to +\infty$. As $\bar{v} \to 0$ (so $\kappa \to 0$), informed traders profit only from information-driven trading, and the limiting values of $\underline{\theta}(q,\kappa)$ and $\bar{\theta}(q,\kappa)$ recover Corollary 4.1.

Finally, we analyze the role of common information by taking the limit $\Omega_0 \to 0$. As common information vanishes,[11] both q and κ approach $+\infty$. In this scenario, both

[11] We analyze the limit rather than setting $\Omega_0 = 0$ directly because the market maker cannot infer information from total order flow in the symmetric subgame if $\Omega_0 = 0$, rendering the price impact undefined.

extra-information-driven trading and price-driven trading become dominant considerations for B. Proposition 4.4 demonstrates that the interval $[\underline{\theta}(q,\kappa), \bar{\theta}(q,\kappa)]$ collapses to the singleton $\{1\}$, as depicted by the intersection in Panel (c) of Figure 3. Whenever the market is inefficient, a *zero-information economy* emerges in which neither extra information nor common information is incorporated into prices. While Nezafat and Schroder (2023) also document a zero-information economy in a strategic trading framework, their mechanism relies on risk aversion and substantial noise trading within an efficient market. In contrast, our result arises from mispricing and persists even with risk-neutral participants.

PROPOSITION 4.4 ***(Vanishing Common Information)*** *Given $\bar{v} > 0$, ceteris paribus, as $\Omega_0 \to 0^+$, B does not choose absolute information acquisition for any $\theta_0 \neq 1$.*

The intuition behind Proposition 4.4 lies in the dominance of market liquidity over the value of information. As the variance of common information, Ω_0, approaches zero, the market becomes infinitely liquid in the absence of information acquisition. While acquiring extra information yields only finite profits, it compels B to forgo the potentially unbounded profits obtainable from exploiting initial mispricing in a frictionless (infinitely liquid) environment without acquiring information. In an efficient market, by contrast, such forgone gains from exploiting mispricing do not arise, and the losses associated with trading on common information vanish as well, making B choose absolute information acquisition.

The intuition inverts the standard logic: when consensus (common information) is scarcest, new private information is in principle most valuable, yet trading on it forces B to correct mispricing in an infinitely liquid market, making the price impact cost prohibitive. A minimal level of common information is therefore a prerequisite for sustaining incentives for information acquisition when prices deviate from efficiency.

4.3 Gains from Information Acquisition

Although severe mispricing deters information acquisition, B can still gain substantially from acquiring when mispricing is moderate. We ask whether these gains are maximized in a perfectly efficient market. As the red line in Panel (a) of Figure 2 shows, the net benefit to B is maximized at some $\theta^* < 1$. The following proposition formalizes this.

PROPOSITION 4.5 *Given $q > 0$, the following statements hold:*

(i) There exists a unique $\theta^*(q, \kappa) \in (0, 1)$ that maximizes $\mathbb{E}\pi_B - \mathbb{E}\hat{\pi}$.

(ii) $\mathbb{E}\pi_B - \mathbb{E}\hat{\pi}$ is strictly decreasing in θ_0 for all $\theta_0 \geq 1$.

To understand the nonmonotonicity of $\mathbb{E}\pi_B - \mathbb{E}\hat{\pi}$, we decompose it into the change in information-driven trading profits and the change in price-driven trading profits.[12]

The change in price-driven trading profits is governed by the exploitation effect. Because acquiring information reveals the true liquidation value to the market maker more quickly via order flow, it erodes B's ability to exploit the initial mispricing. This component is therefore negative and decreases as the magnitude of initial mispricing (whether underpricing or overpricing) grows; it is maximized at zero only in an efficient market.

By contrast, the change in information-driven trading profits behaves asymmetrically. It is maximized under initial underpricing. Although underpricing increases the proportional rise in price impact associated with information acquisition (raising the cost of trading), it also induces a deeper market in absolute terms. With moderate initial underpricing, the effect of increased market depth dominates: the absolute increase in expected profits from trading on private information outweighs the opportunity cost of lost exploitation profits.

Initial overpricing unambiguously reduces the incentives for information acquisition. Moderate initial underpricing, by contrast, increases information-driven profits while imposing only limited exploitation losses, so the net gain from acquisition can exceed the efficient-market benchmark.

This analysis offers a theoretical foundation for the underpricing of initial public offerings (Beatty and Ritter, 1986; Ritter and Welch, 2002). Viewing initial mispricing as a choice variable for issuers setting the offer price, our results identify a rational motive for underpricing: it incentivizes large investors to acquire information and thereby raises price informativeness.[13] Figure 4 shows that the optimal underpricing level $\theta^*(q, \kappa)$ decreases in q and increases in κ. When the quality of private information (q) is lower or the mispricing rent (κ) is higher, issuers should set the offer price closer to fundamental value to motivate information acquisition.

[12] A detailed proof regarding this monotonicity and decomposition is provided in the Supplemental Appendix, Section B.

[13] As shown in Equation (17), initial underpricing increases equilibrium price informativeness.

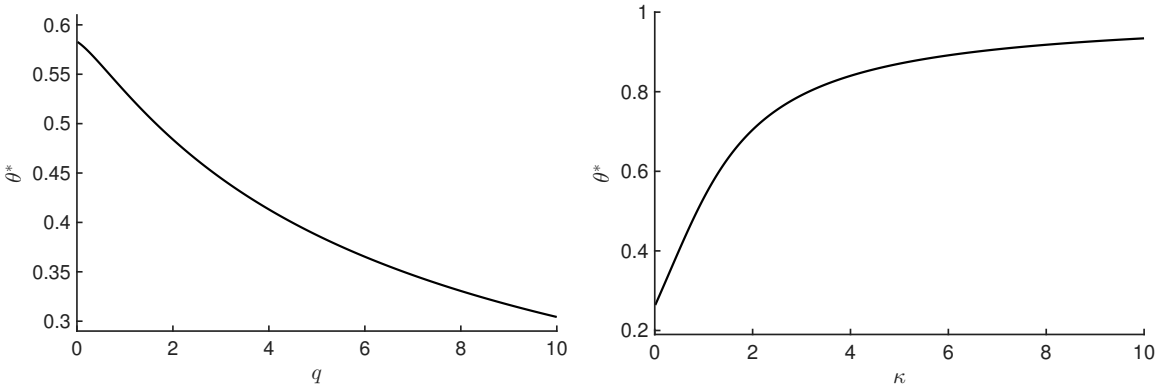

(a) Relation with q (b) Relation with κ

Figure 4: Comparative statics of the profits-maximizing level $\theta^*(q, \kappa)$. Panel (a) plots $\theta^*(q, \kappa)$ against q (with $\kappa = 1$). Panel (b) plots $\theta^*(q, \kappa)$ against κ (with $q = 1$).

5 Multiple Trading Periods

We extend the one-shot model to a dynamic setting with N trading periods. The recursive linear equilibrium has no closed form. The Supplemental Appendix (Propositions C.1 and C.2) gives necessary and sufficient recursive conditions for it, and we compute it numerically by backward induction from terminal values. For $N = 2$ the equilibrium is solved analytically in Appendix C.5: the symmetric subgame has a closed form, and in the asymmetric subgame we prove that L trades the mispricing more aggressively than B in the first round and obtain the difference between B's and L's ex ante profits in closed form (Propositions C.3 and C.4); existence of the asymmetric equilibrium is proved at the benchmark $\Omega_0 = \Lambda_0$, and uniqueness is numerical. For general N, existence and uniqueness are not established, and every statement in this section describes the numerically computed equilibrium under the benchmark parameters of Appendix C.4 unless stated otherwise. In the one-shot model, B always outperforms L when information is acquired, so even a trader motivated solely by relative performance would choose to acquire information. In the dynamic setting this can fail: when initial mispricing is sufficiently large, L outperforms B by exploiting mispricing more aggressively, making information acquisition unattractive even on relative grounds.

5.1 Dynamic Setting

There are N trading periods, with the n-th round of trading taking place at date n for $n = 1, 2, \cdots, N$. The asset has date-$(N+1)$ liquidation value v, as described in Equation (1). Before trading, common information f is revealed to B and L, and extra information e is revealed to B if B costlessly acquires it at date 0.

In the n-th round, the market maker sets the price p_n after observing total order flow $y_n = x_{B,n} + x_{L,n} + z_n$, where $x_{B,n}$ and $x_{L,n}$ are the orders submitted by B and L, and $z_n \sim N(0, \sigma_z^2)$ is the random market order from noise traders. $\{z_n\}_{n=1}^N$ is mutually independent and independent of all other random variables.

The positive mispricing parameter θ_n measures proportional deviation of the equilibrium price p_n from the market maker's rational Bayesian conditional expectation, that is

$$p_n = \theta_n \mathbb{E}\left[v \;\middle|\; y^{(n)} \right], \quad \text{where } \; y^{(n)} = (y_1, \cdots, y_n). \tag{22}$$

Given an exogenous $\theta_0 > 0$ with initial price $p_0 = \theta_0 \bar{v}$, the evolution of θ_n is endogenously determined. Before the n-th round begins, B and L submit market orders $x_{B,n}$ and $x_{L,n}$ to maximize their expected profits from date n to N conditional on their information sets. The profits from date n to N are

$$\pi_{i,n} = \sum_{k=n}^{N} (v - p_k) x_{i,k}, \quad i \in \{B, L\}.$$

5.2 Forecasting Issues

In the symmetric subgame without information acquisition, we use the hat $\hat{}$ sign to mark parameters and omit subscripts due to ex ante homogeneity of B and L. Before the n-th round begins, each informed trader's information advantage relative to the market maker is

$$s_{n-1} = \mathbb{E}\left[v \mid f \right] - \mathbb{E}\left[v \;\middle|\; \hat{y}^{(n)} \right] = \bar{v} + f - \frac{\hat{p}_{n-1}}{\hat{\theta}_{n-1}}.$$

Informed traders trade on their information advantage and on the previous price.

We search for recursive linear equilibrium of the form[14]:

$$\begin{cases} \hat{x}_n = \hat{\gamma}_n s_{n-1} + \hat{\xi}_n \hat{p}_{n-1}, \\ \hat{p}_n = \hat{p}_{n-1} + \hat{\lambda}_n \hat{y}_n, \end{cases} \quad n = 1, 2, \cdots, N. \tag{23}$$

Following Foster and Viswanathan (1996), the remaining information after the n-th round is measured by

$$\begin{aligned} \hat{\Omega}_n &= \mathrm{Var}\left(f \;\middle|\; \hat{y}^{(n)} \right) = \mathrm{Var}\left(s_n \right), \\ \hat{\Sigma}_n &= \mathrm{Var}\left(v \;\middle|\; \hat{y}^{(n)} \right) = \Lambda_0 + \hat{\Omega}_n. \end{aligned} \tag{24}$$

In the asymmetric subgame with information acquisition, the forecasting problem becomes more complex than in the one-shot setting. After the n-th round, L forms a sharper forecast v_n by observing the residual order flow $u_n = y_n - x_{L,n}$, that is,

$$v_{n-1} = \mathbb{E}\left[v \;\middle|\; f, u^{(n-1)} \right], \quad \text{where} \;\; u^{(n-1)} = (u_1, \cdots, u_{n-1}).$$

Before the n-th round begins, L's information advantage relative to the market maker ($s_{L,n-1}$) and B's information advantage relative to L ($s_{B,n-1}$) are, respectively,

$$s_{L,n-1} = v_{n-1} - \mathbb{E}\left[v \;\middle|\; y^{(n)} \right] = v_{n-1} - \frac{p_{n-1}}{\theta_{n-1}} \quad \text{and} \quad s_{B,n-1} = v - v_{n-1}. \tag{25}$$

The recursive linear equilibrium in the asymmetric subgame takes the form

$$\begin{cases} x_{B,n} = \beta_n s_{B,n-1} + \gamma_n s_{L,n-1} + \xi_n p_{n-1}, \\ x_{L,n} = \psi_n s_{L,n-1} + \zeta_n p_{n-1}, \\ p_n = p_{n-1} + \lambda_n y_n, \end{cases} \tag{26}$$

and when B and L submit equilibrium orders, v_n evolves as

$$v_n = v_{n-1} + \phi_n \left(y_n - \mathbb{E}[y_n \mid \mathcal{F}^L_{n-1}] \right), \tag{27}$$

where ϕ_n is L's inference intensity and $\mathcal{F}^L_{n-1}$ is L's information set before the n-

[14]Including price-driven trading explicitly in the order strategy serves to provide economic intuition. Equivalently, one can reformulate the strategy by introducing a time-varying intercept and show that it is proportional to the previous price, with slope $\hat{\xi}_n$ independent of p_{n-1}.

th round of trading. This evolution equation indicates that the increment in v_n is proportional to the unexpected total order flow from L's perspective. Finally, the remaining information after the n-th round is measured by

$$
\begin{aligned}
\Omega_n &= \mathrm{Var}\left(f \;\middle|\; y^{(n)}\right) = \mathrm{Var}\left(s_{L,n}\right), \\
\Lambda_n &= \mathrm{Var}\left(e \;\middle|\; f, u^{(n)}\right) = \mathrm{Var}\left(s_{B,n}\right), \\
\Sigma_n &= \mathrm{Var}\left(v \;\middle|\; y^{(n)}\right) = \Lambda_n + \Omega_n.
\end{aligned} \tag{28}
$$

The Supplemental Appendix (Section C) shows the detailed derivation of the recursive linear equilibria in the symmetric and asymmetric subgames, the numerical method to solve them, and the benchmark parameter setting.

5.3 Price-driven Trading and Correction of Mispricing

The coupled relation between endogenous price-driven trading and the correction of mispricing in the dynamic setting mirrors the one-shot model, and is characterized by the following evolution equations of mispricing in the two subgames:

$$
\frac{\theta_n}{\theta_{n-1}} - 1 = \lambda_n(\xi_n + \zeta_n) \quad \text{and} \quad \frac{\hat{\theta}_n}{\hat{\theta}_{n-1}} - 1 = 2\hat{\lambda}_n\hat{\xi}_n. \tag{29}
$$

Equation (29) shows that the non-zero term $\frac{\theta_n}{\theta_{n-1}} - 1$ creates the price-driven trading opportunity. Price-driven trading with positive (negative) intensity under underpricing (overpricing) moves θ_n toward one, making the market more efficient. The amount of mispricing correction is governed by price impact and total price-driven trading intensity, so each trader's contribution to correction is captured by that trader's price-driven intensity.

Figure 5 reveals richer dynamics than the one-shot model. As before, information acquisition increases price impact and induces more conservative price-driven trading. But L's sharper forecasts from residual order flow generate new patterns.

First, L adopts a more aggressive price-driven trading strategy than B in all but the last trading round,[15] and therefore earns higher price-driven trading profits. In the dynamic setting, informed traders weigh profits in both current and future periods.

[15] L also trades more aggressively on common information, consistent with Foster and Viswanathan (1994); the replication package reports the information-driven intensities.

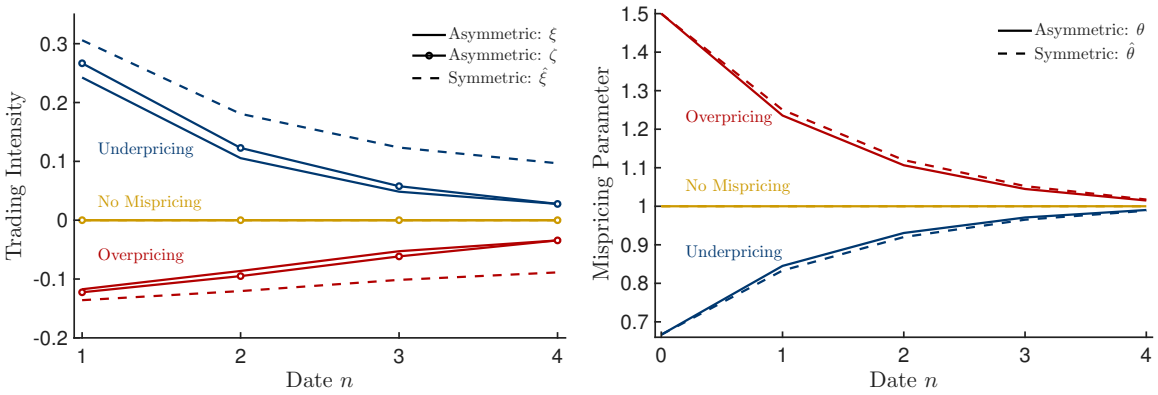

(a) Price-driven trading intensity

(b) Correction of mispricing

Figure 5: This figure shows the evolution of price-driven trading intensity (Panel (a)) and mispricing parameter (Panel (b)). In Panel (a), solid lines without open circles (ξ) and with open circles (ζ) represent trading intensities in the asymmetric subgame; dashed lines ($\hat{\xi}$) represent those in the symmetric subgame. In Panel (b), the solid lines (θ) and dashed lines ($\hat{\theta}$) represent mispricing parameters in the asymmetric and symmetric subgames, respectively. Colors indicate initial mispricing: blue for underpricing, yellow for no mispricing, and red for overpricing.

Extra information gives B additional value through cross-terms involving common and extra information in later-period conditional expected profits. L compensates for the absence of this value by trading the mispricing more aggressively. The one-shot model is the last trading round, where only current profits matter and B and L share the same price-driven trading intensity. This aggressive price-driven trading by L is key to understanding why, in the dynamic setting, L can outperform B when B acquires extra information under large initial mispricing.

Second, competition to exploit mispricing is most intense at the start of trading, when mispricing is largest; as mispricing is corrected, price-driven trading intensities decline. Third, θ_n lies closer to one than $\hat{\theta}_n$: through L's sharper forecast, the higher price informativeness after information acquisition accelerates the correction of mispricing.

5.4 Ex Ante Profits

By standard dynamic programming, the necessary and sufficient conditions for the recursive linear equilibrium are given in the Supplemental Appendix (Propositions C.1 and C.2). Ex ante profits of B and L from date $n+1$ to date N in the symmetric

and asymmetric subgames are

$$\mathbb{E}\hat{\pi}_{n+1} = \underbrace{\hat{A}_n\hat{\Omega}_n}_{\text{Information-driven}} + \underbrace{\left(\bar{v}^2 + \Sigma_0 - \hat{\Sigma}_n\right)\left[\sum_{k=n+1}^{N} \hat{\xi}_k\hat{\theta}_{k-1}\left(1-\hat{\theta}_k\right)\right]}_{\text{Price-driven}} + \hat{\Delta}_n, \tag{30}$$

$$\mathbb{E}\pi_{B,n+1} = \underbrace{A_n\Lambda_n + B_n\Omega_n}_{\text{Information-driven}} + \underbrace{\left(\bar{v}^2 + \Sigma_0 - \Sigma_n\right)\left[\sum_{k=n+1}^{N} \xi_k\theta_{k-1}(1-\theta_k)\right]}_{\text{Price-driven}} + \Delta_n,$$

$$\mathbb{E}\pi_{L,n+1} = \underbrace{G_n\Omega_n}_{\text{Information-driven}} + \underbrace{\left(\bar{v}^2 + \Sigma_0 - \Sigma_n\right)\left[\sum_{k=n+1}^{N} \zeta_k\theta_{k-1}(1-\theta_k)\right]}_{\text{Price-driven}} + \delta_n. \tag{31}$$

Here $\hat{A}_n$ and $\hat{\Delta}_n$ are given in the Supplemental Appendix (Equation C.9), while A_n, B_n, G_n, Δ_n and δ_n are given in Equations C.18 and C.20. In both subgames, ex ante profits decompose into three parts: information-driven trading profits, price-driven trading profits and a residual time-varying term. Profits from trading on common information or on extra information are proportional to the respective remaining information, Ω_n ($\hat{\Omega}_n$) or Λ_n. As private information is revealed through trading, the information advantage shrinks, yielding lower profits from trading on that information.

Positive profits from price-driven trading imply the existence of the exploitation effect in the dynamic setting and it can be further decomposed into two terms. The first term is the same for B and L, and equals the second moment of the hypothetical efficient price from the ex ante perspective:

$$\mathbb{E}\left[\left(\frac{\hat{p}_n}{\hat{\theta}_n}\right)^2\right] = \bar{v}^2 + \Sigma_0 - \hat{\Sigma}_n \quad \text{and} \quad \mathbb{E}\left[\left(\frac{p_n}{\theta_n}\right)^2\right] = \bar{v}^2 + \Sigma_0 - \Sigma_n.$$

Given θ_0, a larger second moment implies a greater absolute deviation of the actual price from the efficient price; hence greater mispricing rent. In the one-shot model, this second moment reduces to $\bar{v}^2$; in the dynamic setting, it also depends on how much remaining information about v is reduced over time.

The second term[16] is a weighted aggregation of price-driven trading intensities in

[16]In the asymmetric subgame, it is $\sum_{k=n+1}^{N} \xi_k\theta_{k-1}(1-\theta_k)$ for B and $\sum_{k=n+1}^{N} \zeta_k\theta_{k-1}(1-\theta_k)$ for

subsequent periods, with weight $\hat{\theta}_{k-1}(1-\hat{\theta}_k)$ in the symmetric subgame and weight $\theta_{k-1}(1-\theta_k)$ in the asymmetric subgame. More aggressive price-driven trading yields a larger contribution to correcting mispricing, and is rewarded with higher price-driven trading profits. Summing the second term for B and L yields

$$\sum_{k=n+1}^{N} \frac{1}{\hat{\lambda}_k}\big(1-\hat{\theta}_k\big)\big(\hat{\theta}_k-\hat{\theta}_{k-1}\big) \quad \text{and} \quad \sum_{k=n+1}^{N} \frac{1}{\lambda_k}(1-\theta_k)(\theta_k-\theta_{k-1}).$$

The sum of the second term is a weighted aggregation of market depth in subsequent periods, with weights given by the product of absolute deviation of $\hat{\theta}_k$ (θ_k) from one and the amount of mispricing correction. This product captures jointly how much mispricing is available to be exploited and the extent of mispricing correction attributable to informed traders. In the one-shot model, the amount of correction at date 1 is proportional to $1-\theta_0$, so B and L's ex ante profits reduce to Corollary 3.1.

5.5 Choice of Information Acquisition

To accommodate B's strategic objective of outperforming L, we introduce a second acquisition criterion on relative grounds. The criterion is relevant when B is evaluated against a peer, as under the relative performance evaluation that governs much delegated portfolio management, and it isolates whether the informational advantage that acquisition confers survives competition in exploiting mispricing.

DEFINITION 5.1 ***(Relative Information Acquisition)*** *B acquires extra information if B's ex ante profits over the entire trading period in the asymmetric subgame exceed those of L.*

The relative information acquisition is always chosen in the one-shot model, but not necessarily in the dynamic setting. By definition, absolute information acquisition is chosen if $\mathbb{E}\pi_{B,1} \geq \mathbb{E}\hat{\pi}_1$, whereas relative information acquisition is chosen if $\mathbb{E}\pi_{B,1} \geq \mathbb{E}\pi_{L,1}$.

The analysis of B's choice of information acquisition begins with the confirmation of the price impact effect in the dynamic setting. In Panel (a) of Figure 6, price impact decreases with date n in both subgames, and is larger with information acquisition than without. In Panel (b), the ratio $\lambda_n/\hat{\lambda}_n$ decreases with θ_0, consistent with

L, whereas in the symmetric subgame, it is $\sum_{k=n+1}^{N} \hat{\xi}_k \hat{\theta}_{k-1}(1-\hat{\theta}_k)$ for each informed trader.

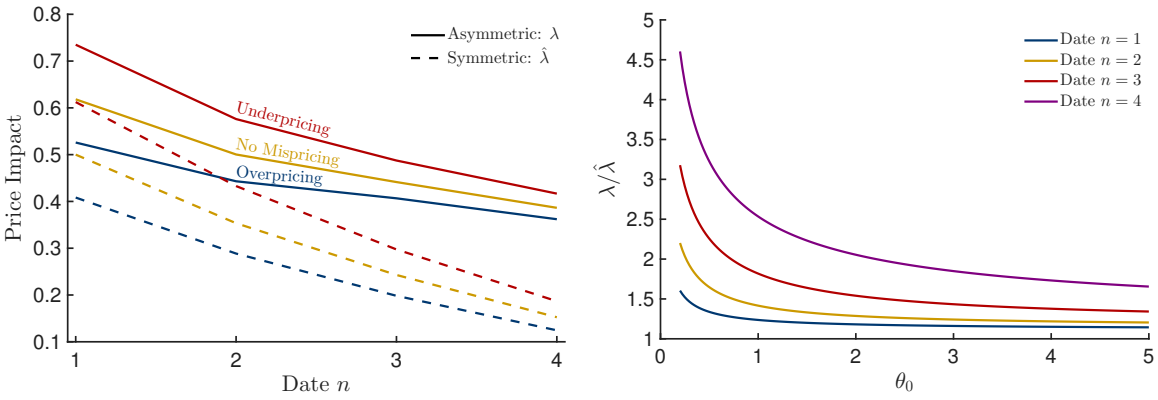

(a) Price impact

(b) Ratio of price impact $\lambda/\hat{\lambda}$

Figure 6: Panel (a) shows the evolution of price impact. Solid lines (λ) and dashed lines ($\hat{\lambda}$) represent price impact in the asymmetric and symmetric subgames, respectively. Colors indicate initial mispricing: blue for underpricing, yellow for no mispricing, and red for overpricing. Panel (b) shows $\lambda/\hat{\lambda}$ as a function of θ_0 at each date n.

Lemma 4.1. With large initial underpricing, the ratio approaches $+\infty$, indicating a very strong price impact effect for small θ_0.

Figure 7 illustrates the choice of information acquisition. In Panel (a), the absolute information acquisition is chosen when the red line lies above zero, and the relative information acquisition is chosen when the purple line lies above zero. When initial mispricing is not large ($\underline{\theta} \leq \theta_0 \leq \bar{\theta}$, the green shaded region), the red line lies above zero, so the absolute information acquisition is chosen, consistent with Proposition 4.2. With larger initial mispricing ($\underline{\theta}' \leq \theta_0 \leq \underline{\theta}$ or $\bar{\theta} \leq \theta_0 \leq \bar{\theta}'$, the blue shaded region), the red line lies below zero, while the purple line remains above zero, so the relative information acquisition is still chosen. With substantially large initial mispricing ($\theta_0 > \bar{\theta}'$ or $\theta_0 < \underline{\theta}'$, the red shaded region), both lines lie below zero, so even the relative information acquisition is deterred.

The deterrence of absolute information acquisition in the dynamic setting operates through the same channels as in the one-shot model: the exploitation effect dominates under large overpricing, and the price impact effect dominates under large underpricing. Relative information acquisition can fail for an additional reason. L earns higher price-driven profits by trading more aggressively, and under very large mispricing these profits exceed B's gains from trading on private information, so L outperforms B.

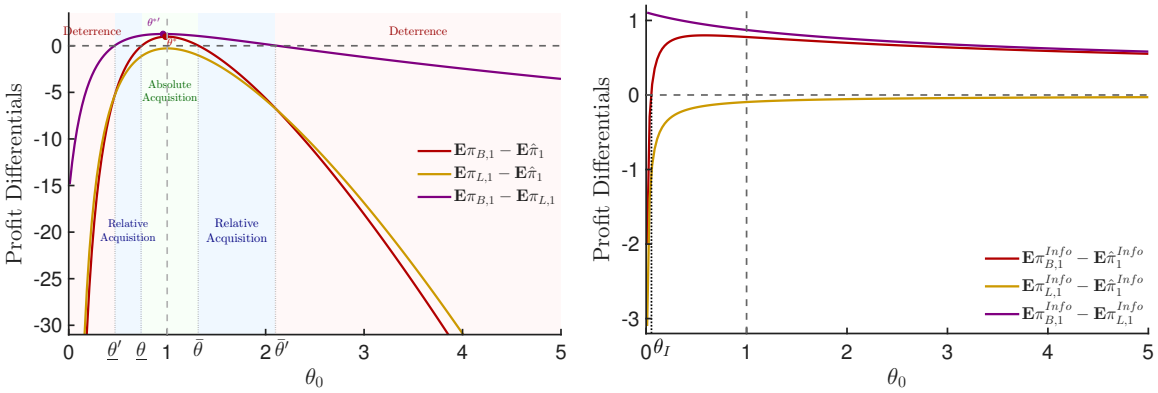

(a) Ex Ante Profits

(b) Information-driven Trading Profits

Figure 7: This figure shows ex ante profit differentials over the entire trading period (Panel (a)) and their information-driven component (Panel (b)) as a function of θ_0. In Panel (a), the red and yellow lines represent profit changes resulting from information acquisition for informed traders B and L, $\mathbb{E}\pi_{B,1} - \mathbb{E}\hat{\pi}_1$ and $\mathbb{E}\pi_{L,1} - \mathbb{E}\hat{\pi}_1$, respectively; the purple line represents the B–L profit difference, $\mathbb{E}\pi_{B,1} - \mathbb{E}\pi_{L,1}$. $\underline{\theta}$ and $\bar{\theta}$ denote the zero-crossing points of the red line and θ^* its maximizer; $\underline{\theta}'$ and $\bar{\theta}'$ denote the zero-crossing points of the purple line and $\theta^{*\prime}$ its maximizer. In Panel (b), the curves show information-driven profit differentials. The red, yellow, and purple lines represent $\mathbb{E}\pi^{\text{Info}}_{B,1} - \mathbb{E}\hat{\pi}^{\text{Info}}_1$, $\mathbb{E}\pi^{\text{Info}}_{L,1} - \mathbb{E}\hat{\pi}^{\text{Info}}_1$, and $\mathbb{E}\pi^{\text{Info}}_{B,1} - \mathbb{E}\pi^{\text{Info}}_{L,1}$, respectively. Here $\mathbb{E}\pi^{\text{Info}}_{B,1} = A_1\Lambda_1 + B_1\Omega_1$, $\mathbb{E}\pi^{\text{Info}}_{L,1} = G_1\Omega_1$, and $\mathbb{E}\hat{\pi}^{\text{Info}}_1 = \hat{A}_1\Omega_1$. θ^I denotes the zero-crossing point of the red line. Parameters: $\bar{v} = 10$ in both panels.

Panel (b) of Figure 7 confirms Corollary 4.1 in the dynamic setting: under very large initial underpricing ($\theta_0 < \theta^I$), B's information-driven trading profits fall with information acquisition. B's information-driven trading profits nonetheless always exceed L's (the purple line lies above zero), so L's aggressive price-driven trading is the primary reason large mispricing deters relative information acquisition.

Finally, when information acquisition is favorable to B, we consider the initial mispricing at which B gains most from information acquisition. The peaks of the blue and red lines in Panel (a) both occur under initial underpricing, implying that moderate initial underpricing can motivate both absolute and relative information acquisition.

6 Application: IPO Underpricing and Information Production

We now apply the model to initial public offerings, in which issuers deliberately set offer prices below expected market value. The framework delivers sharp predictions about how underpricing shapes pre-market information production, and these predictions distinguish it from existing theories of IPO underpricing.

6.1 Model Mapping

In the IPO setting, the initial mispricing parameter $\theta_0 < 1$ measures the degree of deliberate underpricing chosen by the issuer, with a lower θ_0 corresponding to more aggressive underpricing. Institutional investors play the role of the informed traders: they observe the offer price, form beliefs about fundamental value from public information (the common information f), and decide whether to expend resources acquiring private information about the issuing firm (the extra information e). The bookrunner and the secondary-market pricing mechanism play the role of the market maker, and the liquidation value v is the post-IPO fundamental value revealed as trading unfolds.

The mapping is natural because the IPO is the canonical setting in which mispricing is set by design. The issuer chooses the offer price, so underpricing is a deliberate decision variable rather than an equilibrium byproduct of trading frictions. The model therefore isolates how this decision reshapes the informational environment around the offering.

6.2 Testable Predictions

The results of Sections 4 and 5 translate into three predictions that separate our theory from existing explanations. We state them as predictions rather than propositions because they combine proved results with properties established numerically; we indicate the status of each.

PREDICTION 6.1 (Nonmonotone Information Production) *The relationship between IPO underpricing and pre-market information production is inverse-U-shaped:*

(i) moderate underpricing (θ_0 slightly below 1) stimulates information acquisition relative to the efficient-price benchmark $\theta_0 = 1$;

(ii) extreme underpricing ($\theta_0 < \underline{\theta}(q, \kappa)$) deters information acquisition entirely;

(iii) an interior underpricing level $\theta^*(q, \kappa) < 1$ maximizes the gains from information acquisition.

All three parts are proved: (i) and (iii) by Proposition 4.5 and (ii) by Proposition 4.2; Figure 4 illustrates them. The inverse-U shape sets our theory apart from Rock (1986), in which underpricing compensates uninformed investors for adverse selection, and from Benveniste and Spindt (1989), in which underpricing is the price of truthful revelation by investors who are already informed. Those theories concern the *revelation* of information that investors hold; ours concerns the *production* of information by investors who choose whether to become informed, and it is the production margin that turns nonmonotone.

PREDICTION 6.2 (Optimal Underpricing Depends on the Information Environment) *The underpricing level $\theta^*(q, \kappa)$ that maximizes information production satisfies:*

(i) θ^* is decreasing in $q = \Lambda_0/\Omega_0$: when acquirable private information is more valuable relative to common information, the information-maximizing offer price lies further below fundamental value;

(ii) θ^* is increasing in $\kappa = \bar{v}^2/\Omega_0$: when the mispricing rent is larger, the information-maximizing offer price lies closer to fundamental value.

Existence and uniqueness of $\theta^*(q, \kappa)$ are proved in Proposition 4.5; the two monotonicity properties are established numerically over a wide range of (q, κ) and depicted in Figure 4. They carry direct cross-sectional content. Offerings with greater scope for private discovery, such as opaque firms, novel business models, or thinly covered issuers for which q is high, call for more aggressive underpricing to stimulate information production. Offerings for which the exploitation motive is strong, such as issuers with a high prior valuation relative to fundamental uncertainty for which κ is high, call for offer prices nearer fundamental value, because aggressive underpricing would divert large investors toward exploiting the discount rather than acquiring information. These predictions have no counterpart in winner's-curse or information-revelation models.

PREDICTION 6.3 (Self-Reinforcing Extreme Underpricing) *When $\theta_0 < \underline{\theta}(q, \kappa)$, information production ceases; deterred acquisition slows the correction of mispricing and prolongs the environment in which exploitation dominates, generating a negative feedback loop between extreme underpricing and delayed price discovery.*

The first clause is Proposition 4.2. The rest draws on the numerically solved dynamics of Section 5, where θ_n converges to one faster in the asymmetric subgame than in the symmetric one (Figure 5). The prediction implies that IPOs with very large first-day returns exhibit lower post-IPO information production, through fewer analyst initiations and less institutional research, together with slower convergence to fundamental value in the secondary market. The prediction contrasts with "hot-market" explanations (Ljungqvist, 2007), in which large initial returns attract attention and accelerate price discovery.

6.3 Empirical Evidence

Several findings are consistent with these predictions. Anderson (2022) documents that underpricing induced by mutual-fund flows raises the consumption of accounting information, consistent with Proposition 6.1(i). Beatty and Ritter (1986) and Chemmanur (1993) show that IPO underpricing is associated with greater premarket information production by institutional investors, though neither examines the nonmonotone prediction. The cross-sectional variation in Proposition 6.2 is testable using standard IPO databases merged with analyst-coverage and institutional-trading records.

7 Extensions

We extend the baseline model in three directions and confirm the robustness of the main results. For tractability we work in the one-shot setting throughout.

7.1 Multiple Informed Traders and General Forms of Competition

We extend the model to $M \geq 1$ informed traders with correlated private signals, nesting the baseline as the special case $M = 2$ with perfectly correlated private signals.

There are $M \geq 1$ informed traders, each (indexed by $i \in \{1, 2, \cdots, M\}$) endowed with one piece of private signal f_i contained in the asset fundamental, so that the liquidation value decomposes as $v = \bar{v} + \sum_{i=1}^{M} f_i + e$. Here $\bar{v} \geq 0$ is the prior mean, and e is extra information. We assume that informed trader 1 is the unique better-informed trader who can costlessly acquire extra information at date 0, and the other informed traders ($i \in \{2, 3, \cdots, M\}$) observe only their own private signals f_i.

Each private signal has the same prior variance, $f_i \sim N(0, \Omega_0)$ for $i = 1, 2, \cdots, M$, and the covariance between any pair of private signals is the same, $\text{Cov}(f_i, f_j) = C_f$ for $i \neq j$. To avoid unnecessary discussion, we further assume that the private signals are uncorrelated or positively correlated, i.e., $C_f \geq 0$. Extra information follows $N(0, \Lambda_0)$ and is independent of each f_i. The total variance of v is then $\Sigma_0 = \Lambda_0 + M[\Omega_0 + (M-1)C_f]$. The random market order from noise traders is $z \sim N(0, \sigma_z^2)$, independent of all other random variables. Let $\eta_f = 1 + \frac{(M-1)C_f}{\Omega_0}$ and we can summarize the results in the following proposition.

PROPOSITION 7.1 *Given $M \geq 1$ informed traders, $q = \frac{\Lambda_0}{\Omega_0} > 0$ and $\kappa = \frac{\bar{v}^2}{\Omega_0} > 0$, the following statements hold:*

(i) In an efficient market, the absolute information acquisition is always chosen.

(ii) Under sufficiently large initial overpricing, the absolute information acquisition is not chosen.

(iii) If the following condition is satisfied

$$1 + \eta_f + \frac{\eta_f - 1}{2\left(\frac{M}{M+1} - \frac{\eta_f}{\eta_f+1}\right)} > \frac{M\eta_f^2}{\left(\frac{\eta_f}{\eta_f+1}\right)^2 + \frac{\kappa}{(M+1)^2}} + \frac{qM\eta_f^2}{8\left[\left(\frac{\eta_f}{\eta_f+1}\right)^2 + \frac{\kappa}{(M+1)^2}\right]^2}, \tag{32}$$

then under sufficiently large initial underpricing, the absolute information acquisition is not chosen; otherwise, the absolute information acquisition is always chosen under any initial underpricing. The last term on the left-hand side is read as zero when $\eta_f = 1$. When there is a single informed trader ($M = 1$), Equation (32) reduces to $\kappa > \sqrt{1+q}$.

Proposition 7.1 implies that, as in the baseline model, large initial overpricing deters absolute information acquisition through the exploitation effect. Whether large initial underpricing deters acquisition instead depends on Equation (32), which is

governed by the form of competition in private signals. Under perfect competition ($\eta_f = M$), as in the baseline, the price impact effect is strong enough that Equation (32) always holds, so acquisition is not chosen under large initial underpricing. Under less intense competition ($\eta_f < M$), the price impact effect weakens for small θ_0: as $\theta_0 \to 0$, the proportional increase in price impact with acquisition approaches a positive constant rather than diverging. Equation (32) then requires the mispricing rent to be large relative to the quality of extra information, that is, the exploitation effect must be strong. With a small $\bar{v}$, Panel (a) of Figure 8 shows that the change in B's profits from acquisition is negative only under large initial overpricing.

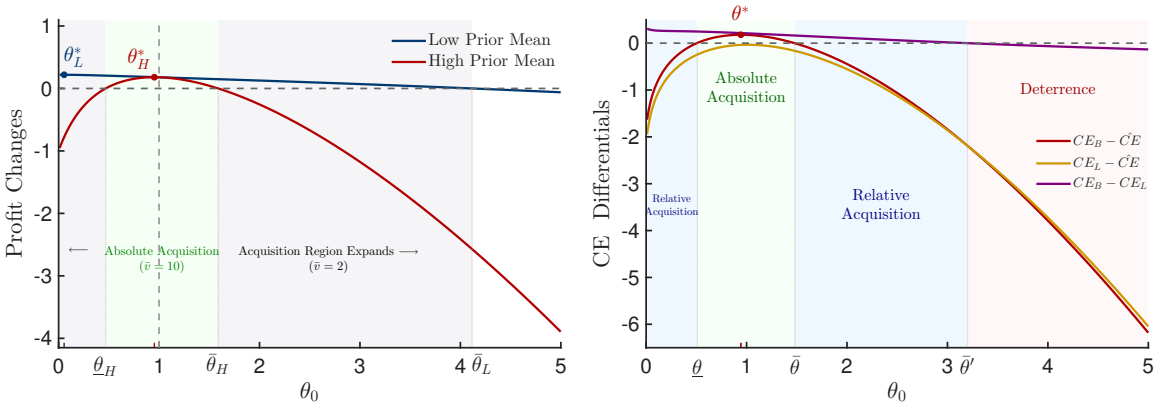

(a) Multiple Informed Traders

(b) Risk-averse Informed Traders

Figure 8: Panel (a) shows the change of ex ante profits resulting from information acquisition for better-informed trader 1, $\mathbb{E}\pi_1 - \mathbb{E}\hat{\pi}$, as a function of θ_0 under low (blue) and high (red) prior mean. $\underline{\theta}_H$ and $\bar{\theta}_H$ denote the zero-crossing points of the red line, while $\bar{\theta}_L$ denotes that of the blue line. Parameters: $\bar{v} = 2$ (blue), $\bar{v} = 10$ (red); $M = 3$, $C_f = 0.3$.
Panel (b) shows certainty-equivalent differentials for risk-averse informed traders as a function of θ_0. The red, yellow, and purple lines represent $CE_B - \hat{CE}$, $CE_L - \hat{CE}$, and $CE_B - CE_L$, respectively. $\underline{\theta}$ and $\bar{\theta}$ denote the zero-crossing points of the red line and θ^* its maximizer; $\bar{\theta}'$ denotes the zero-crossing point of the purple line. Parameters: $\bar{v} = 5$; $\rho_I = 1$.

7.2 Risk-averse Informed Traders

In this extension, we allow the informed traders to be risk-averse, but the market maker remains risk-neutral. Following Grossman and Stiglitz (1980) and Verrecchia (1982), B and L have the constant absolute risk aversion (CARA) utility over trading

profits π, with constant risk aversion coefficient $\rho_I > 0$,

$$U(\pi) = -e^{-\rho_I \pi}. \tag{33}$$

Panel (b) of Figure 8 provides the numerical results for the relationship between θ_0 and the certainty-equivalent differentials, where

$$CE_B = -\frac{\ln(-\mathbb{E}U(\pi_B))}{\rho_I} \text{ and } CE_L = -\frac{\ln(-\mathbb{E}U(\pi_L))}{\rho_I} \text{ and } \hat{CE} = -\frac{\ln(-\mathbb{E}U(\hat{\pi}))}{\rho_I}.$$

When ρ_I is not too large[17], the baseline results hold that large initial mispricing deters information acquisition, and moderate initial underpricing increases the gains from information acquisition. Under sufficiently large initial overpricing, B suffers a larger loss of expected utility than L with information acquisition, so the relative information acquisition is also not chosen. L adopts more aggressive price-driven trading than B under large initial overpricing, echoing the baseline model's multi-period setting with risk-neutral informed traders.

7.3 Algorithmic Trading: Exogenous Feedback Strategies

Following Zhang and Zhang (2025), we incorporate algorithmic traders through a specific form of algorithmic trading: feedback strategies based on the initial price p_0. The results reconcile the "striking conflict" documented by Weller (2018), in which algorithmic trading can simultaneously improve price informativeness with respect to existing information while deterring information acquisition.

To mitigate the discontinuities in price informativeness inherent in binary information acquisition, we generalize the framework to a smooth information acquisition setting, whereby B chooses the precision τ_e of a private signal at a cost. Rather than directly observing the extra information, B receives a costly signal $f_e = e + \varepsilon_e$ prior to trading, where $\varepsilon_e \sim N(0, \tau_e^{-1})$. B selects τ_e at date 0 to maximize its ex ante profit, subject to the cost $C(\tau_e)$. When $\tau_e = +\infty$, the extra information is fully revealed; conversely, when $\tau_e = 0$, B relies solely on common information. The cost function[18]

[17]With large ρ_I, the absolute information acquisition is not chosen under large initial overpricing but is always chosen under initial underpricing. This arises from the more complex price impact in the extension: with initial underpricing, B's trading intensities may increase with information acquisition even if price impact increases, so B need not incur losses from information acquisition.

[18]The cost function satisfies the regularity conditions imposed in Banerjee and Breon-Drish

is specified as $C(\tau_e) = C_e \tau_e^{3/2}$, with $C_e \geq 0$.

For feedback traders, we assume they submit exogenous market orders based on the initial price p_0 while taking the perceived price impact into account, such that

$$z_F = \frac{\xi_z p_0}{\lambda}.$$

Here, the constant $\xi_z > -1$ measures the feedback trading intensity.[19]

PROPOSITION 7.2 *Given $C_e = 0$, $\xi_z > -1$ and $\kappa > 0$, let $\underline{\theta}(q,\kappa)$ and $\bar{\theta}(q,\kappa)$ denote the thresholds in Proposition 4.2. Then the following statements hold:*

(i) If $\theta_0 \in \left[\frac{\underline{\theta}(q,\kappa)}{\xi_z+1}, \frac{\bar{\theta}(q,\kappa)}{\xi_z+1}\right]$, the optimal signal precision is $\tau_e^* = \infty$.

(ii) If $0 < \theta_0 < \frac{\underline{\theta}(q,\kappa)}{\xi_z+1}$ or $\theta_0 > \frac{\bar{\theta}(q,\kappa)}{\xi_z+1}$, the optimal signal precision is $\tau_e^* = 0$.

(iii) Given τ_e, if $\xi_z > 0$, market depth and price informativeness are lower than those in the absence of exogenous feedback traders; otherwise, they are higher.

When the signal is costless ($C_e = 0$), Proposition 7.2 shows that the smooth setting collapses to the binary case, where B's optimization over τ_e reduces to choosing between a perfectly precise signal ($\tau_e = \infty$) and an uninformative one ($\tau_e = 0$). In this setting, large initial mispricing continues to deter information acquisition.

The presence of feedback traders also aligns the acquisition region with that of the baseline model, adjusted by a scaling factor determined by the feedback intensity. Compared to the baseline acquisition region $[\underline{\theta}(q,\kappa), \bar{\theta}(q,\kappa)]$, exogenous feedback trading introduces two changes. First, the region is scaled by $1/(1+\xi_z)$: positive feedback trading narrows it, while negative feedback trading widens it. Second, the acquisition region shifts away from the efficient market, implying that an efficient market does not necessarily incentivize positive information production.

The third statement of Proposition 7.2 abstracts from the endogenous information acquisition channel by conditioning on a given signal precision and characterizes the effect of exogenous feedback trading on market depth and price informativeness with respect to existing information. For any given τ_e, positive (negative) feedback trading amplifies (attenuates) the effect of θ_0 on price impact, resulting in a shallower (deeper) market and lower (higher) price informativeness.

(2022): $C(\tau_e) \geq 0$, $C'(\tau_e) \geq 0$, $C''(\tau_e) > 0$, $C'(0) = 0$, and $\lim_{\tau_e \to \infty} C'(\tau_e) = \infty$.

[19] Requiring $\xi_z > -1$ ensures the existence of a linear equilibrium.

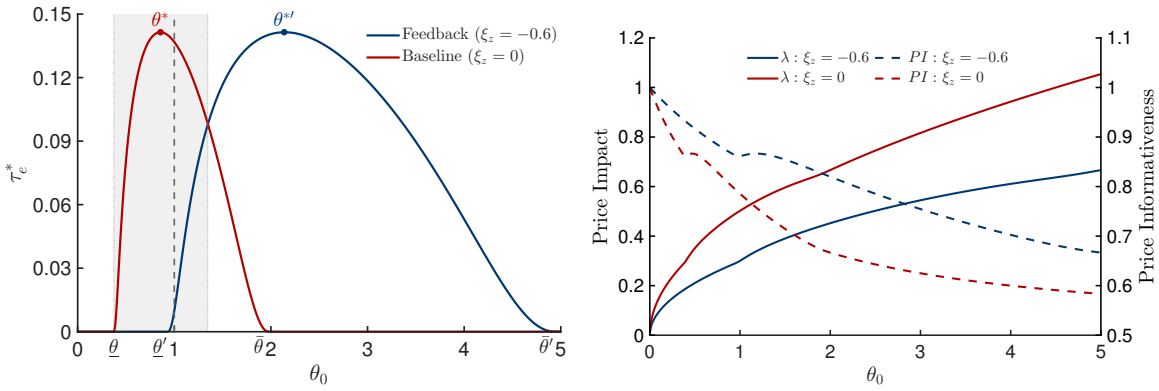

(a) Optimal signal precision

(b) Price impact and price informativeness

Figure 9: Panel (a) shows optimal signal precision τ_e^* as a function of θ_0 in the presence (blue, $\xi_z = -0.6$) and absence (red, $\xi_z = 0$) of exogenous feedback trading. $\underline{\theta}$ ($\underline{\theta}'$) and $\bar{\theta}$ ($\bar{\theta}'$) denote the largest zero-crossing point of the red (blue) line on $(0, \frac{1}{\xi_z+1})$ and the smallest on $(\frac{1}{\xi_z+1}, +\infty)$, respectively, and θ^* ($\theta^{*\prime}$) their maximizers. The shaded region indicates where the red line exceeds the blue line. Panel (b) shows price impact (solid lines, λ) and price informativeness (dashed lines, PI) as a function of θ_0 in the presence (blue) and absence (red) of exogenous feedback trading. Parameters: $\bar{v} = 2$; $C_e = 0.5$.

With costly signals ($C_e > 0$), Figure 9 confirms that large initial mispricing deters information acquisition, while the thresholds of the acquisition region depend on C_e. The maximal τ^* in the absence of exogenous feedback trading is attained under initial underpricing (θ^*), whereas the precision-maximizing level $\theta^{*\prime}$ in the presence of feedback trading depends on ξ_z and may exceed one.

Figure 9 further shows that negative feedback trading (blue lines) offers a plausible explanation for the conflict highlighted by Weller (2018). Negative feedback trading improves market liquidity and induces higher price informativeness with respect to existing information, while simultaneously reducing information acquisition over a wide region around the efficient market (the shaded region in Panel (a)). The aggregate effect on price informativeness hinges on the tension between two opposing forces: the enhanced informativeness with respect to *existing* information and the diminished informativeness with respect to *acquirable* information. In Panel (b), for $\xi_z = -0.6$, the former effect dominates, demonstrating that lower information acquisition and higher price informativeness resulting from algorithmic trading can indeed coexist in equilibrium.

8 Conclusion

Summary Mispricing and information acquisition are commonly studied as cause and effect running in one direction: informed trading reveals private information and corrects mispricing. We show that the arrow also runs in reverse. When an asset is already mispriced, the anticipated profits from exploiting the mispricing compete with the anticipated gains from acquiring new information, and large mispricing deters acquisition even at zero cost. The exploitation effect raises the opportunity cost of information, and it grows with the mispricing rent. The price impact effect is asymmetric: underpricing amplifies the proportional increase in price impact that accompanies acquisition, while overpricing attenuates it. Acquisition is therefore sustained only on a bounded interval of mispricing around the efficient price, whose thresholds depend on information quality and the mispricing rent but, with risk-neutral traders, not on noise trading volume. Within that interval the value of information is nonmonotone and peaks under moderate underpricing, which rationalizes IPO underpricing as an information-production device and delivers cross-sectional predictions for the optimal discount. In the numerically solved multi-period setting, a less-informed trader who exploits the mispricing more aggressively can outperform a better-informed rival, and a feedback loop emerges in which deterred acquisition slows correction and prolongs the environment in which exploitation dominates.

Several directions remain open. First, the pricing rule is a reduced form for an anchored pricing mechanism, and the analysis treats initial mispricing as exogenous. A microfoundation in which the wedge arises from an explicit friction or a biased anchor in the market maker's prior would tie θ_0 to observable primitives, and endogenizing θ_0 as the outcome of a prior-period equilibrium would allow the threat of future information acquisition to discipline the degree of mispricing that survives. The dynamic analysis is analytical only for two rounds; establishing existence and uniqueness of the recursive equilibrium for general N would place the reversal of the informational advantage on the same footing as the one-shot theorems. Second, in a multi-asset economy, the exploration–exploitation trade-off would interact with cross-asset information and correlation, potentially generating predictions about which assets attract more information production and which sustain persistent mispricing. Third, in practice, investors face uncertainty about whether and to what degree an asset is mispriced; incorporating learning about θ_0 would bring the theory

closer to the empirical setting in which mispricing must be estimated from composite indices (Stambaugh et al., 2015) rather than directly observed.

Implications The results bear on how one should read the information choices of sophisticated investors. Exploiting mispricing in actual markets is rarely mechanical: it requires assessing whether the mispricing is tradeable, how one's own trades move prices, and how persistent the underlying pressure is, and this trading-relevant knowledge differs from the predictive signals emphasized in standard information-acquisition models. Sophistication therefore need not translate into fundamental information production. When mispricing is large, the return to understanding the trading environment can exceed the return to an additional signal about fundamentals, so high institutional participation in a mispriced asset may reflect a greater capacity to trade the mispricing rather than stronger information production.

The same logic bears on the persistence of documented anomalies. Once an anomaly is known, the investors most capable of correcting it need not respond by producing more predictive information; large mispricing reduces the incentive to acquire it, slower acquisition slows correction, and slow correction preserves the environment in which exploitation remains attractive. The mechanism helps explain why post-publication return predictability decays gradually rather than disappearing at once (McLean and Pontiff, 2016; Chen and Zimmermann, 2020; Jensen et al., 2023), and it complements the literature on mispricing persistence and limits to arbitrage (Stambaugh et al., 2015; Stambaugh and Yuan, 2017; Jacobs, 2016): persistence may reflect not only risk, constraints, or delegated capital, but also the endogenous allocation of investor effort away from fundamental learning. The loop should be strongest when common information is scarce (Proposition 4.4), so that investors lack a shared valuation anchor against which to interpret new information.

The results also speak to market design. Better liquidity improves the incorporation of existing information into prices, but by making mispricing easier to exploit it raises the opportunity cost of acquiring information not yet in the market, so prices can become more informative about existing information and less informative about acquirable information even as standard liquidity measures improve. The evidence in Weller (2018), where algorithmic trading improves short-run price efficiency while reducing incentives to acquire costly fundamental information, illustrates this possibility. Liquidity provision and information production are distinct policy margins,

and in the spirit of Morris and Shin (2002), more favorable trading conditions need not improve welfare once agents' incentives are endogenous.

References

Anderson, S. S. (2022). *Mispricing and the Demand for Fundamental Information*. PhD thesis, Doctoral dissertation, Massachusetts Institute of Technology.

Back, K., Cao, H., and Willard, G. A. (2000). Imperfect competition among informed traders. *Journal of Finance*, 55(5):2117–2155.

Banerjee, S. and Breon-Drish, B. (2022). Dynamics of research and strategic trading. *Review of Financial Studies*, 35(2):908–961.

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

Beatty, R. P. and Ritter, J. R. (1986). Investment banking, reputation, and the underpricing of initial public offerings. *Journal of Financial Economics*, 15(1-2):213–232.

Benveniste, L. M. and Spindt, P. A. (1989). How investment bankers determine the offer price and allocation of new issues. *Journal of Financial Economics*, 24(2):343–361.

Brennan, M. J. and Wang, A. W. (2010). The mispricing return premium. *Review of Financial Studies*, 23(9):3437–3468.

Chaigneau, P., Edmans, A., and Gottlieb, D. (2018). Does improved information improve incentives? *Journal of Financial Economics*, 130(2):291–307.

Chemmanur, T. J. (1993). The pricing of initial public offerings: A dynamic model with information production. *Journal of Finance*, 48(1):285–304.

Chen, A. Y. and Zimmermann, T. (2020). Publication bias and the cross-section of stock returns. *Review of Asset Pricing Studies*, 10(2):249–289.

de Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance*, 45(2):379–395.

Dyer, T., Guest, N., and Yu, E. (2025). New accounting standards and the performance of quantitative investors. *Journal of Accounting and Economics*, 79(2-3):101740.

Edmans, A., Goldstein, I., and Jiang, W. (2015). Feedback effects, asymmetric trading, and the limits to arbitrage. *American Economic Review*, 105(12):3766–3797.

Farboodi, M. and Veldkamp, L. (2020). Long-run growth of financial data technology. *American Economic Review*, 110(8):2485–2523.

Foster, F. D. and Viswanathan, S. (1994). Strategic trading with asymmetrically informed traders and long-lived information. *Journal of Financial and Quantitative Analysis*, 29(4):499–518.

Foster, F. D. and Viswanathan, S. (1996). Strategic trading when agents forecast the forecasts of others. *Journal of Finance*, 51(4):1437–1478.

Goldstein, I., Xiong, Y., and Yang, L. (2025). Information sharing in financial markets. *Journal of Financial Economics*, 163:103967.

Gromb, D. and Vayanos, D. (2010). Limits of arbitrage. *Annual Review of Financial Economics*, 2(1):251–275.

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

Hébert, B. and La'O, J. (2023). Information acquisition, efficiency, and nonfundamental volatility. *Journal of Political Economy*, 131(10):2666–2723.

Holden, C. W. and Subrahmanyam, A. (1992). Long-lived private information and imperfect competition. *Journal of Finance*, 47(1):247–270.

Jacobs, H. (2016). Market maturity and mispricing. *Journal of Financial Economics*, 122(2):270–287.

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

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

Ljungqvist, A. (2007). IPO underpricing. In Eckbo, B. E., editor, *Handbook of Empirical Corporate Finance*, pages 375–422. Elsevier.

McLean, R. D. and Pontiff, J. (2016). Does academic research destroy stock return predictability? *Journal of Finance*, 71(1):5–32.

Morris, S. and Shin, H. S. (2002). Social value of public information. *American Economic Review*, 92(5):1521–1534.

Nezafat, M. and Schroder, M. (2023). The negative value of private information in illiquid markets. *Journal of Economic Theory*, 210:105664.

Ritter, J. R. and Welch, I. (2002). A review of IPO activity, pricing, and allocations.

Journal of Finance, 57(4):1795–1828.

Rock, K. (1986). Why new issues are underpriced. *Journal of Financial Economics*, 15(1–2):187–212.

Stambaugh, R. F., Yu, J., and Yuan, Y. (2015). Arbitrage asymmetry and the idiosyncratic volatility puzzle. *Journal of Finance*, 70(5):1903–1948.

Stambaugh, R. F. and Yuan, Y. (2017). Mispricing factors. *Review of Financial Studies*, 30(4):1270–1315.

Verrecchia, R. E. (1982). Information acquisition in a noisy rational expectations economy. *Econometrica*, 50(6):1415–1430.

Weller, B. M. (2018). Does algorithmic trading reduce information acquisition? *Review of Financial Studies*, 31(6):2184–2226.

Zhang, L. and Zhang, X. (2025). Mispricing and algorithm trading. *Information Systems Research*, 36(1):21–40.