

Dynamics of Research and Strategic Trading

Snehal Banerjee

Rady School of Management, UC San Diego

Bradyn Breon-Drish

Rady School of Management, UC San Diego

We study how dynamic research affects information acquisition in financial markets. In our strategic trading model, the trader performs costly research to generate private information but does not always succeed. Optimal research activity responds to market conditions and generates novel implications. First, more frequent public disclosures can “crowd in” private information acquisition, increase price informativeness, and harm liquidity, instead of “leveling the playing field.” Second, observed research activity does not necessarily imply that traders are better informed. Finally, improvements in research effectiveness or higher market participation by uninformed investors can simultaneously increase price informativeness and liquidity. (*JEL* D82, D84, G12, G14)

Received April 14, 2019; editorial decision November 21, 2020 by Editor Itay Goldstein.

I could spend all my time thinking about technology for the next year and I wouldn't be the hundredth or the thousandth or the 10,000th smartest guy in the country in looking at those businesses.

—Warren Buffet, 1998 Berkshire Hathaway Annual Meeting¹

We thank Brett Green for numerous, invaluable discussions during an early stage of this project. We thank Itay Goldstein (editor and AFA discussant) and two anonymous referees for their guidance. We also thank Samanvaya Agarwal, Stathi Avdis, Kerry Back, Philip Bond, Jesse Davis, Darrell Duffie, Joey Engelberg, Naveen Gandhi, Jungsuk Han, Mariana Khapko, Igor Makarov, Seymour Malamud, Sophie Moinas, Dmitry Orlov, Christine Parlour, Chris Parsons, Uday Rajan, Michael Reher, Francesco Sangiorgi, Michael Sockin, Allan Timmermann, Alberto Teguia, and Dimitri Vayanos and participants at Collegio Carlo Alberto, UC Davis, LSE, the New Economic School, UC Riverside, University of Colorado Boulder, the Higher School of Economics, Hong Kong University, Chinese University of Hong Kong, Hong Kong University of Science and Technology, Frankfurt School of Economics and Finance, Luxembourg School of Finance, the University of Minnesota Junior Finance Faculty Conference, the FIRS 2017 Meeting, Barcelona GSE Summer Forum, the 2017 Western Finance Association Meeting, the 2017 TAU Finance Conference, the 2018 American Finance Association Meeting, the 2018 Finance Symposium at INSEAD, the 2018 UC Santa Barbara LAEF conference, and the 2019 Southwest Economic Theory Conference for helpful suggestions. This paper subsumes part of an earlier paper titled “Dynamic Information Acquisition and Strategic Trading.” All errors are our own. Send correspondence to Bradyn Breon-Drish, bbreondr@ucsd.edu.

¹ This quote comes from a video recording of the 1998 annual meeting available online (<https://buffett.cnbc.com/video/1998/05/04/morning-session—1998-berkshire-hathaway-annual-meeting.html>).

Research outcomes are uncertain: spending time and effort researching trading opportunities does not guarantee a profitable strategy, even for Warren Buffet.² Moreover, research activity varies over time and with market conditions, as suggested by the recent analysis of measures like Bloomberg search volume and EDGAR queries.³ Yet, the literature on information acquisition in financial markets following Grossman and Stiglitz (1980) largely ignores the uncertainty and dynamics of research. Traditional models treat information acquisition as a static, deterministic decision whereby investors make their information choices before trading begins and are *guaranteed* to receive payoff-relevant information if they pay an appropriate cost.

Understanding these limitations is especially important for policy analysis. In many models of static information acquisition, more public disclosure “crowds out” private information (see Section 1 for further discussion). Consequently, regulators who seek to improve market liquidity often require more public disclosures in an effort to “level the playing field.” An important example in practice is Regulation Fair Disclosure (Reg FD), which was introduced in 2000 to reduce the incidence of selective disclosure.⁴ Surprisingly, however, the policy may have had the opposite effect. For instance, Duarte et al. (2008) show the probability of informed trading (PIN) for NASDAQ stocks *increased* after the introduction of Reg FD.

To better understand how the timing and uncertainty of research affects financial markets, we develop a strategic trading model where the trader dynamically optimizes her research activity in response to market conditions. Research activity evolves dynamically and is not guaranteed to succeed: conditional on devoting time to research, the trader receives payoff-relevant information with some probability that depends on the effectiveness of her research technology. When her research is successful, the trader has an uncertain, but limited, window to exploit her private information before a public disclosure at a random future date eliminates her informational advantage.

We show that accounting for the dynamic and uncertain nature of research has important implications. First, research activity can increase or decrease with the frequency of public disclosures.⁵ As a result, policy efforts to “level the playing field” can be counterproductive because requiring more public disclosure can “crowd in” private information acquisition, which increases price impact and harms liquidity. Second, since research success is uncertain, a

² Not surprisingly, this also applies to academic research!

³ For example, see Drake, Roulstone, and Thornock (2012), Ben-Rephael, Da, and Israelsen (2017), Crane, Crotty, and Umar (2019), and Drake et al. (2020).

⁴ Specifically, Reg FD mandated that firms could not disclose material information to some investors unless it was also disclosed publicly. This limited the ability of large, institutional investors to receive and trade on (effectively private) information before others.

⁵ As we will discuss in Section 1, the channel through which public disclosure can “crowd in” more research by the trader in our setting is distinct from settings in which investors can learn about multiple dimensions of payoffs (e.g., Goldstein and Yang 2015; Banerjee, Davis, and Gondhi 2018; Goldstein and Yang 2019).

trader with access to more effective research technology is, counterintuitively, less likely to conduct research but (unconditionally) more likely to be privately informed. This implies that the empirical relation between research activity (e.g., EDGAR queries, Bloomberg search volume) and informed trading need not always be positive, but instead depends on cross-sectional and time-series variation in investor sophistication. Third, when traders acquire information dynamically, we show that some regulatory changes (e.g., improvements to research technology, or increases in the participation of uninformed traders) can improve both price informativeness and market liquidity at the same time. This suggests that regulators do not always face a trade-off between price informativeness and market liquidity, as might be suggested by traditional models.⁶

Section 2 introduces the model. We augment the standard, continuous-time setting of Kyle (1985) and Back (1992), with (a) stochastic noise trading volume (e.g., Collin-Dufresne and Fos 2016; Banerjee and Breon-Drish 2020), (b) a random trading horizon (e.g., Back and Baruch 2004; Caldentey and Stacchetti 2010), and, most importantly, (c) endogenous information acquisition via a probabilistic research technology. Stochastic noise trading volume, first introduced by Collin-Dufresne and Fos (2016) in a different context, is a tractable and empirically plausible mechanism for generating time variation in the value of being informed.⁷ For instance, Collin-Dufresne and Fos (2015) document that informed traders strategically trade more aggressively on days when uninformed trading volume is higher. Our analysis suggests that such days are also likely to be associated with more research activity and information acquisition, which is consistent with evidence (e.g., Ben-Rephael, Da, and Israelsen 2017). A natural interpretation of the random trading horizon is that the potential trading opportunity goes away due to an unscheduled, public announcement that reveals the previously private information. The assumption ensures that the analysis is tractable, and the interpretation as public disclosure allows us to derive direct empirical predictions and policy implications of our model.

The key innovation in our paper is to model dynamic, stochastic research in this setting. We assume that the strategic trader is not endowed with private information, but instead she must engage in research to acquire it. Specifically, at each instant prior to information arrival, the trader can optimally choose her

⁶ In most static acquisition models, policy makers face a trade-off between improving price informativeness and market liquidity: changes that encourage private information acquisition lead to not only more informative prices but also higher adverse selection and lower liquidity. We show that with dynamic research this trade-off continues to be true along some dimensions (e.g., policies that affect the cost of research), but not always.

⁷ Specifically, we assume the instantaneous noise trade to be $dZ_t = v_t dW_{Z_t}$, where v_t follows a stochastic process. As we will discuss in Section 2, the process v_t drives trading activity in the model. For expositional clarity, we follow Collin-Dufresne and Fos (2016) and refer to v_t as “trading volume.” Furthermore, while we argue that stochastic noise trading volume is a natural driver of the value of being informed, as we will discuss below, our results are robust to other mechanisms that generate a time-varying value of information (e.g., time-varying fundamental uncertainty).

research intensity subject to a flow cost.⁸ The research intensity determines the probability with which a (potentially noisy) signal about the asset value arrives in the next instant. If the signal arrives, the trader optimally begins trading on her information. This is in contrast to standard models of static information acquisition in which the trader is restricted to making a one-time decision, before the market opens, about whether or not to acquire information, and in which the information arrives with certainty if she pays the appropriate cost.

We derive the equilibrium in Section 3, characterizing the trader's optimal research and trading strategy, as well as the equilibrium pricing rule, in closed form. To see how dynamics affect the optimal research strategy, it is useful to first consider the optimal *static* acquisition strategy in our setting. The trader follows a "threshold" rule: she acquires information at the beginning of trading if and only if the initial noise trading volume is greater than an optimal threshold value. Essentially, the trader follows a net present value (NPV) rule because she only acquires information if the expected trading profits from being informed are higher than the explicit cost of information. Consistent with the standard results in the literature, this implies that, all else equal, acquisition occurs when (a) the level of uninformed trading volume is high, (b) uncertainty is high, and (c) the frequency of public disclosures is low (i.e., the trading horizon is longer).

With dynamic research, the trader's optimal research strategy also follows a threshold rule. However, in addition to trading off the expected value of being informed and the explicit cost of research, the optimal *dynamic* research choice also reflects a number of real options: the option to delay research until trading opportunities are more profitable (i.e., the level of uninformed trading volume is even higher), and the option to abandon if research has not been successful. The resultant dynamic research "threshold" is higher than in the static acquisition (NPV) threshold; that is, the trader is willing to wait for higher noise trading volume before engaging in research than in the static acquisition case. More importantly, the optimal threshold in a dynamic research setting responds differently to changes in market conditions than the static acquisition threshold. These differences in optimal acquisition, together with the dynamic and uncertain nature of research itself, yield novel predictions on the behavior of research activity and the relation between research and informed trading, which we explore in Section 4.⁹

First, research activity is stochastic and increases with the volatility of trading volume, even after controlling for the level of volume. This prediction is

⁸ For tractability and clarity, we primarily focus on the case in which the trader faces a proportional cost of research, that is, the flow cost of conducting research at intensity ζ is $C(\zeta)dt$, where $C(\zeta)=c\zeta$, subject to a capacity constraint $\zeta \in [0, \bar{\zeta}]$.

⁹ Our model also generates predictions that are largely common to existing static models. For instance, the likelihood and expected duration of research increases with the investor's prior uncertainty and the level of noise trading volume, and decreases in the cost of research. While these are natural, robust, and important implications, they have been explored theoretically and empirically in the existing literature. As such, we focus our attention on results that are more distinctive to our analysis.

consistent with the positive correlation between *abnormal* trading volume and research activity documented in the empirical literature (e.g., Ben-Rephael, Da, and Israelsen [2017] show this for Bloomberg search volume and Google search activity). Moreover, the prediction is in contrast to acquisition in a static setting, which is unrelated to the volatility of volume and either occurs with certainty prior to trading or never occurs. More surprisingly, we also show that the likelihood of research activity is negatively related to the effectiveness of the research technology (or equivalently, the intensity of research activity). Intuitively, a trader with a more effective technology can afford to wait longer to engage in research because, conditional on performing research, she is very likely to receive information. That is, with a more effective research technology, the opportunity cost of waiting to perform research is lower.

Second, the likelihood of conducting research, and the expected time spent in research, is hump shaped in the frequency of public disclosures. Intuitively, this is because the frequency of public disclosures (or the expected length of the trading horizon) has two offsetting effects. On the one hand, more frequent disclosures decrease the value of acquiring information since, conditional on receiving information, the trader expects to exploit her informational advantage over a shorter horizon. This “trading horizon” effect is what leads more frequent public disclosures to “crowd out” private information in the static acquisition benchmark. On the other hand, more frequent disclosures also increase the opportunity cost of waiting to conduct research, by effectively increasing the trader’s “impatience”: she becomes less willing to wait since she may be preempted by the public disclosure. When public disclosures are rare, the impatience effect dominates: increasing the frequency of disclosures pushes the trader to conduct research more frequently (i.e., it crowds in private research). However, when disclosures are sufficiently frequent, the “trading horizon” effect begins to dominate, and further increases in disclosure frequency crowd out private research.

Third, we are also able to pin down the probability that the trader ever successfully receives information and enters the market (i.e., the joint probability that she ever conducts research and that the research is ever successful). This probability inherits many of the properties of the probability of research. However, the relation between research effectiveness and eventual probability of informed trading delivers one of the most striking predictions from the model. While an increase in research effectiveness reduces the probability that the trader ever conducts research, it increases the *unconditional* probability that she ever receives information.¹⁰ Intuitively, higher-skilled analysts or more efficient information systems may generate more *informed*

¹⁰ While it is not surprising that *conditional on conducting research* a trader with better research technology is more likely to receive information, we emphasize that this result is an unconditional one. Moreover, note that variation in the cost of research (or other parameters) would generate a positive relation between the likelihood of research and eventual success. As such, this result is unique to settings with uncertainty in research.

trading activity (and higher profits), even if they do not appear to generate more *research* activity.

Overall, these results suggest that, in settings in which research takes place dynamically and need not succeed, one should exercise caution when using measures of research activity to proxy for information acquisition or informed trading. Moreover, our analysis helps shed light on why the empirical relation between observed research activity and future performance varies across different types of investors. For instance, Drake et al. (2020) document that while EDGAR search activity by institutional investors is predictive of future performance, similar activity by noninstitutional investors is not. Crane, Crotty, and Umar (2019) show that increased reliance on public information by hedge funds leads to higher performance, while Kacperczyk and Seru (2007) argue that the opposite relation holds for mutual funds. These observations line up with our model's predictions about how the relation between research activity and (informed) trading profits depend on the effectiveness of the research technology and the sophistication of investors.

Section 5 highlights how accounting for research dynamics can lead to novel policy implications for price informativeness and liquidity (price impact). We show that both price informativeness and average price impact (Kyle's λ) are hump shaped in the frequency of public disclosures, and the latter is also hump shaped in noise trading volume. This is in contrast to the monotone relations in static acquisition settings, and suggests that efforts to improve price informativeness or enhance liquidity by "leveling the playing field" (e.g., higher public disclosure requirements or facilitating access for uninformed traders) may be counterproductive. These results may also help shed light on the surprising empirical results of Duarte et al. (2008) discussed above. Finally, the natural tension between price informativeness and liquidity in standard settings need not always obtain in a dynamic setting. Policies that make research more effective (i.e., improve the research technology) or increase uninformed trading activity can sometimes improve both price informativeness and liquidity. Together, these results suggest that insights about price informativeness and liquidity that are derived from standard, static acquisition settings need not hold when investors can engage in research dynamically. Consequently, one should exercise caution when evaluating policy proposals in such settings.

1. Related Literature

Our model builds on two strands of the literature. The first involves the continuous-time Kyle (1985) model with a random horizon (e.g., Back and Baruch 2004; Caldentey and Stacchetti 2010) and stochastic volatility in noise trading (e.g., Collin-Dufresne and Fos 2016). The second is the large literature that studies information acquisition in financial markets, following Grossman and Stiglitz (1980) and Verrecchia (1982).

A primary innovation of our model is allowing the strategic trader to dynamically optimize her research and information acquisition activities. This stands in contrast to most of the existing literature on endogenous information acquisition, which restricts investors to acquire information, or commit to an information acquisition strategy, before trading begins, even if trading takes place dynamically. For instance, Back and Pedersen (1998) and Holden and Subrahmanyam (2002) allow investors to commit to receiving signals at particular dates, while Veldkamp (2006) considers a sequence of one-period information acquisition decisions. Similarly, Kendall (2018), Dugast and Foucault (2018), and Huang and Yueshen (2018) study information acquisition with a “time cost,” but in each of these settings the information acquisition decision is made before trading and therefore the decision is effectively a static one. As we discuss further in the following sections, our results suggest that allowing for dynamics in research has important consequences for both positive and normative analysis (see Sections 4.4 and 5.1, respectively).

Two recent works that do not make such a restriction are Banerjee and Breon-Drish (2020) and Han (2018). In Banerjee and Breon-Drish (2020), a companion to the current paper, we also study dynamic information acquisition by a strategic trader but focus on differences between different types of “standard” information acquisition technologies in a setting in which market entry is not detectable by market participants. We contrast settings in which the trader optimizes the precision of a flow of signals and those in which the trader must optimize the arrival time of a “lump” of information. We show that when the trader chooses a flow of information, the optimal signal precision evolves over time with noise trading activity and leads to novel implications for the dynamic behavior of price informativeness. In contrast, when the trader can acquire “lumpy” information, we show that there generally do not exist equilibria with endogenous information acquisition. In contrast, the current paper studies the empirical and policy implications of uncertain and dynamic research, and establishes that when market entry is detectable, an equilibrium with endogenous research can be sustained even with ultimately “lumpy” information arrival.

Han (2018) studies a competitive dynamic trading model based on Hellwig (1980) (particularly its nonlinear generalization due to Breon-Drish [2015]) in which heterogeneously informed investors dynamically optimize the allocation of attention in response to the evolution of aggregate uncertainty. Investors naturally allocate more attention and thereby acquire more precise information when uncertainty is high. However, this endogenously lowers uncertainty in future periods due to learning from prices, which feeds back into current information acquisition decisions. This differs from our setting, in which fundamental uncertainty is constant over time and research activity is driven by noise trading activity. Furthermore, in our model, research intensity itself has no direct effect on market maker (i.e., public) uncertainty, and the rate of

public learning after information arrival is constant because the trader smooths her trades over time to protect her information advantage.

Our analysis also speaks to the large literature on how public disclosure affects price informativeness and real outcomes (see Bond, Edmans, and Goldstein 2012; Goldstein and Sapra 2014; Goldstein and Yang 2017 for recent surveys). A common insight from this literature is that public disclosure “crowds out” private information acquisition and can lead to lower price informativeness (e.g., Diamond 1985; Colombo, Femminis, and Pavan 2014). Intuitively, the arrival of public information (before, simultaneously with, or after related private information) reduces the trader’s anticipated informational advantage relative to other participants and, therefore, reduces her *ex ante* incentive to acquire private information. More recently, a number of papers have highlighted how more public disclosure can “crowd in” additional private learning about fundamentals when investors can learn about multiple components of payoffs (e.g., Goldstein and Yang 2015; Goldstein and Yang 2019). This “crowding in” can improve price informativeness when investors learn about other components of fundamentals. However, when public disclosure about fundamentals crowds out learning from more informative signals (e.g., Dugast and Foucault 2018) or crowds in more learning about other traders (e.g., Banerjee, Davis, and Gondhi 2018), it can harm price informativeness.

Our model highlights a distinct, but related, channel through which public disclosures can “crowd out” private learning in our setting, which is particularly transparent in the *static* acquisition benchmark in Section 3.2. As we show, when the trader is restricted to acquire information before trading begins, more frequent public disclosures discourage information acquisition. This is because she anticipates being able to exploit her informational advantage over a shorter trading horizon (“trading horizon” effect), which reduces her expected benefit from becoming better informed.

The introduction of research dynamics (Section 3.3) highlights a further novel mechanism by which more frequent public disclosure can “crowd in” learning about fundamentals by affecting the trader’s incentives to wait to conduct research. Our result is complementary to the earlier work by Goldstein and Yang (2015) and Goldstein and Yang (2019), but highlights an economically distinct channel, since the public disclosure and the private information acquisition are about the *same* component of payoffs (indeed, in our model, there is only one component to the asset payoff). Intuitively, when public disclosures are extremely rare, the trader is willing to wait a long time before engaging in research, and so information acquisition and price informativeness are low. Increasing (the frequency of) public disclosure initially increases the opportunity cost of waiting, which leads to more frequent research and higher price informativeness. However, when public disclosures are sufficiently frequent, further increases reduce the incentive to acquire information since the trader does not expect to have enough time to exploit her information

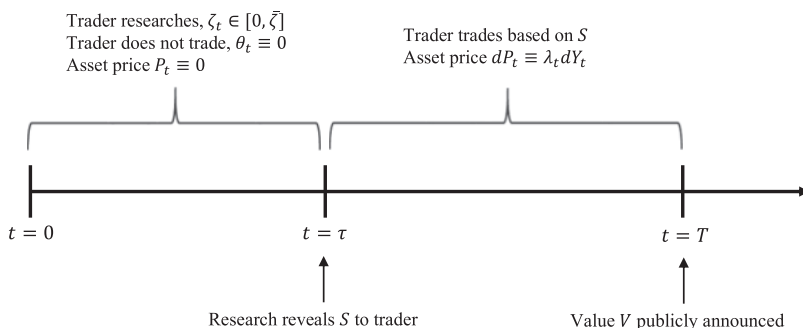


Figure 1

Timeline of the game

The figure illustrates the sequence of moves in the game, in the case that the trader’s research succeeds before the asset value is publicly disclosed, $\tau < T$.

advantage. This suggests that when recommending policy changes that affect the market’s information environment, regulators should account for not just the various dimensions along which investors can acquire information, but also their dynamic incentives to do so.

2. Model

Our framework is similar to Banerjee and Breon-Drish (2020), modified to incorporate dynamic research. Figure 1 presents a timeline of the stages, which we now describe in detail. There are two assets: a risky asset and a risk-free asset with interest rate normalized to zero. The risky asset pays off a terminal value $V \sim N(0, \Sigma_0)$ at random time T , where T is independently exponentially distributed with rate $r > 0$.¹¹ We assume that all market participants have common priors over the distribution of payoffs and signals.

There is a single, risk-neutral strategic trader. Let X_t denote the cumulative holdings of the trader, where we normalize her initial position to $X_0=0$. We consider only absolutely continuous trading strategies, $dX_t = \theta_t dt$, so that the optimal trading problem reduces to choosing the trading rate θ_t .¹² In addition to the strategic trader, there are noise traders whose cumulative holdings Z_t follow

$$dZ_t = v_t dW_{Zt}. \tag{1}$$

¹¹ While we consider the case of a fixed asset value, V , for simplicity, there is no difficulty in accommodating an asset value that evolves over time as a Gaussian process $dV_t = (a - bV_t)dt + \sigma_V dW_{Vt}$ for constants a, b , and σ_V , and independent Brownian motion W_{Vt} . Furthermore, none of our results differs qualitatively in such a setting. While the random date T is important for tractability, we do not expect our results to differ qualitatively in a setting with fixed T . However, such a model will generally be difficult to solve because the optimal research problem effectively reduces to stopping problem on a finite horizon. Such problems are well known to be analytically intractable (e.g., the optimal exercise policy for a finite-maturity American option).

¹² Back (1992) shows that such trading strategies are optimal.

In contrast to benchmark strategic trading models, and following Collin-Dufresne and Fos (2016), the volatility v_t of the noise trading process is stochastic. Specifically, we assume that v_t follows a geometric Brownian motion (GBM)

$$\frac{dv_t}{v_t} = \mu_v dt + \sigma_v dW_{v,t}. \tag{2}$$

We assume that $W_{Z,t}$ and $W_{v,t}$ are independent Brownian motions. Since v drives uninformed “trading volume,” (see, e.g., Collin-Dufresne and Fos 2016), variation in v is a natural and empirically relevant channel of introducing time variation in the value of conducting research as we shall discuss in more detail below.¹³ For ease of exposition, we will follow Collin-Dufresne and Fos (2016) and refer to v as trading volume, μ_v as the drift in trading volume and σ_v as the volatility of trading volume in what follows. We assume that v_t is publicly observable to all market participants. This is without loss since v_t is positive and enters all relevant equilibrium expressions only through v_t^2 , which is the equilibrium order flow volatility and which can be inferred perfectly from the realized quadratic variation of order flow. Moreover, we require $-\infty < \mu_v < r$ to ensure the existence of equilibrium, which we will discuss in more detail below.

In contrast to the existing literature, the strategic trader is not endowed with private information but must engage in research. Specifically, the strategic trader can pay a flow cost $C(\zeta)dt$ to search for information with intensity ζ . That is, conditional on engaging in research at with intensity ζ_t at time- t , information arrives in the next instant with probability $\zeta_t dt$. Formally, the information arrival process is a doubly stochastic Poisson process with rate $\zeta_t \geq 0$ optimally chosen by the trader. Let τ_R denote the time that the trader first engages in research.

Conditional on successful research (i.e., signal arrival) she privately obtains a potentially noisy signal $S = V + \varepsilon$, where $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ is independently distributed. Furthermore, at the time of information arrival the trader (optimally) enters and begins trading in the financial market. τ denotes the time of entry, with $\tau = \infty$ corresponding to no entry (e.g., because research was never successful, or the payoff was revealed before research delivered a signal).¹⁴

¹³ For some intuition, consider a discrete-time analog of the continuous-time setting with length h between periods. Noise trade each period is $Z_{t+h} - Z_t = v_t \sqrt{h} \varepsilon_{t+h}$ for $\varepsilon_{t+h} \sim N(0, 1)$. This implies (expected) trading volume per unit of time is $\frac{\mathbb{E}_t [|Z_{t+h} - Z_t|]}{h} = \frac{v_t \sqrt{\frac{2}{\pi}} h}{h} = v_t \sqrt{\frac{2}{h\pi}}$, which is linear in v_t . It is a well-known artifact of continuous-time settings ($h \rightarrow 0$) in which positions follow diffusions that trading volume is infinite. However, squared volume per unit of time is well-defined and equal to the quadratic variation of the noise trade process, $\frac{\mathbb{E}_t [(Z_{t+h} - Z_t)^2]}{h} = v_t^2$, which is consistent with our treatment of v_t as a natural measure of trading activity.

¹⁴ Treating no entry as an infinite realization for the stopping time, $\tau = \infty$, is standard for arrival/stopping problems. See, for example, the discussion of the first passage time of a Brownian motion with drift on p. 196 of Karatzas and Shreve (1998), or more concretely, the discussion of the exercise time of an American option on p. 341 of Shreve (2004) for applications.

A competitive, risk-neutral market maker sets the price of the risky asset equal to the conditional expected payoff given the public information set. Let \mathcal{F}_t^P denote the public information filtration, which is that induced from observing the aggregate order flow process $Y_t = X_t + Z_t$, stochastic trading volume v_t , and the entry of the informed trader $\mathbf{1}_{\{\tau \leq t\}}$. Note that we assume that the market maker can detect entry by the strategic trader into the financial market, but we do *not* assume that the market maker can observe her research intensity.¹⁵ The price at time $t < T$ is therefore given by

$$P_t = \mathbb{E}[V | \mathcal{F}_t^P]. \tag{3}$$

Let \mathcal{F}_t^I denote the augmentation of the filtration $\sigma(\mathcal{F}_t^P \cup \sigma(S))$. Thus, \mathcal{F}_t^I represents the strategic trader’s information set, post-entry.

We require the trader’s research strategy ζ_t to be weakly positive and adapted to \mathcal{F}_t^P . Importantly, note that we allow the trading intensity to depend on all of the information available to the trader, but as we shall see below, it will only depend on uninformed trading volume (i.e., v_t) in equilibrium. Moreover, conditional on not having received any information so far, the trader is permitted to disengage from research and then re-engage at any time of her choosing.

We require the trader’s pre-entry trading strategy be adapted to \mathcal{F}_t^P and her post-acquisition trading strategy be adapted to \mathcal{F}_t^I . To eliminate doubling-type strategies that accumulate unbounded losses followed by unbounded gains, we also require the following, standard, admissibility condition on trading strategies

$$\mathbb{E} \left[\int_0^\infty e^{-ru} \theta_u^2 du \right] < \infty.$$

Our definition of equilibrium is standard, but modified to account for endogenous entry.

Definition 1. An equilibrium is (a) an admissible research strategy ζ_t and trading strategy X_t for the trader and (b) a price process P_t , such that, given the trader’s strategy the price process satisfies (3) and, given the price process, the research strategy and trading strategy maximize the expected profit

$$\mathbb{E} \left[\int_0^T (V - P_u) \theta_u du \right] = \mathbb{E} \left[\int_0^\infty e^{-ru} (V - P_u) \theta_u du \right]. \tag{4}$$

Let $\hat{V} \equiv \mathbb{E}[V|S]$ and $\Omega \equiv \text{var}(V|S)$ denote the conditional beliefs about V given observation of the signal S . Let $\Psi_t \equiv \text{var}(\hat{V} | \mathcal{F}_t^P)$ denote the market maker’s uncertainty about \hat{V} , and, using the law of total variance, let $\Sigma_t = \Psi_t + \Omega$

¹⁵ In fact, because entry is detectable and the market maker’s pricing problem depends on entry, but not the research intensity, the equilibrium is identical regardless of whether or not the market maker correctly conjectures (or directly observes) the trader’s chosen intensity.

denote her uncertainty about the payoff itself. Then, given the information structure, we have:

$$\hat{V} = \frac{\Sigma_0}{\Sigma_0 + \sigma_\varepsilon^2} S, \tag{5}$$

$$\Omega = \frac{\Sigma_0 \sigma_\varepsilon^2}{\Sigma_0 + \sigma_\varepsilon^2}, \tag{6}$$

and for all times $0 \leq t \leq \tau$, prior to information arrival, we have

$$\Psi_t = \Psi_0 \equiv \Sigma_0 \frac{\Sigma_0}{\Sigma_0 + \sigma_\varepsilon^2}, \quad t \in [0, \tau].$$

The dynamics of Ψ_t after the trader observes information will be pinned down as part of the equilibrium, as the market maker learns about S from observing the order flow. Finally, by the Law of Iterated Expectations, the market-maker's conditional expectation of the asset value can be reduced to her conditional expectation of \hat{V} :

$$P_t = \mathbb{E} \left[\hat{V} \mid \mathcal{F}_t^P \right]. \tag{7}$$

One can interpret Ψ as a measure of the trader's informational advantage over the market, in the event that her research succeeds. If the signal is completely uninformative, $\sigma_\varepsilon^2 \rightarrow \infty$, $\Psi_t \rightarrow 0$. Conversely, when the signal is perfectly informative (as in, e.g., Kyle 1985), we have $\Psi_t = \Sigma_0$ prior to information arrival, so that the market maker's uncertainty about the value estimate reduces to her uncertainty about the value itself.

2.1 Discussion of assumptions

We interpret T as the date of an unscheduled, value-relevant public disclosure or announcement, which eliminates the trader's informational advantage. Intuitively, we think of the model as capturing one "round" of a repeated version of the model in which there are a sequence of public announcements at different, random times. One can show that under standard assumptions, each round of a formal, repeated model is isomorphic to the benchmark model we analyze, which indicates that our assumption of a single announcement is without loss.¹⁶ It is not critical that the asset pays off or that trading stops at

¹⁶ Specifically, consider a repeated version of the model with $0 < N \leq \infty$ "rounds," where each round ends with a public announcement of payoff-relevant news. The asset payoff is the sum of N components, which we interpret as the innovation to the asset value in each round

$$V = \sum_{n=1}^N V_n$$

with $V_n \sim N(0, \Sigma_0)$ i.i.d. The V_n are publicly disclosed at a sequence at random times $T_n = \sum_{j=1}^n \Delta T_j$ with the $\Delta T_j \sim \exp(r)$ i.i.d. One can show that the optimal research and trading strategies and equilibrium price dynamics in such a setting follow from applying the results of our single-announcement model round-by-round to the components V_n . Full details are available on request.

T , but only that the trader's informational advantage disappears. We shall refer to the rate r as the rate, or frequency, of public disclosure. The assumption that T is exponentially distributed is made for tractability, because it makes the trader's problem time stationary. In contrast, a deterministic horizon would be a more appropriate specification for scheduled, periodic disclosures (e.g., earnings announcements) and could generate interesting deadline effects for the trader's research and trading decisions. However, the introduction of an additional state variable (calendar time) makes the trader's optimal research problem analytically intractable, so we leave this analysis for future work.

Allowing for stochastic variation in uninformed trading volume (i.e., stochastic ν) provides a natural, empirically relevant channel through which the value of becoming informed and entering the market evolves over time. For instance, Collin-Dufresne and Fos (2015) show that activist investors, who are plausibly informed, time their trades during periods of high liquidity. Intuitively, conducting research and trading on acquired information is more attractive when there is more trading by liquidity or sentiment traders, which improves market liquidity. We expect similar trade-offs to arise in other settings with a directly time-varying value of private information, for example, in which the asset value V has (publicly observable) stochastic volatility. It is also straightforward to generalize to a general continuous process for ν_t , but at the expense of closed-form solutions to the optimal research problem in most cases.

Because the trader begins without an informational advantage and without any additional motives for trade, her expected profit from following any trading strategy before information arrival is zero. Hence, without loss, we assume that she refrains from trading prior to receiving information.¹⁷ Moreover, in our setting, it is optimal for the trader to begin trading immediately on information arrival instead of waiting to enter the financial market.¹⁸ These features help maintain tractability and allow us to cleanly study the interaction between research dynamics and market outcomes. In practice, investors, especially sophisticated institutions, might trade before acquiring information (e.g., to provide liquidity, make markets, hedge existing positions) or not trade after acquiring information (e.g., to time their trading around scheduled announcements) for many reasons. An analysis of how these forces interact with the incentives to conduct research, though interesting, are beyond the scope of the current paper and left for future work.

We assume that the strategic trader's entry decision is detected by the market maker. This is motivated by theoretical and empirical considerations. First, as Banerjee and Breon-Drish (2020) show in a related setting with

¹⁷ No trade would be strictly optimal if trading involved even an arbitrarily small trading cost, either explicit or implicit, prior to information arrival.

¹⁸ Immediate trading is a result in our model since there is no benefit to waiting to trade after information is acquired, but there is a cost: the trading opportunity disappears if the information is publicly disclosed.

dynamic information acquisition, when there is a fixed component to the cost of information acquisition, there does not exist an equilibrium when entry is not observable by the market maker. Intuitively, in this case, the strategic trader can always deviate from any conjectured equilibrium strategy by either preemptively entering or delaying entry. Second, in practice, entry by large investors into new markets is publicly scrutinized by the financial media. The addition of star traders, portfolio managers, and executives also garners significant media attention. Even if not covered by the popular press, participation by large traders is often known to other market participants.¹⁹ Third, many institutional investors are subject to regulatory reporting requirements, and disclosures about trading positions and capital adequacy can provide (potentially noisy) information about trading activity. Finally, it is worth noting that if given the choice, a trader might prefer to strategically disclose the extent of her informational advantage and the timing of her entry. For instance, Xiong and Yang (2020) show that in a setting with static information acquisition and immediate entry, a strategic trader may prefer to publicly disclose the precision of her private signal (which they refer to as “overt acquisition”) rather than keeping this choice private (“secret acquisition”).

Finally, for expositional clarity, we assume that the strategic trader does not explicitly discount future profits. First, this facilitates direct comparison of our results with those of benchmark Kyle-style models, which do not typically feature discounting. Second, one can show that in a setting in which the trader maximizes discounted expected trading profits with explicit subjective discount rate $\delta \geq 0$, the structure of the equilibrium (specifically, Propositions 1 and 2) is identical and simply replaces the implicit discount rate r with the sum of implicit and explicit rates, $\hat{r} \equiv r + \delta$. The remaining formal results on research dynamics, price informativeness, and liquidity that we present are also straightforward extensions, and yield qualitatively similar implications.²⁰

3. Equilibrium

In this section we construct an overall equilibrium of the model by working backward. In Section 3.1, we characterize the financial market equilibrium given an entry time τ . In 3.2, we solve for the benchmark static research problem in which the trader commits whether or not to conduct research at time $t=0$. Finally, in Section 3.3, we characterize the optimal, dynamic research strategy of the trader.

¹⁹ For instance, major broker-dealers that provide block trading services (or “upstairs trading desks”) directly observe trading demand from institutional investors, and so can detect “entry” or increased participation. Similarly, prime brokers observe the cash and securities positions of their clients, and counterparties in over-the-counter (OTC) derivative transactions disclose their interests to each other through ISDA agreements.

²⁰ Details are available on request.

3.1 Preliminary analysis

In the following result, we characterize the financial market equilibrium, conditional on an arbitrary acquisition time.

Proposition 1. Fix an information arrival time $\tau \in \mathcal{T}$ and for $t \geq \tau$ let $G_t = \frac{e^{-2r(t-\tau)}}{2(r-\mu_v)} v_t^2$. As long as $\mu_v < r$, there exists a post-entry equilibrium (i.e., for $t \in [\tau, T)$) in which the price of the risky asset follows $dP_t = \lambda_t dY_t$ and optimal trading rate follows $dX_t = \theta_t dt = \beta_t (\hat{V} - P_t) dt$, where

$$\lambda_t = e^{-r(t-\tau)} \sqrt{\frac{\Psi_t}{G_t}} = e^{-(r-\mu_v)(t-\tau)} \sqrt{\frac{2(r-\mu_v)\Psi_0}{v_t^2}}, \quad \beta_t = \frac{\lambda_t v_t^2}{\Psi_t} \quad (8)$$

the market maker's conditional variance of the trader's valuation $\Psi_t \equiv \text{var}(\hat{V} | \mathcal{F}_t^P)$ is given by

$$\Psi_t = e^{-\int_t^\tau e^{-2r(s-\tau)} \frac{v_s^2}{G_s} ds} \Psi_0 = e^{-2(r-\mu_v)(t-\tau)} \Psi_0, \quad (9)$$

and the trader's value function is given by

$$J(\hat{V}, P_t, \Psi_t, \lambda_t) = \frac{1}{2\lambda_t} \left((\hat{V} - P_t)^2 + \Psi_t \right). \quad (10)$$

The unconditional, gross expected trading profit from becoming informed at any time t prior to information arrival is

$$\bar{J}(\cdot) \equiv \mathbb{E} \left[J(\hat{V}, P_t, \Psi_0, \lambda_\tau) | \mathcal{F}_t^P \right] = \sqrt{\frac{\Psi_0}{2(r-\mu_v)}} v_t. \quad (11)$$

A few comments about the result are in order. First, the restriction $\mu_v < r$ is required to guarantee the existence of equilibrium. If it did not hold, the expected growth in noise trading would cause the trader to optimally refrain from trading to wait for a noisier market in the future. In such an equilibrium, the market maker would optimally set price impact to zero, since she does not face any informed trading. However, this cannot be an equilibrium, since the informed trader could profitably deviate from her prescribed strategy and trade against the unresponsive pricing rule.

Second, conditional on the information arrival time, r and μ_v enter the financial market equilibrium only through $r - \mu_v$. Intuitively, this is because in the trading game, changes in r and μ_v both effectively change the trader's degree of patience with respect to her trading activity. For instance, when r increases, the trader becomes less patient because she is more concerned about being preempted by the public announcement before she can trade a sufficient quantity, and her expected profit decreases. When μ_v increases, the trader becomes more patient because all else equal she desires to time her trades to times of high noise trading activity, and her expected profit increases.

Third, both price impact λ_t and trading aggressiveness β_t take intuitive forms. The price impact is higher (trading aggressiveness lower) if the trader's informational advantage is stronger (i.e., Ψ_t is higher) or there is less noise trade (i.e., ν_t is lower). Fourth, if the public announcement never arrives, $T = \infty$, all of the trader's information is incorporated into the price in the limit, that is, $\Psi_t \rightarrow 0$ as $t \rightarrow \infty$. Finally, the expected gross trading profits conditional on information arrival, \bar{J} , drives the trader's research activity. The above result highlights that the value of observing the signal S increases with its informativeness (i.e., Ψ), uninformed trading volume (i.e., ν_t), the drift in trading volume (i.e., μ_ν) and decreases in the frequency of public disclosure (i.e., r) which captures the expected duration of the trader's information advantage.

3.2 Benchmark: Static information choice

As a benchmark for comparison, we first characterize the optimal information acquisition choice when it is a one-shot, static decision. Specifically, suppose the strategic trader chooses whether or not to pay a fixed cost C to immediately acquire signal S at time zero. The following result characterizes the optimal acquisition strategy in this case.

Corollary 1. (Static Information Acquisition) Suppose the investor chooses whether to pay a fixed, lump cost C to immediately acquire signal S at date zero. Then, she chooses to acquire information if and only if $v_0 \geq v_S$, where

$$v_S = C \sqrt{\frac{2(r - \mu_\nu)}{\Psi_0}}. \quad (12)$$

The optimal acquisition boundary v_S increases in C and r , decreases in μ_ν and Ψ_0 , and is invariant to σ_ν .

The above result is intuitive and follows immediately from comparing the expected benefit of acquiring information in Equation (11) at time zero to the cost C . The trader acquires information only when uninformed trading volume is sufficiently high, and the acquisition threshold v_S increases in the cost C and the disclosure rate r , but decreases in the trader's informational advantage Ψ_0 and the drift of trading volume μ_ν . Intuitively, the optimal strategy is analogous to following the NPV rule - only acquire information if the expected value of being informed is higher than the expected cost.

The volatility of trading volume σ_ν has no effect on static information acquisition. Because increases in σ_ν symmetrically increase the likelihood of high and low noise trade activity, it does not affect the expected trading profit, and therefore does not affect the trader's acquisition decision. As we will see, this contrasts with our dynamic setting, in which the volatility of trading volume plays an important role in the value of conducting research.

Moreover, note also that with static acquisition, increasing the rate of public disclosures r reduces the ability of the trader to exploit his informational

advantage because the (expected) trading horizon is shorter and therefore tends to discourage information acquisition. This is consistent with the standard intuition from models with static information acquisition, where more public disclosure “crowds out” private learning by investors. As we show below, this monotonic relation between information acquisition and trading horizon does not hold when the trader can choose the timing of her research activity.

3.3 Optimal research strategy

In this section, we characterize the optimal research strategy of the trader. The trader’s research intensity problem is:

$$J^U \equiv \sup_{\{\zeta_s\}} \mathbb{E}_t \left[1_{\{\tau < T\}} \bar{J}(v_\tau) - \int_t^{\min\{\tau, T\}} C(\zeta_s) ds \right]. \tag{13}$$

That is, the trader dynamically optimizes her research intensity, ζ_t , to maximize the expected gross trading profit, which she enjoys if and only if information arrives before the payoff is revealed, $\tau < T$. However, she also pays a flow research cost $C(\zeta)$, which depends on the chosen intensity, up to the point that information arrives or the payoff is revealed, whichever comes first. The corresponding HJB equation, which is the PDE that characterizes the trader’s value function J^U conditional on still being uninformed (i.e., on not yet having observed a signal nor the payoff being revealed), is given by

$$0 = \sup_{\zeta \geq 0} r(0 - J^U) + \zeta(\bar{J} - J^U) - C(\zeta) + \mathcal{D}J. \tag{14}$$

where $\mathcal{D}J = \mu_v v J_v + \frac{1}{2} \sigma_v^2 v^2 J_{vv}$ is the generator of the v process, applied to the function J . Conditional on not being informed, over the next instant, either the trader’s value becomes zero with probability $r dt$ (if the value is publicly disclosed) or she acquires information with probability ζdt if she is currently conducting research (i.e., if $\zeta > 0$) by paying flow cost $C(\zeta) dt$.

In principle, given a cost specification, one proceeds by jointly characterizing (a) the optimal choice ζ^* and (b) the uninformed trader’s value function J^U . For instance, if we assume $C(\cdot)$ is strictly increasing, strictly convex, twice continuously differentiable and has $C(0) = C'(0) = 0$, then we can characterize the optimal intensity ζ^* using the first-order condition:

$$C'(\zeta_t^*) = (\bar{J} - J^U). \tag{15}$$

Since $C(\cdot)$ is strictly convex, we can define the inverse cost $f(\cdot) = [C']^{-1}(\cdot)$. Then, plugging into the HJB equation above yields

$$0 = r(0 - J^U) + f(\bar{J} - J^U)(\bar{J} - J^U) - c(f(\bar{J} - J^U)) + \mathcal{D}J, \tag{16}$$

which is a partial differential equation that characterizes J^U .

Unfortunately, an analytical solution for $J^U(\cdot)$ is not available for arbitrary cost specifications. For tractability, we focus on the special case of a

proportional cost of research. Specifically, suppose the trader can choose research intensity $\zeta \in [0, \bar{\zeta}]$ at a flow cost $C(\zeta)dt = c\zeta dt$, where $c > 0$ is the cost of each unit of intensity.²¹ In this case, we show that the optimal research strategy takes a particularly transparent, intuitive form. The trader engages in research if and only if trading volume exceeds a threshold v^* , as characterized by the next result.

Proposition 2. (Optimal research intensity) Suppose that the investor chooses research intensity $\zeta \in [0, \bar{\zeta}]$ at a flow cost $C(\zeta)dt = c\zeta dt$. Then, the optimal research strategy is characterized by a threshold: the trader engages in research at rate $\bar{\zeta}$ if and only if uninformed trading volume $v_t \geq v^*$ and does not engage in research otherwise, where

$$v^* = \frac{\beta_1}{\beta_1 - 1} \frac{\beta_2}{\beta_2 - 1} \frac{r + \bar{\zeta} - \mu_v}{r + \bar{\zeta}} c \sqrt{\frac{2(r - \mu_v)}{\Psi_0}}, \tag{17}$$

$$\beta_1 = \frac{1}{2} - \frac{\mu_v}{\sigma_v^2} + \sqrt{\left(\frac{1}{2} - \frac{\mu_v}{\sigma_v^2}\right)^2 + \frac{2r}{\sigma_v^2}}, \quad \text{and} \tag{18}$$

$$\beta_2 = \frac{1}{2} - \frac{\mu_v}{\sigma_v^2} - \sqrt{\left(\frac{1}{2} - \frac{\mu_v}{\sigma_v^2}\right)^2 + \frac{2(r + \bar{\zeta})}{\sigma_v^2}}. \tag{19}$$

The optimal threshold v^* is increasing in c , $\bar{\zeta}$, σ_v , and μ_v , decreasing in Ψ , and U shaped in r . Moreover, trader’s value function, prior to the arrival of the signal, is given by

$$J^U(v_t) = \begin{cases} \frac{c\bar{\zeta}\beta_2}{(\beta_1 - 1)(\beta_2 - \beta_1)(r + \bar{\zeta})} \left(\frac{v_t}{v^*}\right)^{\beta_1} & \text{when } v_t < v^* \\ \frac{c\bar{\zeta}\beta_1}{(\beta_2 - 1)(\beta_2 - \beta_1)(r + \bar{\zeta})} \left(\frac{v_t}{v^*}\right)^{\beta_2} + \frac{\bar{\zeta}v_t}{r - \mu_v + \bar{\zeta}} \sqrt{\frac{\Psi_0}{2(r - \mu_v)}} - \frac{c\bar{\zeta}}{r + \bar{\zeta}} & \text{when } v_t \geq v^* \end{cases} \tag{20}$$

The optimal threshold v^* reflects the two sources of optionality embedded in research: (1) the option to delay and (2) the option to abandon. To gain some intuition, it is useful to compare the expression for v^* to the threshold for static information acquisition v_S and to the threshold for a hypothetical, infinite intensity research technology (i.e., the threshold as $\bar{\zeta} \rightarrow \infty$) under which the signal arrives with probability one the instant that the trader engages in research.

The static information acquisition threshold $v_S = C \sqrt{\frac{2(r - \mu_v)}{\Psi_0}}$, characterized in Corollary 1, reflects the fact that the trader has neither the ability to delay research nor the ability to abandon it, once started. Specifically, this is the

²¹ One could also interpret this as a first-order approximation to more flexible cost functions that are well-behaved, as earlier described. Due to the linear cost function, our equilibrium is also isomorphic to one that would arise in a model with a single available research intensity $\bar{\zeta}$.

threshold that the trader would optimally choose when she must make a one-shot decision to immediately acquire or never acquire information. Next, consider the threshold that the trader would optimally choose if she had access to a dynamic research technology with infinite intensity. In this case, the trader has the option to delay conducting research, but when she chooses to do so, information arrives immediately with probability one, with no opportunity to abandon or reverse the research decision. The following corollary characterizes the optimal threshold for such dynamic information acquisition.

Corollary 2. Suppose the investor can pay a fixed, lump cost C to immediately acquire signal S at a time of her choosing. Then, the optimal information acquisition strategy is characterized by a threshold: the trader acquires information if and only if uninformed trading volume $v_t \geq v_\infty^*$, where

$$v_\infty^* = \frac{\beta_1}{\beta_1 - 1} C \sqrt{\frac{2(r - \mu_v)}{\Psi_0}}, \text{ and } \beta_1 = \frac{1}{2} - \frac{\mu_v}{\sigma_v^2} + \sqrt{\left(\frac{1}{2} - \frac{\mu_v}{\sigma_v^2}\right)^2 + \frac{2r}{\sigma_v^2}}. \quad (21)$$

The optimal threshold v_∞^* is increasing in C , σ_v , and μ_v , decreasing in Ψ , and U shaped in r .

The “infinite intensity” dynamic acquisition threshold can be written as $v_\infty^* = \frac{\beta_1}{\beta_1 - 1} v_S$, which differs from the static acquisition threshold v_S by a coefficient $\frac{\beta_1}{\beta_1 - 1} > 1$. The coefficient reflects the trade-off embedded in the option to delay: at any moment, the trader can acquire information now and begin exploiting her informational advantage, or wait until the value of information increases (when v_t increases). The value of this option to wait leads to an optimal acquisition threshold that is higher than in the case of static acquisition (i.e., $v_\infty^* > v_S$): the higher threshold reflects not only the expected value of acquiring information (as in the “NPV” decision rule) but also the additional real option value of being able to wait and acquire at a later date. Intuitively, the trader optimally waits past the threshold at which the “naive” expected value of being informed is equal to the explicit cost of information because acquiring information involves the additional opportunity cost of giving up the option to wait longer. Naturally, she waits until the “all in” expected payoff, accounting for this opportunity cost, is weakly positive.

The value of the option to wait, and therefore the threshold, increases with both the drift μ_v and volatility σ_v of uninformed trading volume. These results are intuitive. A higher drift μ_v implies future trading volume is likely to be higher, which increases the option value of waiting. A higher μ_v also increases the expected trading profit, conditional on information arrival, which tends to reduce the threshold, but this effect is dominated by the option value effect for the range of values considered in the plot.²² Similarly, a higher volatility

²² In general, the dependence of the threshold on μ_v can be nonmonotonic (U shaped).

σ_v of trading volume makes the option to wait more valuable. These results stand in contrast to the static benchmark, in which there is no optionality in the information acquisition decision, and therefore the optimal acquisition threshold (i.e., v_S) decreases with the drift and does not depend on the volatility of trading volume.

The optimal threshold with dynamic information acquisition is nonmonotonic in the frequency of public disclosure r . This is in contrast to the static acquisition case, where the marginal effect of increasing the frequency of public disclosure is to reduce the expected duration of the trader's information advantage and, consequently, increase the acquisition threshold. We refer to this as the "trading horizon" effect. An additional, novel effect of changes in public disclosure emerges with dynamic acquisition. Intuitively, the marginal effect of a higher disclosure frequency (higher r) is to increase the cost of waiting to engage in research because it increases the likelihood that the information is publicly revealed before the investor has a chance to acquire and trade on it. This effect tends to push the acquisition threshold downward because it incentivizes the trader to conduct research sooner. We refer to this as the "impatience" effect.

The shape of the overall relation between disclosure frequency and optimal research threshold depends on the relative magnitudes of these two marginal effects. When the disclosure frequency is very low, the "impatience" effect dominates: an increase in r leads to a decrease in the threshold because the value of waiting grows more rapidly than the conditional expected trading profits shrink. However, when the disclosure frequency is sufficiently high, the "trading horizon" effect dominates and further increases in r increase the threshold.

Note that while the optimal trading strategy (and equilibrium price) and the expected profit in the event of successful research depend on $r - \mu_v$ (see Proposition 1), the impact of r and μ_v for the optimal research threshold v_S^* and the value function J^U is slightly richer but qualitatively similar. Specifically, holding the expected profit conditional on information arrival (i.e., $\sqrt{\frac{\Psi_0}{2(r - \mu_v)}}$) fixed, changes in r and μ_v affect the trader's dynamic incentives to engage in research because they effectively change the trader's patience. Intuitively, higher frequency of disclosures makes the trader less patient (since the public disclosure is more likely to arrive over the next instant), but a higher drift in noise trading volume μ_v makes the trader more patient (since information becomes more valuable when there is more noise trading).

The "infinite intensity" threshold v_S^* reflects the effect of the option to delay, relative to the static benchmark. The general research technology we consider embeds an additional option to abandon: if the trader has engaged in research but it has not been successful (i.e., the signal has not yet arrived), she can always choose to stop conducting research at any point in time. The option to abandon reduces the expected cost of initiating research and, consequently, lowers the optimal threshold. That is, knowing that she can walk away from research at

any time, the trader’s incentive to wait is lower than in the infinite intensity case.

To see why, note that the expected total cost of conducting research (at the time the trader first initiates research τ_R) is equal to²³

$$\mathbb{E} \left[\int_{\tau_R}^{\min\{\tau, T\}} c\bar{\zeta} 1_{\{v_s \geq v^*\}} ds \mid v_{\tau_R} = v^* \right] = \frac{\beta_2}{\beta_2 - \beta_1} \frac{c\bar{\zeta}}{r + \bar{\zeta}}. \tag{22}$$

That is, the trader pays flow cost $c\bar{\zeta}$ at any time s that she conducts research $v_s > v^*$, as long as information has not yet arrived and the asset value has not been disclosed, $s < \min\{\tau_I, T\}$. Here, $\frac{c\bar{\zeta}}{r + \bar{\zeta}}$ is the total expected flow cost of conducting research without the option to abandon (i.e., if the trader would have to pay a flow cost of $c\bar{\zeta} dt$ at all times after initiating research until it was either successful or the value was publicly realized). The coefficient $0 < \frac{\beta_2}{\beta_2 - \beta_1} < 1$ reflects the effect of the option to abandon.²⁴ Intuitively, the expected total cost of starting research is lower because the trader can choose to stop conducting research at a time of her choosing.

Finally, note that as the research intensity becomes arbitrarily large, the total expected cost of research is

$$\lim_{\bar{\zeta} \rightarrow \infty} \frac{\beta_2}{\beta_2 - \beta_1} \frac{c\bar{\zeta}}{r + \bar{\zeta}} = c, \tag{23}$$

and the optimal research threshold approaches the “infinite intensity” one from below, that is,

$$\lim_{\bar{\zeta} \rightarrow \infty} v^*(c) = \frac{\beta_1}{\beta_1 - 1} c \sqrt{\frac{2(r - \mu_v)}{\Psi_0}} = v_\infty^*(c). \tag{24}$$

This implies that the “infinite intensity” technology of Corollary 2 is the limit of the general research technology as we make the research intensity arbitrarily high, but keep fixed the per unit flow cost, c , of research intensity. The above also implies that for any finite intensity research technology with identical per-unit flow cost, the additional option to abandon lowers the optimal research threshold, that is,

$$v^*(c) = \frac{\beta_2}{\beta_2 - 1} \frac{r + \bar{\zeta} - \mu_v}{r + \bar{\zeta}} v_\infty^*(c) \leq v_\infty^*(c). \tag{25}$$

²³ This expression follows from Proposition 4 below.

²⁴ More precisely, it reflects the option to abandon and then restart research at any point in time. If the abandon decision was irreversible (i.e., if we do not allow the trader to resume research if she has abandoned it once), then the multiplier would instead be $\frac{\beta_2}{\beta_2 - 1}$, which, naturally, is smaller as the expected cost is lower if she can never reinitiate research.




Figure 2

Total expected cost of research $\frac{\beta_2 - c\bar{\zeta}}{\beta_2 - \beta_1 r + c\bar{\zeta}}$ and dynamic research threshold v^* as a function of maximum intensity $\bar{\zeta}$

This figure plots the expected total cost of research $\frac{\beta_2 - c\bar{\zeta}}{\beta_2 - \beta_1 r + c\bar{\zeta}}$ and the dynamic research threshold v^* as a function of the maximum research intensity $\bar{\zeta}$. Unless specified, other parameters are set to $\mu=0.05$, $\Psi_0=1$, $c=C=1$, $r=0.2$, and $\sigma_v=0.1$.

The extent to which $v^*(c)$ is lower than $v_\infty^*(c)$ is driven by the maximal research intensity $\bar{\zeta}$: as research intensity is higher, $v^*(c)$ increases and approaches $v_\infty^*(c)$. At first glance, it may appear surprising that when the trader has a more effective research technology, she actually waits for more extreme noise trading before engaging in research. However, the intuition follows directly from the discussion above. When $\bar{\zeta}$ increases the option to abandon becomes less valuable. That is, as $\bar{\zeta}$ increases, in the event that the trader initiates research, it becomes more likely that information will arrive before she has a chance to abandon. Figure 2 plots the threshold v^* as a function of the maximum intensity $\bar{\zeta}$ for various values of the disclosure frequency r .

Figure 3 illustrates the various thresholds and how they depend on some of the underlying parameters. Specifically, the left panels plot the static information acquisition threshold v_S as a function of the drift μ_v and volatility σ_v of uninformed trading volume, and the frequency of public disclosure r , whereas the right panels plot the corresponding dynamic information acquisition threshold v_∞^* and the dynamic research thresholds v^* for various values of research intensity ($\bar{\zeta} \in \{0.25, 1, 4\}$). Consistent with Lemma 1, the optimal static acquisition threshold is decreasing in μ_v , independent of σ_v , and increasing in r . In contrast, allowing for dynamic information acquisition implies that the acquisition threshold is increasing in μ_v , increasing in σ_v and U shaped in r , as established in Theorem 2 and Corollary 2. As discussed above, these qualitative differences arise because of the dynamic nature of the information acquisition and research decisions and reflect the effect of the trader’s ability to delay acquiring research. Finally, note that the threshold is increasing in the research intensity parameter $\bar{\zeta}$, which reflects the effect of the secondary option embedded in research, namely, the trader’s ability to abandon at a time of her choosing.

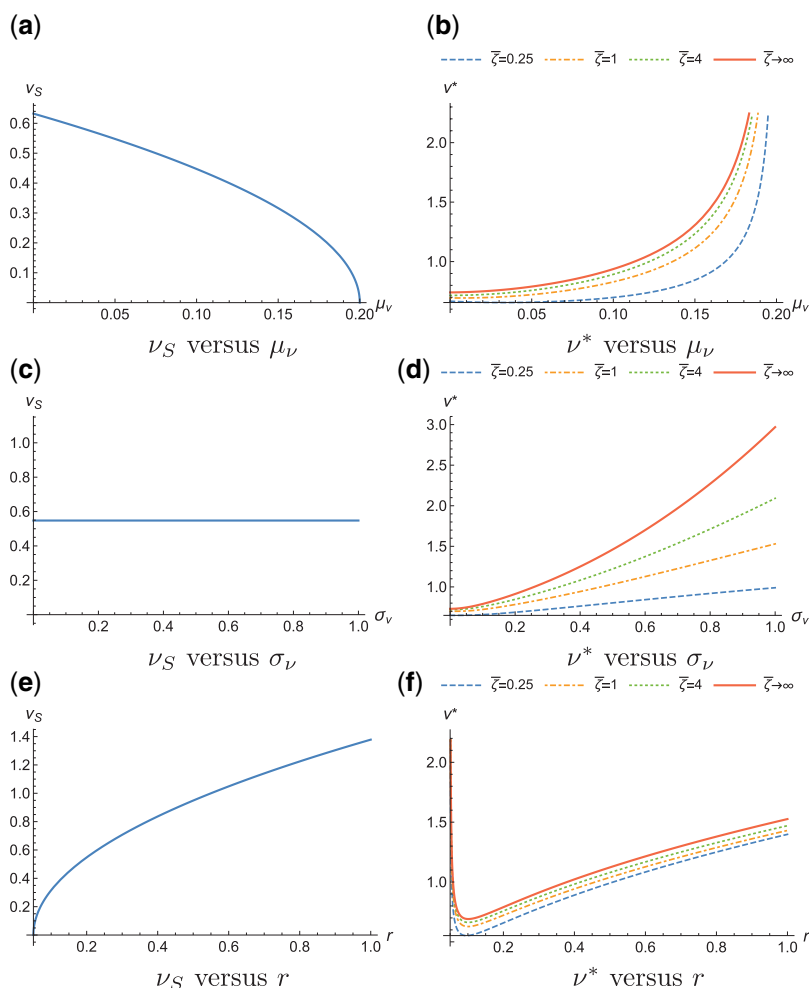


Figure 3
Optimal static acquisition threshold, v_S , and dynamic research threshold, v^*
 The figure plots the optimal, static acquisition threshold v_S and the optimal, dynamic research threshold v^* as a function of public disclosure frequency (i.e., r), and the drift (i.e., μ_v) and volatility (i.e., σ_v) of trading volume. Unless specified, other parameters are set to $\mu=0.05$, $\Psi_0=1$, $c=C=1$, $r=0.2$, and $\sigma_v=0.1$.

4. Dynamics of Research

In this section, we characterize the dynamics of research and information arrival. In the first subsection, we begin by studying the properties of the time that the trader first engages in research, and show that if the initial level of noise trading is sufficiently low, there is positive probability that she never engages in research. In Section 4.2, we study the expected total amount of time that the trader devotes to research, and establish novel predictions that are not easily

captured by standard models of static acquisition. Finally, Section 4.3 properties of the information arrival time (i.e., the time that research is successful, if ever) and the likelihood that the trader ever becomes informed, and Section ?? outlines some testable implications of the model and highlights connections to the empirical literature on research and informed trading.

4.1 Time to first research

We begin with a characterization of τ_R , which is the first time that the trader initiates research.

Proposition 3. Let $\tau_R = \inf \{t \in [0, T) : v_t \geq v^*\}$ be the first time that the trader initiates research, with $\tau_R = \infty$ denoting no research. Supposing that $v_0 < v^*$ so that the trader does not begin research immediately, the probability density of τ_R is

$$\mathbb{P}(\tau_R \in dt) = e^{-rt} \frac{\log(v^*/v_0)}{\sigma_v \sqrt{2\pi t^3}} \exp \left\{ -\frac{1}{2t} \frac{(\log(v^*/v_0) - (\mu_v - \frac{1}{2}\sigma_v^2)t)^2}{\sigma_v^2} \right\}. \quad (26)$$

For any initial trading volume v_0 , the probability that the trader ever engages in research is

$$\mathbb{P}(\tau_R < \infty) = \begin{cases} \left(\frac{v_0}{v^*}\right)^{\beta_1} & 0 \leq v_0 < v^* \\ 1 & v_0 \geq v^* \end{cases}. \quad (27)$$

In standard, static models of information acquisition, the trader engages in research with probability one or zero. In contrast, unless initial trading volume is sufficiently high to trigger immediate investment in research (i.e., $v_0 \geq v^*$), the trader engages in research with a probability that is strictly between zero and one in our dynamic setting. We illustrate the implications of the above result in Figure 4, which plots the probability that the trader ever conducts research as a function of underlying parameters. Panel A plots the probability as a function of the initial level of noise trading v_0 for different values of maximal research intensity $\bar{\zeta}$, while panel B plots the probability as a function of disclosure frequency r .

Panel A shows that the likelihood of initiating research increases with initial trading volume v_0 . This is intuitive: the higher is v_0 , the more likely noise trade is to reach the research threshold v^* . Panel B shows that the probability of initiating research is hump shaped in r . The research probability inherits the nonmonotonicity of the research threshold v^* because of the “trading horizon” and “impatience” effects. For low disclosure frequencies, the impatience effect dominates and increases in r increase the probability that the trader conducts research, while for high values of r , the trading horizon effect dominates and further increases in r reduce the probability that she conducts research. These results stand in contrast to those from a static model, in which only the trading horizon effect operates, and increases in disclosure frequency r

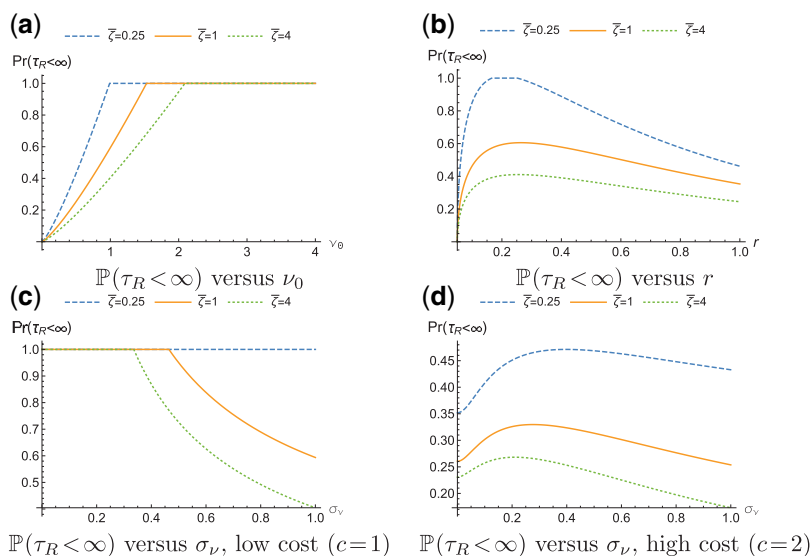


Figure 4
Probability that the trader conducts research, $\mathbb{P}(\tau_R < \infty)$

The figure plots the unconditional probability that the trader ever conducts research as a function of initial noise trading (i.e., ν_0), public disclosure frequency (i.e., r) and maximum research intensity (i.e., $\bar{\zeta}$). Unless specified, other parameters are set to $\mu_\nu=0.05$, $\Psi_0=1$, $c=C=1$, $r=0.2$, $\bar{\zeta}=1$, and $\sigma_\nu=0.1$.

always weakly decrease the probability of information acquisition. Intuitively, this is the “crowding out” effect of public information on private information acquisition and research.

Panels C and D plot the likelihood of research as a function of the volatility of uninformed trading volume, σ_ν , for both low and high costs of research. There are two effects of increasing σ_ν . First, it increases the threshold ν^* , as discussed above, which tends to reduce the likelihood of conducting research. We refer to this as a “threshold” effect. However, when the volatility of ν_t increases, it increases the likelihood that trading volume hits any given level, which tends to increase the likelihood of research. We refer to this as the “volatility” effect. As can be seen in panel C, for sufficiently low cost, the volatility effect dominates, which generates an unambiguously negative relation between σ_ν and the probability of research. However, for sufficiently high cost, the threshold effect dominates for low σ_ν and the volatility effect dominates for high σ_ν , which generates a hump-shaped relation with the likelihood of research.²⁵ These patterns are in sharp contrast to research in a static setting, where neither effect is present, and information acquisition does not depend on the volatility of trading volume.

²⁵ A similar tension between the “threshold” and “drift” effects arises when one studies changes in the drift, μ_ν , of trading volume.

Finally, note that all four panels imply that, holding other parameters fixed, higher research intensity can decrease the likelihood of engaging in research. This follows from the behavior of the research threshold v^* as a function of $\bar{\zeta}$: increases in $\bar{\zeta}$ increase the threshold, which decreases the probability that noise trade ever hits the threshold and triggers research activity. Note that these implications are distinctive consequence of the dynamic nature of research.

4.2 Expected time conducting research

Next, we study the amount of time that the trader expects to devote to research. Define $R(v) = \mathbb{E} \left[\int_t^{\min\{\tau, T\}} \mathbf{1}_{\{v_s \geq v^*\}} ds \mid v_t = v \right]$ as expected amount of time devoted to research from time t onward. Because of time stationarity, this quantity does not explicitly depend on t , only the value of v_t . The following Proposition characterizes the function R as a function of the current level of uninformed trading volume.

Proposition 4. Suppose that information has not yet arrived and the asset value has not yet been revealed, $t < \min\{\tau, T\}$. Then, the expected amount of time devoted to research from time- t onward is

$$R(v) = \begin{cases} \frac{\beta_2}{\beta_2 - \beta_1} \frac{1}{r + \bar{\zeta}} \left(\frac{v}{v^*} \right)^{\beta_1} & 0 \leq v < v^* \\ \frac{1}{r + \bar{\zeta}} - \frac{\beta_1}{\beta_1 - \beta_2} \frac{1}{r + \bar{\zeta}} \left(\frac{v}{v^*} \right)^{\beta_2} & v \geq v^* \end{cases} \quad (28)$$

We can relate the unconditional expected research time to the probability of conducting research by writing

$$\begin{aligned} R(v_0) &= \mathbb{P}(\tau_R < \infty) \mathbb{E} \left[\int_t^{\min\{\tau, T\}} \mathbf{1}_{\{v_s \geq v^*\}} ds \mid \tau_R < \infty \right] \\ &= \mathbb{P}(\tau_R < \infty) \times \begin{cases} \frac{1}{r + \bar{\zeta}} \frac{\beta_2}{\beta_2 - \beta_1} & 0 \leq v_0 < v^* \\ \frac{1}{r + \bar{\zeta}} \left(1 - \frac{\beta_1}{\beta_1 - \beta_2} \left(\frac{v_0}{v^*} \right)^{\beta_2} \right) & v_0 \geq v^* \end{cases} \end{aligned}$$

Hence, the expected time devoted to research is the product of the probability of research and the expected time devoted to research conditional on conducting research. Figure 5 illustrates the unconditional expected time devoted to research as a function of key underlying parameters. Not surprisingly, the effects of parameters on the expected time in research closely track those on the likelihood of initiating research. Specifically, as the plots suggest, the expected time in research is increasing in the initial uninformed trading volume v_0 , hump shaped in disclosure frequency r , and decreasing in the quality of research technology as measured by the maximal research intensity $\bar{\zeta}$.

The implications for disclosure frequency and research intensity are distinctive implications of the dynamic nature of research. If the investor were restricted to make a one-shot, “research or not” decision at the beginning and forced to commit to conducting research until acquiring information, we




Figure 5
Expected time devoted to research, $R(v_0) = \mathbb{E} \left[\int_0^{\min\{\tau, T\}} 1_{\{v_s \geq v^*\}} ds \mid v_0 \right]$
 The figure plots the unconditional expected time devoted to research as a function of initial trading volume (i.e., v_0), public disclosure frequency (i.e., r) and maximum research intensity (i.e., $\bar{\zeta}$). Unless specified, other parameters are set to $\mu = 0.05$, $\Psi_0 = 1$, $c = C = 1$, $r = 0.2$, $\bar{\zeta} = 1$ and $\sigma_v = 0.1$.

would expect that expected time conducting research would be (a) decreasing in disclosure frequency and (b) decreasing in the maximal research intensity $\bar{\zeta}$. These results highlight the importance of accounting for the dual optionality embedded in the research decision.

4.3 Time to information arrival

We now characterize properties of the information arrival time τ in the following Proposition.

Proposition 5. The unconditional probability that the trader ever receives information is

$$\mathbb{P}(\tau < \infty) = \begin{cases} \frac{\bar{\zeta}}{r + \bar{\zeta}} \frac{\beta_2}{\beta_2 - \beta_1} \left(\frac{v_0}{v^*}\right)^{\beta_1} & 0 \leq v_0 < v^* \\ \frac{\bar{\zeta}}{r + \bar{\zeta}} \left(1 - \frac{\beta_1}{\beta_1 - \beta_2} \left(\frac{v_0}{v^*}\right)^{\beta_2}\right) & v_0 \geq v^* \end{cases} \quad (29)$$

Figure 6 illustrates the unconditional probability that information arrives. Information arrival is driven by two effects: the probability that the trader engages in research, and the probability that the trader receives information, conditional on ever engaging in research

$$\mathbb{P}(\tau < \infty) = \mathbb{P}(\tau_R < \infty) \mathbb{P}(\tau < \infty \mid \tau_R < \infty) \quad (30)$$

$$= \mathbb{P}(\tau_R < \infty) \times \begin{cases} \frac{\bar{\zeta}}{r + \bar{\zeta}} \frac{\beta_2}{\beta_2 - \beta_1} & 0 \leq v_0 < v^* \\ \frac{\bar{\zeta}}{r + \bar{\zeta}} \left(1 - \frac{\beta_1}{\beta_1 - \beta_2} \left(\frac{v_0}{v^*}\right)^{\beta_2}\right) & v_0 \geq v^* \end{cases} \quad (31)$$

Hence, the effects of parameter changes on the probability of information arrival are partially driven by the research probability, illustrated in Figure 4 above.

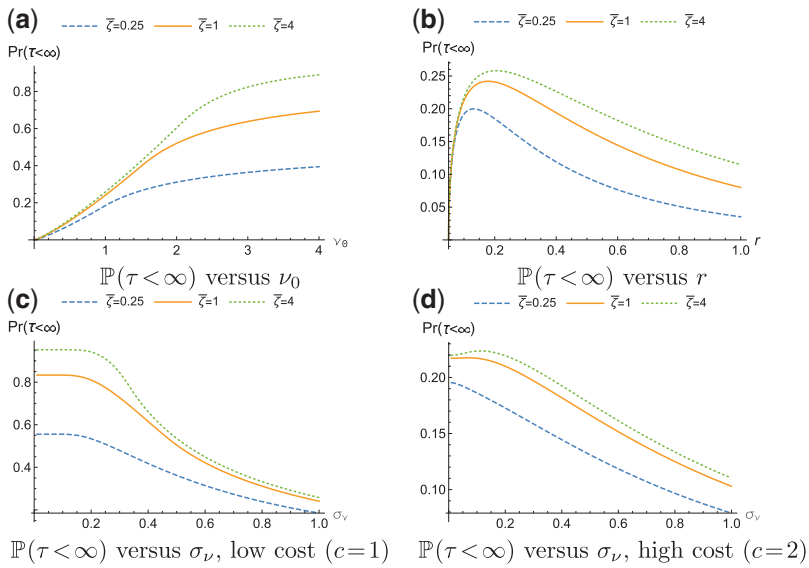


Figure 6

Probability that trader receives information, $\mathbb{P}(\tau < \infty)$

The figure plots the unconditional probability that the trader ever receives information as a function of initial trading volume (i.e., ν_0), public disclosure frequency (i.e., r), and maximum research intensity (i.e., $\bar{\zeta}$). Unless specified, other parameters are set to $\mu=0.05$, $\Psi_0=1$, $c=C=1$, $r=0.2$, $\bar{\zeta}=1$, and $\sigma_\nu=0.1$.

Differences in behavior from the probability of engaging in research are driven by the conditional probability of information arrival $\mathbb{P}(\tau < \infty | \tau_R < \infty)$.

Panel A plots the effect of changing the initial trading volume ν_0 . In this case, the conditional probability of information arrival, given research ν_0 , operates in the same direction (weakly) as the probability of research. If the trader is not initially engaged in research (i.e., $\nu_0 < \nu^*$), then the only effect of increasing ν_0 is to increase the probability that she ever conducts research, and the information arrival probability is increasing in ν_0 . As ν_0 increases above ν^* , the probability of ever conducting research is one, but the probability of remaining in the research region, and therefore of information arrival, strictly increases. The overall effect is an unambiguous positive relation between ν_0 and the probability of information arrival.

Panel B plots the probability of information arrival as a function of the disclosure frequency r . In all cases, the probability is hump shaped, an effect that is driven by the hump shape of the probability of ever conducting research. Indeed, one can demonstrate that the conditional probability of information arrival $\mathbb{P}(\tau < \infty | \tau_R < \infty)$ decreases in r in all cases. Intuitively, conditional on conducting research, an increase in disclosure frequency increases the probability that the payoff information will be publicly revealed before the trader's research activity bears fruit. However, the hump-shaped relation

between research probability and disclosure frequency dominates, and the information arrival probability is also hump shaped.

Panels C and D plot the probability of information arrival as a function of the volatility of trading volume σ_v . These effects are largely inherited from the effects of σ_v on the probability of research; however, there is an additional effect that, conditional on conducting research, more volatile trading volume tends to make it more likely that v_t will drop below the research threshold at times, which tends to reduce the probability of information arrival. Similarly to how σ_v affects the probability of research in Figure 4, for sufficiently low cost, there is an unambiguously negative relation between σ_v and the probability of information arrival. However, for sufficiently high cost, the fact that that highly volatile volume makes it more likely that the trader will abandon research can lead to a negative relation, even though the analogous relation between research activity and σ_v was hump shaped.

Finally, note that across the panels, the probability of information arrival increases with the maximum research intensity $\bar{\zeta}$. This is in sharp contrast to the earlier results, which implied that the likelihood of initiating research (and the expected time conducting research) were decreasing in $\bar{\zeta}$. Intuitively, the result is driven by the fact that conditional on conducting research, a higher research intensity increases the probability that information arrives. This effect is sufficiently strong to overcome the fact that the probability of conducting research $\mathbb{P}(\tau_R < \infty)$ is decreasing, as illustrated in Figure 4.

Combined with the observations from the previous sections, this implies that higher likelihood of research initiations and higher average time in research do not necessarily correspond to higher likelihood of information acquisition or informed trading. As we discuss next, this has potentially important implications for how we interpret existing evidence and suggests novel testable hypotheses.

4.4 Empirical predictions

Empirically testing our model predictions is challenging because it may be difficult for an econometrician to measure research by traders or detect market entry, even when these are detectable by other market participants.²⁶ A growing empirical literature suggests that one may be able to directly proxy for research activity using search activity or downloads of regulatory filings, for example, Bloomberg queries (Ben-Rephael, Da, and Israelsen 2017), online requests to the EDGAR system (e.g., DeHaan, Shevlin, and Thornock 2015; Loughran and McDonald 2017), and even Google or Yahoo Finance searches (e.g., Da, Engelberg, and Gao 2011;

²⁶ Importantly, note that our model and, consequently, the implications we will discuss below, presume only that other market participants can detect entry/increased participation by large investors, not necessarily that outside econometricians can do so. Our model makes no assumptions at all about the detectability of research activity. As we discussed in Section 2.1, market participants can often detect entry or increased participation by others (e.g., market makers, prime brokers, and OTC counterparties observe trading demand from institutional investors).

Drake, Roulstone, and Thornock 2012; Lawrence, Ryans, and Sun 2017). To the extent that increased demand for such public information is accompanied by private research activity (or essentially represents private information generation via more effective processing/interpretation of public releases), these measures provide a noisy proxy for research activity. Other, less directly visible research activity could include FOIA requests for regulatory filings (Gargano, Rossi, and Wermers 2016); increased access to corporate management through investor conferences (Bushee, Jung, and Miller 2011; Bushee, Jung, and Miller 2017), private analyst/investor days (Kirk and Markov 2016), or roadshows (Bushee, Gerakos, and Lee 2018); or even easier access to headquarters locations through flight or high-speed rail introductions (Ellis, Madureira, and Underwood 2020; Xu 2019). Finally, a number of recent papers (e.g., Gargano and Rossi 2018; Fedyk 2017) are able to directly link individual- investor-level proxies of news consumption to trading behavior.

One could similarly use regulatory filings to proxy for the entry and participation of large investors in the financial market. For instance, Schedule 13D filings can be used to identify trading by large investors who have acquired more than 5% of any class of securities of a publicly traded company (e.g., as in Collin-Dufresne and Fos 2015; Brav et al. 2008). Similarly, changes in the panel of quarterly Schedule 13F filings (required for any institution with at least \$100 million under management) can be used to estimate large position changes associated with information acquisition by large institutional investors such as hedge funds (e.g., Griffin and Xu 2009; Agarwal et al. 2013), and one can use Thomson and CRSP data to do the same for mutual funds (e.g., Wermers 2000). Further, by distinguishing between initiation of new positions and changes in existing holdings, such filings allow one to, at least partially, separate entry from changes in trading intensity.

Our model's predictions about research activity and information arrival are of one of two types. One category consists of predictions that are common to existing, static models of information acquisition. For instance, the likelihood and duration of research activity increases with uninformed trading volume and prior uncertainty, and decreases with the cost of conducting research. This implies research can be triggered by public news that increases uncertainty about the payoff of the risky asset or induces more trading by retail or other uninformed investors. Moreover, shocks that make research more costly, either directly or indirectly, are likely to reduce the incidence and duration of research. Given that much of the existing theoretical literature has focused on variants of static information acquisition, existing empirical work has largely focused on these types of predictions.

The other category of predictions is a consequence of dynamics and helps distinguish our analysis from earlier work. First, depending on the relative cost of research, the likelihood of research is either decreasing or hump shaped in the volatility of trading volume, after controlling for the level of trading volume

itself. This is a consequence of the trader's ability to wait to acquire information, and is absent from models of static information acquisition.

Second, sorting stocks on the frequency of (unscheduled) disclosures, should reveal a nonmonotonic relation with research activity: stocks with the fewest or the most disclosures per year should be associated with less research by investors than those in the middle. This is in contrast to settings with static acquisition, where a higher frequency of public disclosures should be associated with less research activity. Standard models predict that when public disclosures provide information about the same component of payoffs that traders can learn about, more public disclosure "crowds out" information acquisition. Our analysis characterizes conditions under which more frequent disclosures can "crowd in" information acquisition by changing the dynamic incentives to engage in research. Notably, this channel is distinct from the effect in settings in which investors can learn about multiple components of payoffs (e.g., Goldstein and Yang 2015; Banerjee, Davis, and Gondhi 2018; Goldstein and Yang 2019), and in which more public disclosure along one dimension can endogenously can "crowd in" information acquisition along another dimension. It is also different from the observation that in some cases, public disclosures can increase (and not decrease) uncertainty and, therefore trigger more research.²⁷

Third, a particularly stark prediction of our model is that higher likelihood or duration of research activity driven by changes in research technology ($\bar{\zeta}$) need not be associated with higher likelihood of informed trading. This is easily seen by contrasting Figure 4 (or Figure 5) with Figure 6. This implication helps us better understand the conflicting empirical results about the research activities of mutual funds and hedge funds. For instance, Kacperczyk and Seru (2007) suggest that for mutual funds, reliance on information from public sources tends to be associated with low performance, while Crane, Crotty, and Umar (2019) demonstrate the opposite is true for hedge funds. Our model can help reconcile these findings if we assume that mutual funds access to relatively less effective research technology (i.e., lower $\bar{\zeta}$ technology) relative to hedge funds. In this case, the model predicts that while mutual funds should be more likely to conduct research, it is less likely that their research will succeed. Similarly, Drake et al. (2020) find that research by sophisticated investors predicts future firm performance, but not research by less sophisticated investors. Overall, our model suggests caution when using measures of research activity to proxy for information acquisition. In a setting in which research takes place dynamically and need not succeed, simply observing research activity (e.g., EDGAR searches, Bloomberg queries) does not necessarily imply that traders are better informed. However, research activity followed by evidence of trading/information arrival does do so.

²⁷ In some models, more information can increase uncertainty (e.g., Veronesi (1999)) about fundamentals, and increased uncertainty could generate more information acquisition.

5. Price Informativeness and Market Liquidity

In this section, we turn our attention to the consequences of dynamic research for market quality, focusing on price informativeness and market liquidity. These are important quantities of interest to regulators, who have the ability to influence reporting frequency (i.e., r) and at least indirectly, the cost and effectiveness of conducting research (i.e., c and $\bar{\zeta}$) by altering disclosure rules.

We measure price informativeness by the fraction of total information that the market has learned about the asset by date t

$$PI_t = \frac{\Sigma_0 - \Sigma_t}{\Sigma_0} \tag{32}$$

and liquidity by price impact λ_t . We study two properties of price informativeness and liquidity. First, we characterize the dynamic behavior of PI_t and λ_t and how it depends on underlying parameters. While these properties are economically intuitive, they may be difficult to identify empirically from a regulator’s perspective. Hence, we also study a situation in which the regulator targets the average or “typical” PI and λ for firms. These quantities are (arguably) easier to measure and more consistent with policy decisions in practice, which are typically “unconditional” and without reference to specific firms or specific points in time. We characterize the behavior of these expected values, allowing for the possibility (but not requiring) that the regulator discounts the future.

The following proposition characterizes the dynamic behavior of the market quality variables over time.

Proposition 6. The evolution of price informativeness PI_t is given by

$$dPI_t = \begin{cases} 0 & 0 \leq t < \tau \\ \frac{\Psi_0}{\Psi_0 + \Omega} (1 + 2(r - \mu_v)e^{-2(r - \mu_v)(t - \tau)}) dt & \tau \leq t < T \end{cases}$$

and the evolution of price impact λ_t is given by

$$\frac{d\lambda_t}{\lambda_t} = \begin{cases} 0 & 0 \leq t < \tau \\ (\sigma_v^2 - r) dt - \sigma_v dW_{vt} & \tau \leq t < T. \end{cases}$$

The proposition has a number of implications. First, price informativeness is weakly increasing in time, and strictly increasing after the trader observes information. This is intuitive, and reflects the fact that more information is impounded into the price over time. Second, PI_t increases at a decreasing rate, which follows from the fact that a fixed amount of information can be revealed, and the rate at which it does so becomes small once the market has already learned a great deal. However, the rate of increase at each instant t may

be hump shaped in disclosure frequency r . Specifically, while the frequency of disclosure that maximizes the rate of increase in price informativeness is intermediate, it depends on firm specifics (i.e., μ_v) and on calendar time t , which makes it difficult to target from a policy perspective.

Third, the price impact is stochastic, and its drift depends on the difference between the variance of trading volume σ_v^2 and the disclosure frequency r . An increase in noise trading volatility ($\sigma_v^2 \uparrow$) or a decrease in disclosure frequency ($r \downarrow$) makes the trader more patient and creates an incentive to trade less aggressively. To sustain equilibrium, the resultant price impact must increase over time to give the trader an incentive to trade today.²⁸ Moreover, note that the drift in λ can either be positive or negative, depending on the relative magnitude of r and σ_v^2 . This is similar to the setting in Banerjee and Breon-Drish (2020), but in contrast to many existing models in the literature.²⁹ The model suggests that all else equal, more frequent disclosures are associated with declining price impact over time.

Finally, it is interesting to note that, given the information arrival time τ , the evolution of price informativeness or price impact do not depend on the cost of conducting research. While a decrease in the cost c makes research more attractive, it does not affect the rate at which information is impounded into prices once research is successful. As such, it may be difficult to evaluate the effectiveness of policies that change the cost of information acquisition by studying the impact on the dynamic behavior of price informativeness or price impact.

While the dynamics of price informativeness and price impact have economically intuitive interpretations, they may be difficult to target with policy and so not practical from a regulatory perspective. Arguably, it may be more natural to consider how the average values of these market quality variables respond to underlying parameters, and hence changes in regulatory policy.³⁰ With this view in mind, suppose that the regulator applies a discount rate $\rho \geq 0$ (the case of no discounting, $\rho = 0$ is allowed) to future PI_t and λ_t . We are interested in studying the average price informativeness and price impact, given by³¹

²⁸ More precisely, this property of λ follows from “trade timing indifference,” a property of continuous-time Kyle (1985) style frameworks that was first noted by Back (1992). That is, over any interval of time, the trader must always be indifferent between following her prescribed equilibrium trading strategy or refraining from trading over the interval and then following her prescribed strategy. Economically, timing indifference arises because, to sustain equilibrium, there cannot be any predictability in the level or slope of the price function if the trader refrains from trading. If there were, then she could deviate from her proposed equilibrium strategy by waiting and then exploit such predictability.

²⁹ For example, λ is constant in Kyle (1985), a martingale in Back (1992), a supermartingale in Back and Baruch (2004), and a submartingale in Collin-Dufresne and Fos (2016).

³⁰ This is also consistent with the treatment of market quality variables in one-shot models, in which the single value of, for example, PI and/or λ is typically interpreted as capturing an average or typical value.

³¹ It may seem more natural to use market depth $1/\lambda_t$ as our liquidity measure. However, because depth is infinite before the strategic trader enters the market, defining a summary measure of average or expected depth becomes difficult.

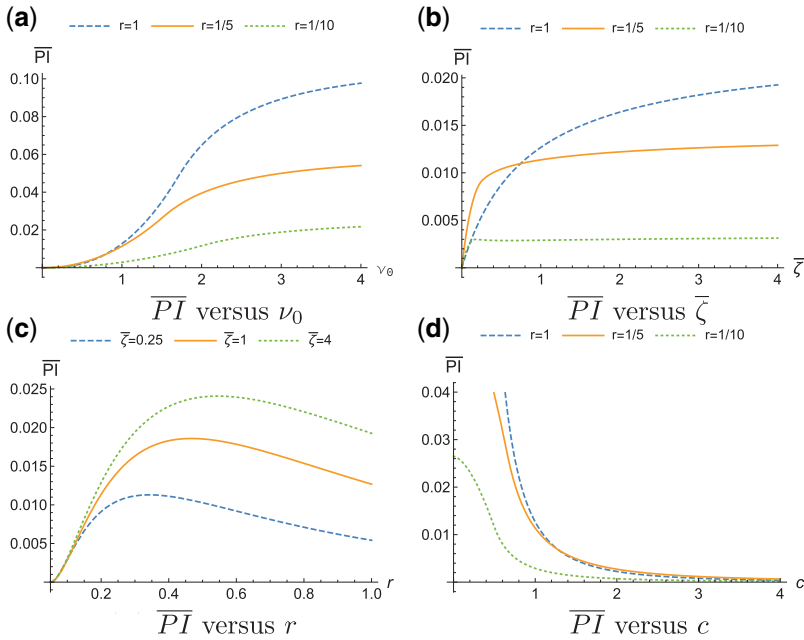


Figure 7

Expected discounted price informativeness, \overline{PI}

The figure plots the discounted, expected price informativeness \overline{PI} as a function of the initial trading volume (i.e., ν_0), public disclosure frequency (i.e., r), maximum research intensity (i.e., $\bar{\zeta}$), and flow cost of research (i.e., c). Unless specified, other parameters are set to $\mu=0.05$, $\Psi_0=1$, $\Omega=1/2$, $c=C=1$, $r=0.2$, $\sigma_v=0.1$, and $\rho=1$.

$$\overline{PI} = \mathbb{E} \left[(r + \rho) \int_0^T e^{-\rho s} P I_s ds \right], \tag{33}$$

$$\bar{\lambda} = \mathbb{E} \left[(r + \rho) \int_0^T e^{-\rho s} \lambda_s ds \right]. \tag{34}$$

Note that the multiplicative $r + \rho$ term in these expressions ensures that the discount factors, both explicit (because of ρ) and implicit (because of r), integrate to one so that these expressions behave as averages rather than sums (e.g., for a constant k , we have $\mathbb{E} \left[(r + \rho) \int_0^T e^{-\rho s} k ds \right] = k$). While both expressions are available in closed form, they are rather complicated, so we relegate the expressions to Proposition A.2 in the appendix and here focus on the key economic forces.

Figure 7 illustrates how price informativeness varies with key parameters. Note that we can decompose it into pre- and post-information arrival terms

$$\overline{PI} = \mathbb{E} \left[(r + \rho) \int_0^{\min(\tau, T)} e^{-\rho s} P I_s ds \right] + \mathbb{E} \left[(r + \rho) \int_{\min(\tau, T)}^T e^{-\rho s} P I_s ds \right], \tag{35}$$

$$= \mathbb{E} \left[(r + \rho) e^{-\rho \min(\tau, T)} \int_{\min(\tau, T)}^T e^{-\rho(s - \min(\tau, T))} P I_s ds \right], \tag{36}$$

where the first equality splits the integral at the information arrival time (or the disclosure date, if it occurs first), and the second equality uses the fact that price informativeness is zero prior to information arrival to set the first term to zero and rearranges the discount factor in the second term. This expression shows that two forces drive average price informativeness: (a) the arrival of information and its timing, captured by the discount factor $e^{-\rho \min(\tau, T)}$, and (b) the average price informativeness between information arrival and the disclosure date $\int_{\min(\tau, T)}^T e^{-\rho(s - \min(\tau, T))} P I_s ds$. Hence, (a) any change in parameters that increases the likelihood or speed of information arrival tends to increase the average price informativeness and (b) any change in parameters that incentivizes the trader to trade more aggressively when she is informed, and thereby speed up the incorporation of her information into prices, also tends to increase average price informativeness. Intuitively, as the trader becomes more impatient in either research or trading, price informativeness increases. This is illustrated in Figure 7, which plots \overline{PI} as a function of initial noise trade, disclosure frequency, maximum research intensity, and the flow cost of research. The effects of changes in these parameters is qualitatively similar to their effects on the probability of information arrival (see, e.g., Figure 6 above).

Figure 8 illustrates how expected discounted price impact varies with parameters. As with price informativeness, we can decompose $\overline{\lambda}$ into pre- and post-arrival terms and express it as

$$\overline{\lambda} = \mathbb{E} \left[(r + \rho) e^{-\rho \min(\tau, T)} \int_{\min(\tau, T)}^T e^{-\rho(s - \min(\tau, T))} \lambda_s ds \right], \tag{37}$$

which indicates that we can also interpret average price impact as due to the likelihood and timing of information arrival, and the average price impact conditional on arrival. However, changes in trader impatience can have opposite effects on these two terms, and so the overall effect on price impact can be more subtle. Consider panel A, which plots the average price impact as a function of the initial amount of noise trade. An increase in v_0 has two effects: (1) it makes it more likely that the trader conducts research and therefore that a signal arrives (and arrives quickly), which tends to increase average price impact, but (2) it increases the expected amount of volume during the post-arrival stage, which tends to reduce price impact. For low values of v_0 , the first effect dominates, while for high values of v_0 the second effect dominates, leading to a hump-shaped relation between $\overline{\lambda}$ and initial noise trade.

Similarly, consider Panel (b), which plots $\overline{\lambda}$ against the maximum research intensity $\overline{\zeta}$. As discussed above, an increase in $\overline{\zeta}$: (a) tends to make the trader wait longer to conduct research (and therefore tends to delay information arrival, all else equal), which tends to reduce average price impact, but (b) conditional

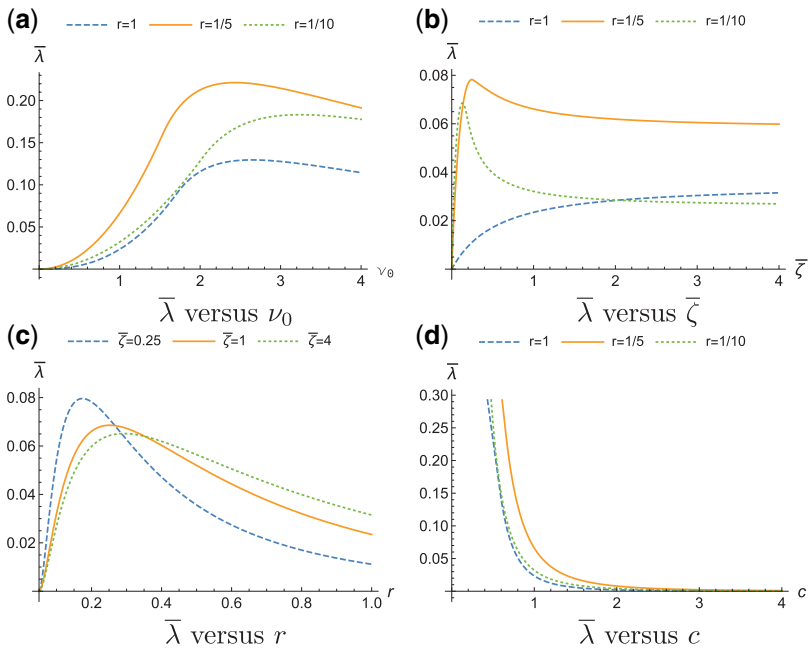


Figure 8

Expected discounted price impact, $\bar{\lambda}$

The figure plots the discounted, expected price impact $\bar{\lambda}$ as a function of the initial trading volume (i.e., ν_0), public disclosure frequency (i.e., r), maximum research intensity (i.e., $\bar{\zeta}$), and flow cost of research (i.e., c). Unless specified, other parameters are set to $\mu=0.05$, $\Psi_0=1$, $\Omega=1/2$, $c=C=1$, $r=0.2$, $\sigma_v=0.1$, and $\rho=1$.

on initiating research it tends to speed up information arrival, which tends to increase average price impact. For low values of $\bar{\zeta}$ the second effect dominates, and for high values of $\bar{\zeta}$ the first effect may (but need not) dominates. This produces either a hump-shaped or monotonically increasing relation between research intensity and price impact.

Note that for both price informativeness and price impact, the nonmonotonic relation with disclosure frequency r remains (panels C in Figures 7 and 8). This is again due to tension between the “trading horizon” and “impatience” effects. The impatience effect dominates for low values of r , while the trading horizon effect dominates for high values of r , leading to hump-shaped relations between disclosure frequency and price informativeness \overline{PI} and price impact $\bar{\lambda}$.

5.1 Policy implications

Our analysis of price informativeness and market liquidity above highlights a number of important implications from a policy perspective.

First, the nonmonotonic relation between the frequency of public disclosures (i.e., r) and price informativeness highlights a novel channel through which changes to disclosure requirements can affect price informativeness. As

discussed above, in settings in which the public disclosure pertains to the component of fundamentals that investors can learn about, the existing literature usually documents a “crowding out” effect, whereby more public disclosure decreases private information acquisition and lowers price informativeness. Our analysis highlights that this is a consequence of restricting attention to static information acquisition, and that the impact of public disclosures can be more nuanced in settings in which research is dynamic and information can be acquired over time.³² This result also sheds light on how regulatory changes that are intended to reduce information asymmetry by requiring more public disclosure (e.g., Regulation Fair Disclosure, or Reg FD, in 2000) may be counterproductive. For instance, our results are consistent with the evidence in Duarte et al. (2008), who document that the introduction of Reg FD was actually associated with higher likelihood of informed trading (as measured by the Easley, Hvidkjaer, and O’hara 2002 PIN measure), especially for NASDAQ stocks.

Second, note that unlike conventional wisdom, increasing the frequency of public disclosures (i.e., increasing r) or encouraging more trading (i.e., increasing ν_0 by making it easier for retail/uninformed investors to trade) need not improve liquidity. The standard intuition (from many static settings) is that greater public disclosure “levels the playing field” by reducing adverse selection between informed and uninformed investors, thereby improving liquidity. Similarly, increasing access for liquidity traders tends to lower price impact in Kyle (1985) and related settings. However, with endogenous research and dynamic information acquisition, these types of changes can have the opposite effect on liquidity, as suggested by panels A and C of Figure 8.

Finally, the analysis underscores the inherent tension between price informativeness and market liquidity that is common to other settings with asymmetrically informed traders. Specifically, consider a regulator who would like to improve both price informativeness and market liquidity. On the one hand, to improve price informativeness, the regulator would like to encourage research and information acquisition by the trader. On the other hand, doing so harms market liquidity since it increases the extent of adverse selection. As a result, regulatory changes that improve price informativeness will have a negative impact on market liquidity and vice versa.

This is generally true in our setting as well. Comparing Figures 7 and 8, note that changes in disclosure horizon (i.e., r) and the cost of research (i.e., c) change price informativeness and market impact in the same direction, and so these changes have opposite effects on informativeness and liquidity. However, this is not always the case for trading volume (i.e., ν_0) and the quality of research technology, as measured by the maximal research intensity $\bar{\zeta}$. Specifically,

³² For example, Figure 3 implies that when the trader is restricted to a static information choice, the acquisition threshold increases in the frequency of disclosure, r , which decreases the likelihood of acquisition and, consequently, price informativeness.

note that when v_0 is sufficiently high, further increases in trading volume can lead to higher price informativeness and higher liquidity (panel A in the two plots). Similarly, while improving research technology (i.e., making $\bar{\zeta}$ higher) always increases price informativeness (panel B in Figure 7), it can sometimes also improve liquidity by lowering the price impact (panel B of Figure 8). As discussed above, these features are a consequence of the dynamic nature of research and information acquisition, and so are likely to be missing from analysis that focuses on information acquisition as a static choice.³³ To the extent that research and information gathering are inherently dynamic activities, our analysis implies that policy changes that affect the level of retail trading volume or the ability of investors to conduct research can have more nuanced impact on price informativeness and liquidity than previous work suggests.

6. Concluding Remarks

We develop a model of research by a strategic trader that captures two important features: research is dynamic and probabilistic. As a result, the optimal choice of information acquisition embeds a number of real options: the option to delay research, the option to abandon research if unsuccessful, and the option to restart when trading opportunities improve. Relative to traditional settings in which information acquisition is a static choice, we show that incorporating these features has important consequences for our understanding of how markets produce and reflect information.

First, the likelihood and duration of conducting research is increasing in the volatility of trading volume, even after controlling for the level of volume. Second, an increase in the frequency of public disclosures can “crowd in” or “crowd out” private information acquisition: when the initial frequency is low, an increase encourages more research, but if the initial frequency is sufficiently high, it has the opposite effect. Finally, our analysis recommends caution in interpreting empirical evidence of research activity (e.g., Bloomberg search volume or EDGAR queries) as evidence of information acquisition, since research success is probabilistic. In fact, we show that when variation in research activity is driven by the effectiveness of traders’ research technology, observed research activity can be negatively related to likelihood of information acquisition.

Our analysis also highlights important trade-offs from a policy perspective. Contrary to the conventional wisdom (derived from static acquisition settings), we show that more frequent public disclosures can increase price

³³ Avdis and Banerjee (2019) derive a related result in the context of a multitrader, strategic trading model. They show that an increase in “clarity” (or a decrease in receiver specific noise) leads to higher correlation in the private information across traders, which leads to more aggressive trading via competition and, consequently, greater liquidity. While our result is complementary to theirs, the forces that underlie the positive relation between liquidity and price informativeness in our model are distinct because they are not driven by competition across traders (we have only one trader).

informativeness by encouraging research, even when such disclosures reflect information that traders can learn about. Second, we show that efforts to improve liquidity by “leveling the field” through more public disclosure or encouraging trading by uninformed investors can be counterproductive, because they can encourage research activity and, consequently, lead to higher equilibrium price impact. Finally, we show that unlike standard settings in which policy changes tend to have opposite effects on price informativeness and liquidity, improving the effectiveness of research or encouraging uninformed trading can lead to improvements along both dimensions in a dynamic setting.

Our paper is an early step in exploring the implications of research and information acquisition dynamics on trading market behavior, and suggests a number of natural avenues for future research. A natural direction is to explore other specifications of research and information acquisition technologies. While we believe that our model captures important, economically relevant features of research in markets, investors have access to many different types of information sources and research technologies. For instance, in a companion paper (Banerjee and Breon-Drish, 2020), we characterize how information acquisition and trading vary in a setting in which the trader can choose the precision of a flow of information. In that setting, the optimal precision varies stochastically in response to uninformed trading volume, but price impact and market uncertainty are unaffected by the choice of precision. It would be interesting to explore the effect of richer information acquisition technologies on equilibrium market dynamics.

One could also explore alternative assumptions about the timing and nature of public disclosures (i.e., about the terminal date, T). We model disclosure as an unscheduled public announcement, but one could also entertain the notion of a fixed disclosure time corresponding to, for example, an earnings announcement or other scheduled disclosure.³⁴ While we expect similar tension between “trading horizon” and “impatience” effects in such a setting, a full exploration could potentially provide richer insights. It would also be interesting to consider the problem of a regulator who endogenously chooses a disclosure policy to maximize some social objective function, or a firm manager who chooses a firm-level disclosure policy.

It also would be interesting to explore how our results change if the amount of fundamental uncertainty, the precision of private information, and the cost of conducting research depend on the trading horizon.³⁵ One might expect that less frequent, longer horizon disclosures are associated with more uncertainty about fundamentals. A natural parametrization to capture this feature is to scale the variance of fundamentals and errors (i.e., Σ_0 and σ_ε^2) by $1/r$ and the flow

³⁴ As we discussed above in Section 2, a model with a fixed terminal date T will generally be intractable analytically, since the optimal research problem effectively reduces to an optimal stopping problem on a finite horizon.

³⁵ We thank an anonymous referee for this suggestion.

cost of research (i.e., c) by $\sqrt{1/r}$ (i.e., with the volatility of fundamentals).³⁶ In this case, then our findings remain qualitatively unchanged and most of the formal results are literally identical.³⁷ This suggests that trading-horizon-dependent parameters will not change our basic conclusions, at least for plausible parameterizations. However, we leave a full exploration of such issues to future work.

Another important extension would be to consider competition among multiple strategic traders. In a recent paper, Xiong and Yang (2020) explore the implications of endogenous precision choice in a static model with multiple strategic traders. We expect the results from a dynamic model with imperfect competition to share many qualitative features of the current analysis, as long as private information is not perfectly correlated across traders (see Foster and Viswanathan [1993]; Back, Cao, and Willard [2000] for the effect of imperfect competition in strategic trading settings), while generating a set of novel implications for strategic research dynamics. However, extending the analysis to a dynamic setting is quite challenging, and so we leave it for future work.

Appendix A. Proofs

This appendix collects together all of the formal results that are not presented in the main text.

A.1. A Preliminary Result

“Discounted expectations” of the form $\mathbb{E}\left[e^{-uT_J}(a+bX_{T_J})|X_t=x\right]$ for various Poisson jump times T_J , constants a, b and u , and where X_t is a geometric Brownian motion, play a role in a number of our key results. In this section, we derive a general expression for expectations of this form to streamline later proofs.

Proposition A.1. Fix $x^* \in \mathbb{R}$ and let T_J be the first jump time of a doubly stochastic Poisson process with piecewise constant intensity

$$\zeta(X_t) = c_0 + c_L \mathbf{1}_{\{X_t < x^*\}} + c_H \mathbf{1}_{\{X_t \geq x^*\}}, \tag{A1}$$

³⁶ The linear scaling for fundamentals and errors can be motivated by the fact that variances scale linearly with time when uncertainty is modeled as a constant-coefficient diffusion process. Hence, if one considers V as arising from an (unmodeled) diffusion process, the linear scaling is natural. Whether a scaling for c is appropriate is ultimately a question of whether it is an economically reasonable order of magnitude for dependence on r . Scaling with the square root appears to be economically reasonable because a Kyle trader’s gross expected profit upon receipt of information scales with fundamental volatility (not variance) (e.g., in our model, we have $\bar{J} = \sqrt{\frac{\Psi_0}{2(r-\mu_V)}} v_t$ earlier. See also the discussion of the value of information in Banerjee and Breon-Drish 2020), and it is effectively this profit to which she compares costs when making research decisions.

³⁷ More generally, consider an arbitrary, increasing scaling function $f(1/r)$, where we specify $V \sim N(0, \Sigma_0 f(1/r))$, $\varepsilon \sim N(0, \sigma_\varepsilon^2 f(1/r))$ and $C(\zeta) = \zeta c \sqrt{f(1/r)}$. Then it follows immediately that the optimal research threshold v^* is identical to the one in Proposition 2, and as a result, the characterization of research dynamics in Section 4 are identical to our benchmark model. Similarly, the average price informativeness \bar{PI} is also identical in this setting. Moreover, we can show that $\bar{\lambda}$, while not identical, is still a hump-shaped function of r . The reason is that when fundamental uncertainty increases with horizon, it has the effect of weakening the standard “trading horizon” effect (since, in the event of information arrival, price impact is higher when the information has a high prior variance) but leaving the novel “impatience effect” unchanged. This increases the relative importance of the impatience effect and preserves the hump-shaped dependence on r .

where the process X_t is a geometric Brownian motion

$$dX_t = \mu X_t dt + \sigma X_t dW_t. \tag{A2}$$

Then, for any $a, b \in \mathbb{R}$ and $u \geq 0$, we have

$$\mathbb{E} \left[e^{-uT_J} (a + bX_{T_J}) \mid X_0 = x \right] = \begin{cases} \frac{c_0 + c_L}{c_0 + c_L + u} a + \frac{c_0 + c_L}{c_0 + c_L + u - \mu} bx + A_L x^{\gamma_L} & 0 \leq x < x^* \\ \frac{c_0 + c_H}{c_0 + c_H + u} a + \frac{c_0 + c_H}{c_0 + c_H + u - \mu} bx + B_H x^{\gamma_H} & x \geq x^* \end{cases}, \tag{A3}$$

where

$$A_L \equiv \frac{c_H - c_L}{\gamma_H - \gamma_L} \left(\frac{u\gamma_H}{(c_0 + c_L + u)(c_0 + c_H + u)} a + \frac{(\gamma_H - 1)(u - \mu)}{(c_0 + c_L + u - \mu)(c_0 + c_H + u - \mu)} bx^* \right) \left(\frac{1}{x^*} \right)^{\gamma_L}, \tag{A4}$$

$$B_H \equiv \frac{c_H - c_L}{\gamma_H - \gamma_L} \left(\frac{u\gamma_L}{(c_0 + c_L + u)(c_0 + c_H + u)} a + \frac{(\gamma_L - 1)(u - \mu)}{(c_0 + c_L + u - \mu)(c_0 + c_H + u - \mu)} bx^* \right) \left(\frac{1}{x^*} \right)^{\gamma_H}, \tag{A5}$$

$$\gamma_L = \frac{1}{2} - \frac{\mu}{\sigma^2} + \sqrt{\frac{2(c_0 + c_L + u)}{\sigma^2} + \left(\frac{1}{2} - \frac{\mu}{\sigma^2} \right)^2}, \text{ and} \tag{A6}$$

$$\gamma_H = \frac{1}{2} - \frac{\mu}{\sigma^2} - \sqrt{\frac{2(c_0 + c_H + u)}{\sigma^2} + \left(\frac{1}{2} - \frac{\mu}{\sigma^2} \right)^2}. \tag{A7}$$

Proof of Proposition A.1. We would like to compute $\mathbb{E} [e^{-uT_J} (a + bX_{T_J})]$. Note that we can express it as

$$\mathbb{E} [e^{-uT_J} (a + bX_{T_J})] = \mathbb{E} \left[\mathbb{E} \left[e^{-uT_J} (a + bX_{T_J}) \mid \{X_t\}_{t \geq 0} \right] \right] \tag{A8}$$

$$= \mathbb{E} \left[\int_0^\infty \mathbb{P}(T_J \in ds \mid \{X_t\}_{t \geq 0}) e^{-us} (a + bX_s) \right] \tag{A9}$$

$$= \mathbb{E} \left[\int_0^\infty \zeta(X_s) e^{-\int_0^s \zeta(X_r) dr} e^{-us} (a + bX_s) ds \right] \tag{A10}$$

$$= \mathbb{E} \left[\int_0^\infty e^{-\int_0^s (u + \zeta(X_r)) dr} \zeta(X_s) (a + bX_s) ds \right] \tag{A11}$$

where the first equality uses the law of iterated expectations and conditions on the entire history of $\{X_t\}$; the second equality uses the fact that conditional on the history of $\{X_t\}$ the only random variable is T_J , and computing the conditional expectation therefore reduces to integrating over T_J ; the third equality uses the fact that conditional on $\{X_t\}$, the path of the intensity is $\{\zeta_t\}$ is known and therefore $\mathbb{P}(T_J \geq s \mid \{X_t\}) = e^{-\int_0^s \zeta(X_r) dr}$ from which $\mathbb{P}(T_J \in ds \mid \{X_t\}) = \zeta(X_s) e^{-\int_0^s \zeta(X_r) dr} ds$; and the final equality gathers terms. Define the function

$$f(t, x) = \mathbb{E} \left[\int_t^\infty e^{-\int_t^s (u + \zeta(X_r)) dr} \zeta(X_s) (a + bX_s) ds \mid X_t = x \right]. \tag{A12}$$

The Feynman-Kac theorem yields a differential equation that f must follow, which, owing to the time stationarity of X_t and time homogeneity of $\zeta(\cdot)$, is an ordinary differential equation (ODE) in x

$$0 = \zeta(x)(a + bx) - (\zeta(x) + u)f + \mu x f_x + \frac{1}{2} \sigma^2 x^2 f_{xx}. \tag{A13}$$

That is, there is no explicit time dependence in the function f .³⁸ Because the “killing rate” $u + \zeta(x)$ is discontinuous at x^* , the solution to this equation will, in general, be only continuously differentiable at the point x^* and twice continuously differentiable elsewhere.³⁹ We now proceed with the solution to the ODE. It naturally splits into two regions, depending on whether $x \geq x^*$.

$$0 = \begin{cases} (c_0 + c_L)(a + bx) - (c_0 + c_L + u)f + \mu x f_x + \frac{1}{2} \sigma^2 x^2 f_{xx} & x < x^* \\ (c_0 + c_H)(a + bx) - (c_0 + c_H + u)f + \mu x f_x + \frac{1}{2} \sigma^2 x^2 f_{xx} & x \geq x^* \end{cases} \quad (A14)$$

We can solve the ODE by first characterizing the solution within each region and then ensuring that the resultant function is sufficiently well-behaved across the boundary. Within each region the ODE is linear, so the solution is equal to the sum of any particular solution and the general solution to the homogenous version of the equation. Hence, we proceed by characterizing particular and general solutions in each region, imposing the boundary conditions to eliminate certain terms, and then conclude by imposing that the overall solution must be continuously differentiable across the boundary x^* . It is easy to verify that particular solutions in each region are

$$\begin{cases} \frac{c_0 + c_L}{c_0 + c_L + u} a + \frac{c_0 + c_L}{c_0 + c_L + u - \mu} bx & x < x^* \\ \frac{c_0 + c_H}{c_0 + c_H + u} a + \frac{c_0 + c_H}{c_0 + c_H + u - \mu} bx & x \geq x^* \end{cases} \quad (A15)$$

and general solutions are of the form

$$\begin{cases} A_L x^{\gamma_{L+}} + B_L x^{\gamma_{L-}} & 0 \leq x < x^* \\ A_H x^{\gamma_{H+}} + B_H x^{\gamma_{H-}} & x \geq x^* \end{cases} \quad (A16)$$

where

$$\gamma_{L\pm} = \frac{1}{2} - \frac{\mu}{\sigma^2} \pm \sqrt{\frac{2(c_0 + c_L + u)}{\sigma^2} + \left(\frac{1}{2} - \frac{\mu}{\sigma^2}\right)^2} \quad (A17)$$

$$\gamma_{H\pm} = \frac{1}{2} - \frac{\mu}{\sigma^2} \pm \sqrt{\frac{2(c_0 + c_H + u)}{\sigma^2} + \left(\frac{1}{2} - \frac{\mu}{\sigma^2}\right)^2} \quad (A18)$$

The overall solution is therefore of the form

$$f = \begin{cases} A_L x^{\gamma_{L+}} + B_L x^{\gamma_{L-}} + \frac{c_0 + c_L}{c_0 + c_L + u} a + \frac{c_0 + c_L}{c_0 + c_L + u - \mu} bx & 0 \leq x < x^* \\ A_H x^{\gamma_{H+}} + B_H x^{\gamma_{H-}} + \frac{c_0 + c_H}{c_0 + c_H + u} a + \frac{c_0 + c_H}{c_0 + c_H + u - \mu} bx & x \geq x^* \end{cases} \quad (A19)$$

where the A 's and B 's are constants to be determined. Natural boundary conditions for the ODE are

$$\lim_{x \rightarrow 0} f = \lim_{x \rightarrow 0} \left(\frac{c_0 + c_L}{c_0 + c_L + u} a + \frac{c_0 + c_L}{c_0 + c_L + u - \mu} bx \right) = \frac{c_0 + c_L}{c_0 + c_L + u} a \quad (A20)$$

$$\lim_{x \rightarrow \infty} f_x = \lim_{x \rightarrow \infty} \frac{\partial}{\partial x} \left(\frac{c_0 + c_H}{c_0 + c_H + u} a + \frac{c_0 + c_H}{c_0 + c_H + u - \mu} bx \right) = \frac{c_0 + c_H}{c_0 + c_H + u - \mu} b \quad (A21)$$

³⁸ Note that by reversing the steps above, it is easily seen that this function characterizes $\mathbb{E} \left[e^{-u(T_J - t)} (a + bX_{T_J}) \middle| X_t = x \right]$. That is, the discounted conditional expectation, where we discount back only to time t . Because of time stationarity of X_t and the fact that T_J is the jump time of a Poisson process, this expression does not depend on calendar time.

³⁹ See, for example, theorem 4.9 in chap. 4 of Karatzas and Shreve (1998). The theorem establishes this for the case in which X_t is a standard Brownian motion (and to which one can reduce our case by appropriate transformations of the X_t process), or the general treatment of Glau (2016).

since in each limit it becomes arbitrarily unlikely that X_t will ever again transition to the other region. Hence, the discounted conditional expectation should converge to that for a constant arrival intensity.⁴⁰ The boundary conditions (A20) and (A21) imply $B_L = 0$ and $A_H = 0$. Next, the solution should be continuously differentiable across x^* , which gives two equations to pin down A_L and B_H

$$A_L x^{\gamma_L+} + \frac{c_0+c_L}{c_0+c_L+u} a + \frac{c_0+c_L}{c_0+c_L+u-\mu} b x = B_H x^{\gamma_H-} + \frac{c_0+c_H}{c_0+c_H+u} a + \frac{c_0+c_H}{c_0+c_H+u-\mu} b x \tag{A23}$$

$$\gamma_L+ A_L x^{\gamma_L+ - 1} + \frac{c_0+c_L}{c_0+c_L+u-\mu} b = \gamma_H- B_H x^{\gamma_H- - 1} + \frac{c_0+c_H}{c_0+c_H+u-\mu} b \tag{A24}$$

Solving this system yields

$$A_L = \frac{c_H - c_L}{\gamma_H - \gamma_L+} \left(\frac{u\gamma_H-}{(c_0+c_L+u)(c_0+c_H+u)} a + \frac{(\gamma_H- - 1)(u-\mu)}{(c_0+c_L+u-\mu)(c_0+c_H+u-\mu)} b x^* \right) (x^*)^{-\gamma_L+} \tag{A25}$$

$$B_H = \frac{c_H - c_L}{\gamma_H - \gamma_L+} \left(\frac{u\gamma_L+}{(c_0+c_L+u)(c_0+c_H+u)} a + \frac{(\gamma_L+ - 1)(u-\mu)}{(c_0+c_L+u-\mu)(c_0+c_H+u-\mu)} b x^* \right) (x^*)^{-\gamma_H-} \tag{A26}$$

and letting $\gamma_H = \gamma_H-$ and $\gamma_L = \gamma_L-$ to condense notation gives us the result. ■

A.2. Proofs of the Main Results

Proof of Proposition 1. The post-entry financial market equilibrium is essentially a special case of the “flow signals” setting of Banerjee and Breon-Drish (2020). Nominally, their result applies to settings in which (a) the trader can dynamically optimize a flow of signals after observing her initial signal and (b) both the trader and market maker learn the true value at the limit when the economy continues without end ($T \rightarrow \infty$). To apply their results, (a) let the cost of ongoing flow signals tend to infinity $c(\cdot) \uparrow \infty$, which implies that the trader’s only information source is her initial signal, and (b) use risk neutrality and the fact that the market maker has coarser information than the trader to redefine the asset value V as the conditional expected value given observation of the signal $\hat{V} = \mathbb{E}[V|S]$. The equilibrium now immediately follows with the market maker’s conditional variance of \hat{V} , $\Psi_t = \text{var}(\hat{V}|\mathcal{F}_t^P)$, taking the place of her variance of the asset value itself. ■

Proof of Proposition 2. Let $\bar{J}(\cdot) = K v_t$, with $K \equiv \sqrt{\frac{\psi_0}{r-\mu v}}$, be the unconditional expected profit given information arrival at time t . Given an arbitrary strategy $\{\zeta_s\}$ that is adapted to $\sigma(\{v_t\})$, her overall expected profit is

$$\mathbb{E} \left[\mathbf{1}_{\{\tau < \infty\}} \bar{J}(v_\tau) - \int_0^{\min\{\tau, T\}} c(\zeta_s) ds \right]. \tag{A27}$$

We would like to express this in a more standard form. To do this, it is helpful to “delink” the information arrival time τ and the public disclosure date T . Let T_J be the first jump time of a Poisson process with the posited intensity ζ_s for all $s \geq 0$ (i.e., the arrival time if one were to follow the strategy $\{\zeta_s\}$ regardless of whether the disclosure has occurred). We can express

$$\mathbb{E}[\mathbf{1}_{\{\tau < \infty\}} \bar{J}(v_\tau)] = \mathbb{E}[\mathbf{1}_{\{T_J < T\}} \bar{J}(v_{T_J})] \tag{A28}$$

⁴⁰ In the case of a constant intensity c the time T_J is independently, exponentially distributed with rate c , and by integrating out T_J the expectation that we desire to compute reduces to

$$\mathbb{E} \left[e^{-uT_J} (a + bX_{T_J}) \right] = \mathbb{E} \left[\int_0^\infty c e^{-cs} e^{-us} (a + bX_s) ds \right] = \int_0^\infty \left(c e^{-(c+u)s} a + c e^{-(c+u)s} b X_0 \right) ds, \tag{A29}$$

which is a sum of two growing perpetuities. Computing the values of these perpetuities produces the expressions in the text.

$$= \mathbb{E} \left[e^{-rT_I} \bar{J}(v_{T_I}) \right] \tag{A29}$$

$$= \mathbb{E} \left[\int_0^\infty e^{-\int_0^s (r+\zeta u) du} \zeta_s \bar{J}(v_s) ds \right], \tag{A30}$$

where the second equality uses the fact that T is independently exponentially distributed, and the last equality follows from steps analogous to those used in the derivation of Equation (A11) above. Similarly,

$$\mathbb{E} \left[\int_0^{\min\{\tau, T\}} c(\zeta_s) ds \right] = \mathbb{E} \left[\int_0^{\min\{T_I, T\}} c(\zeta_s) ds \right] \tag{A31}$$

$$= \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{s \leq \min\{T_I, T\}\}} c(\zeta_s) ds \right] \tag{A32}$$

$$= \mathbb{E} \left[\int_0^\infty e^{-\int_0^s (r+\zeta u) du} c(\zeta_s) ds \right], \tag{A33}$$

where the second equality writes the limits of integration using indicator function notation and the final equality follows from steps analogous to those used in the derivation of Equation (A11) above, using the fact that $\min\{T_I, T\}$ is the first jump time of the sum of two independent Poisson processes, one with intensity ζ that tracks information arrival and one with intensity r that tracks the public disclosure. Since the research strategy was arbitrary the above expressions allow us to state the trader’s research problem in a compact, standard form

$$J^U \equiv \sup_{\{\zeta_s\}} \mathbb{E} \left[\int_0^\infty e^{-\int_0^s (r+\zeta u) du} (\zeta_s \bar{J}(v_s) - c(\zeta_s)) ds \right]. \tag{A34}$$

We now proceed to characterize the optimal research strategy. We begin by making use of the usual HJB equation to derive a candidate optimal research strategy and the corresponding value function. We then verify that the candidate strategy and value function so constructed do in fact characterize an optimum. With linear cost function $c(\zeta) = c\zeta$ and constraint $\zeta \in [0, \bar{\zeta}]$, the HJB equation for this problem is

$$0 = \sup_{\zeta \in [0, \bar{\zeta}]} \left\{ r(0 - J^U) + \zeta(\bar{J} - J^U) + J_v^U \mu_v v + \frac{1}{2} J_{vv}^U \sigma_v^2 v^2 - c\zeta \right\}. \tag{A35}$$

We conjecture that the optimal research strategy follows a threshold rule. That is, there exists some $v^* \geq 0$, to be determined, such that

$$\zeta^*(v) = \begin{cases} 0 & 0 \leq v < v^* \\ \bar{\zeta} & v \geq v^* \end{cases} \tag{A36}$$

We further conjecture that the value function takes the time homogenous form of

$$J^U(v) = \begin{cases} Av^{\beta_1} & 0 \leq v < v^* \\ Bv^{\beta_2} + \frac{\bar{\zeta}}{r - \mu_v + \bar{\zeta}} K v - \frac{c\bar{\zeta}}{r + \bar{\zeta}} & v \geq v^* \end{cases} \tag{A37}$$

for constants A and B to be determined, and where β_1 and β_2 are defined in eqs. (18) and (19). It remains to characterize the candidate threshold v^* and constants A and B . To do this, we enforce that the candidate value function should be *twice* continuously differentiable across the threshold v^* . Equating the level and first derivative are analogous to the standard value-matching and smooth-pasting conditions of optimal stopping. Furthermore, because the state can transition freely back across the threshold v^* (i.e., this is effectively a *switching* problem, not a *stopping* problem), one also requires a “high-contact” condition that the second derivative be continuous. We emphasize

that these conditions, while intuitive, are simply conjectures for how the value function should behave, and we later verify that the function so constructed is in fact optimal. We have

$$Av^{\beta_1} = Bv^{\beta_2} + \frac{\bar{\zeta}}{r - \mu_v + \bar{\zeta}} K v - \frac{c\bar{\zeta}}{r + \bar{\zeta}} \quad (\text{value matching}) \quad (A38)$$

$$\beta_1 Av^{\beta_1-1} = B\beta_2 v^{\beta_2-1} + \frac{\bar{\zeta}}{r - \mu_v + \bar{\zeta}} K \quad (\text{smooth pasting}) \quad (A39)$$

$$\beta_1(\beta_1 - 1)Av^{\beta_1-2} = B\beta_2(\beta_2 - 1)v^{\beta_2-2} \quad (\text{high contact}) \quad (A40)$$

Solving for the three unknowns yields

$$v^* = \frac{\beta_1}{\beta_1 - 1} \frac{\beta_2}{\beta_2 - 1} \frac{c}{K} \frac{r + \bar{\zeta} - \mu_v}{r + \bar{\zeta}} \quad (A41)$$

$$A = \frac{c\bar{\zeta}}{r + \bar{\zeta}} \frac{\beta_2}{(\beta_1 - 1)(\beta_2 - \beta_1)} \frac{1}{(v^*)^{\beta_1}} \quad (A42)$$

$$B = -\frac{c\bar{\zeta}}{r + \bar{\zeta}} \frac{\beta_1}{(\beta_2 - 1)(\beta_1 - \beta_2)} \frac{1}{(v^*)^{\beta_2}}. \quad (A43)$$

We can now check that the conjectured research intensity in fact achieves the maximum on the right-hand side of Equation (A35) and the conjectured J^U solves the HJB equation. To confirm that the conjectured ζ^* maximizes the right-hand side of Equation (A35), it suffices to show that $\bar{J} - J^U - c$ is strictly increasing in v and is equal to zero at v^* . Given the constraint on ζ , it will then follow immediately that the posited threshold strategy achieves the maximum. For $v < v^*$, we have

$$\frac{\partial}{\partial v}(\bar{J} - J^U - c) = K - A\beta_1 v^{\beta_1-1} \quad (A44)$$

$$= K - \frac{c\bar{\zeta}}{r + \bar{\zeta}} \frac{\beta_1}{\beta_1 - 1} \frac{\beta_2}{\beta_2 - \beta_1} \frac{1}{(v^*)^{\beta_1}} v^{\beta_1-1} \quad (A45)$$

$$> K - \frac{c\bar{\zeta}}{r + \bar{\zeta}} \frac{\beta_1}{\beta_1 - 1} \frac{\beta_2}{\beta_2 - \beta_1} \frac{1}{(v^*)^{\beta_1}} (v^*)^{\beta_1-1} \quad (A46)$$

$$= K - \frac{c\bar{\zeta}}{r + \bar{\zeta}} \frac{\beta_1}{\beta_1 - 1} \frac{\beta_2}{\beta_2 - \beta_1} \frac{1}{v^*} \quad (A47)$$

$$= K - K \frac{\bar{\zeta}}{r + \bar{\zeta} - \mu_v} \quad (A48)$$

$$> 0 \quad (A49)$$

and analogous calculations hold for $v \geq v^*$. Straightforward, but tedious, algebra also verifies that $\bar{J} - J^U - c = 0$ at v^* . Hence, the conjectured research strategy maximizes the right-hand side of (A35). Given the research strategy, we can also verify that the conjectured value function satisfies the HJB equation (A35). For instance, for $v < v^*$ we have

$$r(0 - J^U) + \zeta^*(v)(\bar{J} - J^U) + J_v^U + J_v^U \mu_v v + \frac{1}{2} J_{vv}^U \sigma_v^2 v^2 - c\zeta^*(v) \quad (A50)$$

$$= -rJ^U + J_v^U \mu_v v + \frac{1}{2} J_{vv}^U \sigma_v^2 v^2 \quad (A51)$$

$$= -rAv^{\beta_1} + A\beta_1v^{\beta_1-1}\mu_vv + \frac{1}{2}A\beta_1(\beta_1-1)v^{\beta_1-2}\sigma_v^2v^2 \tag{A52}$$

$$= Av^{\beta_1} \left(-r + \mu_v\beta_1 + \frac{1}{2}\sigma_v^2\beta_1(\beta_1-1) \right) \tag{A53}$$

$$= 0 \tag{A54}$$

where the last line follows from the fact that β_1 solves $0 = -r + \mu_v\beta_1 + \frac{1}{2}\sigma_v^2\beta_1(\beta_1-1)$ by construction. A similar calculation confirms that the equation holds for $v \geq v^*$. To verify that the conjectured research strategy is, in fact, optimal, consider an arbitrary strategy $\{\hat{\xi}_s\} \in [0, \bar{\xi}]$ and apply Ito's Lemma to the function $e^{-\int_0^t (r+\hat{\xi}_u)du} J^U(v_t)$

$$e^{-\int_0^t (r+\hat{\xi}_u)du} J^U(v_t) - J^U(v_0) \tag{A55}$$

$$= \int_0^t e^{-\int_0^s (r+\hat{\xi}_u)du} \left(-(r+\hat{\xi}_s)J^U + J_v^U \mu_v v + \frac{1}{2} J_{vv}^U \sigma_v^2 v^2 \right) ds \tag{A56}$$

$$+ \int_0^t e^{-\int_0^s (r+\hat{\xi}_u)du} J_v^U v \sigma_v dW_{v_s} \tag{A57}$$

$$\leq - \int_0^t e^{-\int_0^s (r+\hat{\xi}_u)du} \hat{\xi}_s (\bar{J} - c) ds + \int_0^t e^{-\int_0^s (r+\hat{\xi}_u)du} J_v^U v \sigma_v dW_{v_s} \tag{A58}$$

where the inequality follows from Equation (A35) and the fact that the strategy $\hat{\xi}_u$ is arbitrary. Rearranging and using the fact that $J^U(v_t) \geq 0$ we have

$$\int_0^t e^{-\int_0^s (r+\hat{\xi}_u)du} \hat{\xi}_s (\bar{J} - c) ds \leq J^U(v_0) + \int_0^t e^{-\int_0^s (r+\hat{\xi}_u)du} J_v^U v \sigma_v dW_{v_s} \tag{A59}$$

Because $\left| e^{-\int_0^t (r+\hat{\xi}_u)du} \hat{\xi}_s (\bar{J} - c) \right| \leq e^{-rs} \bar{\xi} (Kv_s + c)$ and $\mathbb{E} \left[\int_0^\infty e^{-rs} \bar{\xi} (Kv_s + c) ds \right] < \infty$ (because of $\mu_v < r$), the dominated convergence theorem implies

$$\mathbb{E} \left[\int_0^\infty e^{-\int_0^s (r+\hat{\xi}_u)du} \hat{\xi}_s (\bar{J} - c) ds \right] = \lim_{t \rightarrow \infty} \mathbb{E} \left[\int_0^t e^{-\int_0^s (r+\hat{\xi}_u)du} \hat{\xi}_s (\bar{J} - c) ds \right] \tag{A60}$$

Combined with Equation (A59), this implies

$$\mathbb{E} \left[\int_0^\infty e^{-\int_0^s (r+\hat{\xi}_u)du} \hat{\xi}_s (\bar{J} - c) ds \right] \leq J^U(v_0) + \lim_{t \rightarrow \infty} \mathbb{E} \left[\int_0^t e^{-\int_0^s (r+\hat{\xi}_u)du} J_v^U v \sigma_v dW_{v_s} \right] \tag{A61}$$

$$= J^U(v_0) \tag{A62}$$

where the final equality follows from the fact that for each $t \geq 0$ we have

$$\mathbb{E} \left[\int_0^t \left(e^{-\int_0^s (r+\hat{\xi}_u)du} J_v^U v_s \sigma_v \right)^2 ds \right] < \mathbb{E} \left[\int_0^t e^{-2rs} (Kv_s \sigma_v)^2 ds \right] < \infty \tag{A63}$$

and therefore the stochastic integral is a martingale (see, e.g., corollary 3 in chap. II.6 of Protter 2003). Owing to the representation for the expected trading profits of an arbitrary strategy in Equation (A34), it follows from Equation (A62) that our candidate value function J^U is an upper bound for expected profit for the trader under any arbitrary strategy $\{\hat{\xi}_t\} \in [0, \bar{\xi}]$ and associated arrival time τ

$$J^U(v_0) \geq \mathbb{E} \left[\mathbf{1}_{\{\tau < \infty\}} \bar{J}(v_\tau) - \int_0^{\min\{\tau, T\}} c(\hat{\xi}_s) ds \right]. \tag{A64}$$

It remains to show that the candidate strategy yields expected profit equal to $J^U(v_0)$. Return to Equation (A58), in which the inequality now becomes an equality because our candidate strategy attains the maximum in the HJB equation. Rearranging, we have

$$\int_0^t e^{-\int_0^s (r+\xi_u)du} \zeta_s (\bar{J} - c) ds = J^U(v_0) + \int_0^t e^{-\int_0^s (r+\xi_u)du} J^U_v v \sigma_v dW_{v_s} - e^{-\int_0^t (r+\xi_u)du} J^U(v_t) \tag{A65}$$

Again, taking expectations and considering $t \rightarrow \infty$, to demonstrate the optimality of the strategy, it suffices to show $\limsup_{t \rightarrow \infty} \mathbb{E} \left[e^{-\int_0^t (r+\xi_u)du} J^U(v_t) \right] = 0$ since $J^U \geq 0$. However, this follows immediately since

$$e^{-\int_0^t (r+\xi_u)du} J^U(v_t) \leq e^{-rt} \bar{J} = e^{-rt} K v_t \tag{A66}$$

where $J^U \leq \bar{J}$ follows from the fact that $\bar{J} - J^U$ is strictly increasing, as shown above, and $\bar{J}(0) - J^U(0) = 0$. Therefore,

$$\limsup_{t \rightarrow \infty} \mathbb{E} \left[e^{-\int_0^t (r+\xi_u)du} J^U(v_t) \right] \leq \limsup_{t \rightarrow \infty} K v_0^{(\mu_v - r)t} = 0 \tag{A67}$$

since $\mu_v < r$, which establishes the optimality of the posited research strategy. ■

Proof of Proposition 3. Research occurs the first time that v hits v^* from below, as long as this occurs before time T when the asset payoff is revealed. We begin by characterizing the distribution of the first hitting time $T_R = \inf\{t \geq 0 : v_t \geq v^*\}$ and then account for the possibility that payoff is publicly revealed before this time. Note that

$$v_t \geq v^* \Leftrightarrow \log v_t \geq \log v^* \tag{A68}$$

$$\Leftrightarrow \log(v_0) + \left(\mu_v - \frac{1}{2} \sigma_v^2 \right) t + \sigma_v W_{v_t} \geq \log(v^*) \tag{A69}$$

$$\Leftrightarrow \frac{1}{\sigma_v} \left(\mu_v - \frac{1}{2} \sigma_v^2 \right) t + W_{v_t} \geq \frac{1}{\sigma_v} \log(v^*/v_0) \tag{A70}$$

so that it is equivalent to find the first time that a standard Brownian motion with drift $\frac{1}{\sigma_v} (\mu_v - \frac{1}{2} \sigma_v^2)$ hits $\frac{1}{\sigma_v} \log(v^*/v_0)$. It follows from Karatzas and Shreve (1998, p. 197, equation (5.12)) that

$$\mathbb{P}(T_R \in dt) = \frac{\left| \frac{1}{\sigma_v} \log(v^*/v_0) \right|}{\sqrt{2\pi t^3}} \exp \left\{ -\frac{1}{2t} \left(\frac{1}{\sigma_v} \log(v^*/v_0) - \frac{1}{\sigma_v} \left(\mu_v - \frac{1}{2} \sigma_v^2 \right) t \right)^2 \right\} \tag{A71}$$

$$= \frac{\log(v^*/v_0)}{\sigma_v \sqrt{2\pi t^3}} \exp \left\{ -\frac{1}{2t} \frac{(\log(v^*/v_0) - (\mu_v - \frac{1}{2} \sigma_v^2) t)^2}{\sigma_v^2} \right\}, \tag{A72}$$

where the second equality rearranges some terms and uses the fact that $v^* \geq v_0$ by assumption, so $\log(v^*/v_0) \geq 0$. Furthermore, from Karatzas and Shreve (1998, p. 197, equation (5.13)) or by integrating the density above, we have

$$\mathbb{P}(T_R < \infty) = \exp \left\{ \frac{1}{\sigma_v^2} \left(\mu_v - \frac{1}{2} \sigma_v^2 \right) \log(v^*/v_0) - \left| \frac{1}{\sigma_v^2} \left(\mu_v - \frac{1}{2} \sigma_v^2 \right) \log(v^*/v_0) \right| \right\} \tag{A73}$$

$$= \exp \left\{ \frac{1}{\sigma_v^2} \left(\left(\mu_v - \frac{1}{2} \sigma_v^2 \right) - \left| \mu_v - \frac{1}{2} \sigma_v^2 \right| \right) \log(v^*/v_0) \right\} \tag{A74}$$

$$= (v^*/v_0)^{\frac{1}{\sigma_v^2} \left((\mu_v - \frac{1}{2} \sigma_v^2) - \left| \mu_v - \frac{1}{2} \sigma_v^2 \right| \right)} \tag{A75}$$

$$= \begin{cases} (v^*/v_0) \frac{2}{\sigma_v^2} (\mu v - \frac{1}{2} \sigma_v^2) & \mu v - \frac{1}{2} \sigma_v^2 < 0 \\ 1 & \mu v - \frac{1}{2} \sigma_v^2 \geq 0 \end{cases} \tag{A76}$$

Now, accounting for the possibility that the payoff is publicly revealed before v_t reaches v^* , we can write the first research time as

$$\tau_R = T_R \mathbf{1}_{\{T_R \leq T\}} + \infty \times \mathbf{1}_{\{T_R > T\}} \tag{A77}$$

Hence, for $t \in [0, \infty)$,

$$\mathbb{P}(\tau_R \in dt) = \mathbb{P}(T_R \in dt, T \geq T_R) \tag{A78}$$

$$= \mathbb{P}(T \geq T_R | T_R \in dt) \mathbb{P}(T_R \in dt) \tag{A79}$$

$$= e^{-rt} \mathbb{P}(T_R \in dt) \tag{A80}$$

which delivers the expression in the proposition after substituting for $\mathbb{P}(T_R \in dt)$ from above. Similarly,

$$\mathbb{P}(\tau_R < \infty) = \mathbb{P}(T > T_R) = \int_0^\infty \mathbb{P}(T > s) \mathbb{P}(T_R \in ds) = \int_0^\infty e^{-rs} \mathbb{P}(T_R \in ds), \tag{A81}$$

and the expression in the proposition now follows from straightforward, but tedious, integration. Note that because

$$\mathbb{P}(\tau_R < \infty) = \mathbb{P}(T_R < T) = \mathbb{E}[\mathbb{P}(T_R < T | T_R)] = \mathbb{E}[e^{-rT_R}], \tag{A82}$$

one can also derive the expression for $\mathbb{P}(\tau_R < \infty)$ as a limiting case of Proposition A.1, where $a = 1$, $b = 0$, $u = r$, and the underlying stochastic process $X_t \equiv v_t$. To do this, set $x^* = v^*$, and allow the arrival intensity to become arbitrarily large for $v_t \geq v^*$ and equal zero elsewhere: $c_0 = c_L = 0$ and $c_H \rightarrow \infty$. ■

Proof of Proposition 4. Let $\zeta_s = \bar{\zeta} \mathbf{1}_{\{v_s \geq v^*\}}$ denote the optimal research strategy, and T_I be the jump time for a Poisson process with this intensity for all $s \geq 0$ (i.e., if one were to follow the strategy regardless of the disclosure date T). The amount of time devoted to research can be expressed as

$$R(v) = \mathbb{E} \left[\int_0^{\min\{T_I, T\}} \mathbf{1}_{\{v_s \geq v^*\}} ds \right] \tag{A83}$$

$$= \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{s \leq \min\{T_I, T\}\}} \mathbf{1}_{\{v_s \geq v^*\}} ds \right] \tag{A84}$$

$$= \mathbb{E} \left[\mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{s \leq \min\{T_I, T\}\}} \mathbf{1}_{\{v_s \geq v^*\}} ds \mid \{v_t\}_{t \geq 0} \right] \right] \tag{A85}$$

$$= \mathbb{E} \left[\int_0^\infty e^{-\int_0^s (r + \zeta_u) du} \mathbf{1}_{\{v_s \geq v^*\}} ds \right] \tag{A86}$$

$$= \frac{1}{\bar{\zeta}} \mathbb{E} \left[\int_0^\infty e^{-\int_0^s (r + \zeta_u) du} \zeta_s ds \right], \tag{A87}$$

where the first equality is the definition of expected research time, the second equality writes the upper limit of the integral using indicator function notation, and the third equality uses the law of iterated expectations and conditions on the entire history of $\{v_t\}$. The fourth equality uses the fact that $\min\{T, T_I\}$ is the first jump time of the sum of two independent Poisson processes, one with intensity ζ that tracks the information arrival date and one with intensity r that tracks the disclosure date. The final equality simply multiplies and divides by $\bar{\zeta}$ and uses the definition of the research intensity process ζ_s . We can now evaluate the expectation as a special case of Proposition A.1 with $a = 1$, $b = 0$, $u = r$, and the underlying stochastic process $X_t \equiv v_t$. To do this, set $x^* = v^*$, and set the intensity parameters as $c_0 = c_L = 0$ and $c_H = \bar{\zeta}$. ■

Proof of Proposition 5. As in the previous proof, let $\zeta_s = \bar{\zeta} \mathbf{1}_{\{v_s \geq v^*\}}$ denote the optimal research strategy, and T_I be the jump time for a Poisson process with this intensity for all $s \geq 0$ (i.e., if one were to follow the strategy regardless of the disclosure date T). The probability of information arrival can be expressed as

$$\mathbb{P}(\tau < \infty) = \mathbb{P}(T_I < T) \tag{A88}$$

$$= \mathbb{E}[\mathbb{P}(T_I < T | T_I)] \tag{A89}$$

$$= \mathbb{E}[e^{-rT_I}]. \tag{A90}$$

The explicit expression in the statement of the proposition now follows as a special case of Proposition A.1 with $a=1, b=0, u=r$, and the underlying stochastic process $X_t \equiv v_t$. To do this, set $x^* = v^*$, and set the intensity parameters as $c_0 = c_L = 0$ and $c_H = \bar{\zeta}$. ■

Proof of Proposition 6. Note that

$$P I_t = \frac{\Sigma_0 - \Sigma_t}{\Sigma_0} = \frac{(\Psi_0 + \Omega) - (\Psi_t + \Omega)}{\Psi_0 + \Omega} = \frac{\Psi_0 - \Psi_t}{\Psi_0 + \Omega}.$$

The result now follows from applying Ito's lemma to the expressions for Ψ_t and λ_t from Proposition 1. ■

A.3. Expressions for Price Informativeness and Market Liquidity

This subsection derives expressions for discounted expected price informativeness and price impact, and for the dynamics of price informativeness and price impact, that are used for the plots in Section 5. As in earlier proofs, we let $\zeta_s = \bar{\zeta} \mathbf{1}_{\{v_s \geq v^*\}}$ denote the optimal research strategy and let T_I denote the first jump time of a Poisson process with intensity ζ_s for all $s \geq 0$. We formalize the results in the following Proposition:

Proposition A.2. The discounted, expected price informativeness is:

$$\bar{P I} \equiv \mathbb{E} \left[(r + \rho) \int_0^T e^{-\rho s} P I_s ds \right] \tag{A91}$$

$$= \frac{\Psi_0}{\Psi_0 + \Omega} \frac{2(r - \mu_v)}{r + \rho + 2(r - \mu_v)} \mathbb{E} \left[e^{-(r+\rho)T_I} \right], \tag{A92}$$

where $\mathbb{E} \left[e^{-(r+\rho)T_I} \right]$ can be evaluated as a special case of Proposition A.1 with $a=1, b=0, u=r+\rho$ and $X_t \equiv v_t$, with $x^* = v^*, c_0 = c_L = 0$, and $c_H = \bar{\zeta}$. And the discounted, expected price impact is

$$\bar{\lambda} \equiv \mathbb{E} \left[(r + \rho) \int_0^T e^{-\rho s} \lambda_s ds \right] \tag{A93}$$

$$= \sqrt{2(r - \mu_v)} \frac{\Psi_0}{2r + \rho - \sigma_v^2} \mathbb{E} \left[e^{-(r+\rho)T_I} \frac{1}{v_{T_I}} \right], \tag{A94}$$

where $\mathbb{E} \left[e^{-(r+\rho)T_I} \frac{1}{v_{T_I}} \right]$ can be evaluated as a special case of Proposition A.1 with $a=0, b=1, u=r+\rho$ and $X_t \equiv \frac{1}{v_t}$, with $x^* = \frac{1}{v^*}, c_0 = c_H = 0$, and $c_L = \bar{\zeta}$.

Proof of Proposition A.2. Consider first price informativeness. Note that prior to information arrival, $t \leq \min\{T_I, T\}$, we have $P I_t = 0$, so it suffices to compute

$$\mathbb{E} \left[(r + \rho) \int_0^T e^{-\rho s} P I_s ds \right] = (r + \rho) \mathbb{E} \left[\mathbf{1}_{\{T_I \leq T\}} \int_{T_I}^T e^{-\rho s} P I_s ds \right] \tag{A95}$$

$$= (r + \rho) \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{T_I \leq T\}} \mathbf{1}_{\{T_I \leq s\}} \mathbf{1}_{\{s \leq T\}} e^{-\rho s} P I_s ds \right] \tag{A96}$$

$$= (r + \rho) \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{T_I \leq s\}} \mathbf{1}_{\{s \leq T\}} e^{-\rho s} P I_s ds \right] \tag{A97}$$

$$= (r + \rho) \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{T_I \leq s\}} e^{-(r+\rho)s} P I_s ds \right] \tag{A98}$$

$$= (r + \rho) \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{T_I \leq s\}} e^{-(r+\rho)s} P I_s ds \right] \tag{A99}$$

Recall that $P I_s = \frac{\Sigma_0 - \Sigma_s}{\Sigma_0}$ and

$$\Sigma_t = \Omega + \Psi_t \tag{A100}$$

$$= \begin{cases} \Omega + \Psi_0 & 0 \leq t < \tau \\ \Omega + e^{-2(r-\mu_v)(t-\tau)} \Psi_0 & t \geq \tau \end{cases} \tag{A101}$$

$$= \begin{cases} \Omega + \Psi_0 & 0 \leq t < \min\{T, T_I\} \\ \Omega + e^{-2(r-\mu_v)(t-T_I)} \Psi_0 & t \geq T_I, \text{ if } T_I \leq T \end{cases} \tag{A102}$$

Hence,

$$(r + \rho) \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{T_I \leq s\}} e^{-(r+\rho)s} P I_s ds \right] \tag{A103}$$

$$= (r + \rho) \frac{\Psi_0}{\Sigma_0} \mathbb{E} \left[\int_{T_I}^\infty e^{-(r+\rho)s} \left(1 - e^{-2(r-\mu_v)(s-T_I)} \right) ds \right] \tag{A104}$$

$$= (r + \rho) \frac{\Psi_0}{\Sigma_0} \mathbb{E} \left[e^{-(r+\rho)T_I} \int_{T_I}^\infty e^{-(r+\rho)(s-T_I)} \left(1 - e^{-2(r-\mu_v)(s-T_I)} \right) ds \right] \tag{A105}$$

$$= (r + \rho) \frac{\Psi_0}{\Sigma_0} \mathbb{E} \left[e^{-(r+\rho)T_I} \frac{2(r-\mu_v)}{(r+\rho)(r+\rho+2(r-\mu_v))} \right] \tag{A106}$$

$$= \frac{\Psi_0}{\Sigma_0} \frac{2(r-\mu_v)}{r+\rho+2(r-\mu_v)} \mathbb{E} \left[e^{-(r+\rho)T_I} \right] \tag{A107}$$

Consider now discounted price impact $\mathbb{E} \left[(r + \rho) \int_0^T e^{-\rho s} \lambda_s ds \right]$. Recall that

$$\lambda_t = \begin{cases} 0 & 0 \leq t < \tau \\ e^{-(r-\mu_v)(t-\tau)} \sqrt{\frac{2(r-\mu_v)\Psi_0}{v_t^2}} & t \geq \tau \end{cases} \tag{A108}$$

$$= \begin{cases} 0 & 0 \leq t < \min\{T, T_I\} \\ e^{-(r-\mu_v)(t-T_I)} \sqrt{\frac{2(r-\mu_v)\Psi_0}{v_t^2}} & t \geq T_I, \text{ if } T_I \leq T \end{cases} \tag{A109}$$

We have

$$\mathbb{E} \left[\int_0^T e^{-\rho s} \lambda_s ds \right] = \mathbb{E} \left[\mathbf{1}_{\{T_I \leq T\}} \int_{T_I}^T e^{-\rho s} \lambda_s ds \right] \tag{A110}$$

$$= \sqrt{2(r - \mu_v)} \Psi_0 \mathbb{E} \left[\mathbf{1}_{\{T_I \leq T\}} \int_{T_I}^T e^{-\rho s} e^{-(r - \mu_v)(s - T_I)} \frac{1}{v_s} ds \right] \tag{A111}$$

$$= \sqrt{2(r - \mu_v)} \Psi_0 \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{T_I \leq T\}} \mathbf{1}_{\{T_I \leq s\}} \mathbf{1}_{\{s \leq T\}} e^{-\rho s} e^{-(r - \mu_v)(s - T_I)} \frac{1}{v_s} ds \right] \tag{A112}$$

$$= \sqrt{2(r - \mu_v)} \Psi_0 \mathbb{E} \left[\int_0^\infty \mathbf{1}_{\{T_I \leq s\}} e^{-(r + \rho)s} e^{-(r - \mu_v)(s - T_I)} \frac{1}{v_s} ds \right] \tag{A113}$$

$$= \sqrt{2(r - \mu_v)} \Psi_0 \mathbb{E} \left[e^{-(r + \rho)T_I} \int_{T_I}^\infty e^{-(r + \rho + (r - \mu_v)(s - T_I))} \frac{1}{v_s} ds \right]. \tag{A114}$$

The process $\frac{1}{v_t}$ is a geometric Brownian motion

$$d\left(\frac{1}{v_t}\right) = (\sigma_v^2 - \mu_v) \frac{1}{v_t} dt - \sigma_v \frac{1}{v_t} dW_{v,t}, \tag{A115}$$

so, for $s \geq T_I$,

$$\frac{1}{v_s} = \frac{1}{v_{T_I}} e^{(\sigma_v^2 - \mu_v - \frac{1}{2}\sigma_v^2)(s - T_I) - \sigma_v(W_{v,s} - W_{v,T_I})}, \tag{A116}$$

and therefore

$$\mathbb{E} \left[e^{-(r + \rho)T_I} \int_{T_I}^\infty e^{-(r + \rho + (r - \mu_v)(s - T_I))} \frac{1}{v_s} ds \right] \tag{A117}$$

$$= \mathbb{E} \left[e^{-(r + \rho)T_I} \frac{1}{v_{T_I}} \int_{T_I}^\infty e^{-(r + \rho + (r - \mu_v)(s - T_I))} e^{(\sigma_v^2 - \mu_v)(s - T_I)} ds \right] \tag{A118}$$

$$= \frac{1}{2r + \rho - \sigma_v^2} \mathbb{E} \left[e^{-(r + \rho)T_I} \frac{1}{v_{T_I}} \right]. \tag{A119}$$

■

References

Agarwal, V., W. Jiang, Y. Tang, and B. Yang. 2013. Uncovering hedge fund skill from the portfolio holdings they hide. *Journal of Finance* 68:739–83.

Avdis, E., and S. Banerjee. 2019. Clear and liquid: The interaction of firm disclosure and trader competition. Working Paper, University of Alberta.

Back, K. 1992. Insider trading in continuous time. *Review of Financial Studies* 5:387–409.

Back, K., and S. Baruch. 2004. Information in securities markets: Kyle meets Glosten and Milgrom. *Econometrica* 72:433–65.

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

Back, K., and H. Pedersen. 1998. Long-lived information and intraday patterns. *Journal of Financial Markets* 1:385–402.

Banerjee, S., and B. M. Breon-Drish. 2020. Strategic trading and unobservable information acquisition. *Journal of Financial Economics* 138:458–82.

- Banerjee, S., J. Davis, and N. Gondhi. 2018. When transparency improves, must prices reflect fundamentals better? *Review of Financial Studies* 31:2377–414.
- Ben-Rephael, A., Z. Da, and R. D. Israelsen. 2017. It depends on where you search: Institutional investor attention and underreaction to news. *Review of Financial Studies* 30:3009–47.
- Bond, P., A. Edmans, and I. Goldstein. 2012. The real effects of financial markets. *Annual Review of Financial Economics* 4:339–60.
- Brav, A., W. Jiang, F. Partnoy, and R. Thomas. 2008. Hedge fund activism, corporate governance, and firm performance. *Journal of Finance* 63:1729–75.
- Breon-Drish, B. 2015. On existence and uniqueness of equilibrium in a class of noisy rational expectations models. *Review of Economic Studies* 82:868–921.
- Bushee, B. J., J. Gerakos, and L. F. Lee. 2018. Corporate jets and private meetings with investors. *Journal of Accounting and Economics* 65:358–79.
- Bushee, B. J., M. J. Jung, and G. S. Miller. 2011. Conference presentations and the disclosure milieu. *Journal of Accounting Research* 49:1163–92.
- . 2017. Do investors benefit from selective access to management? *Journal of Financial Reporting* 2:31–61.
- Caldentey, R., and E. Stacchetti. 2010. Insider trading with a random deadline. *Econometrica* 78:245–83.
- Collin-Dufresne, P., and V. Fos. 2015. Do prices reveal the presence of informed trading? *Journal of Finance* 70:1555–82.
- . 2016. Insider trading, stochastic liquidity and equilibrium prices. *Econometrica* 84:1441–75.
- Colombo, L., G. Femminis, and A. Pavan. 2014. Information acquisition and welfare. *Review of Economic Studies* 81:1438–83.
- Crane, A., K. Crotty, and T. Umar. 2019. Hedge funds and public information acquisition. Working Paper, Rice University.
- Da, Z., J. Engelberg, and P. Gao. 2011. In search of attention. *Journal of Finance* 66:1461–99.
- DeHaan, E., T. J. Shevlin, and J. R. Thornock. 2015. Market (in) attention and the strategic scheduling and timing of earnings announcements. *Journal of Accounting and Economics* 60:36–55.
- Diamond, D. W. 1985. Optimal release of information by firms. *Journal of Finance* 40:1071–94.
- Drake, M. S., B. A. Johnson, D. T. Roulstone, and J. R. Thornock. 2020. Is there information content in information acquisition? *Accounting Review* 95:113–39.
- Drake, M. S., D. T. Roulstone, and J. R. Thornock. 2012. Investor information demand: Evidence from google searches around earnings announcements. *Journal of Accounting Research* 50:1001–40.
- Duarte, J., X. Han, J. Harford, and L. Young. 2008. Information asymmetry, information dissemination and the effect of regulation fd on the cost of capital. *Journal of Financial Economics* 87:24–44.
- Dugast, J., and T. Foucault. 2018. Data abundance and asset price informativeness. *Journal of Financial Economics* 130:367–91.
- Easley, D., S. Hvidkjaer, and M. O'Hara. 2002. Is information risk a determinant of asset returns? *Journal of Finance* 57:2185–221.
- Ellis, J., L. Madureira, and S. Underwood. 2020. The causal effects of proximity on investment: Evidence from flight introductions. *Journal of Financial and Quantitative Analysis* 55:1978–2004.
- Fedyk, A. 2017. Disagreement after news: Gradual information diffusion or differences of opinion? Working Paper, UC Berkeley.

- Foster, F. D., and S. Viswanathan. 1993. The effect of public information and competition on trading volume and price volatility. *Review of Financial Studies* 6:23–56.
- Gargano, A., and A. G. Rossi. 2018. Does it pay to pay attention? *Review of Financial Studies* 31:4595–649.
- Gargano, A., A. G. Rossi, and R. Wermers. 2016. The freedom of information act and the race toward information acquisition. *Review of Financial Studies* 30:2179–228.
- Glau, K. 2016. A Feynman–Kac-type formula for Lévy processes with discontinuous killing rates. *Finance and Stochastics* 20:1021–39.
- Goldstein, I., and H. Sapra. 2014. Should banks' stress test results be disclosed? an analysis of the costs and benefits. *Foundations and Trends in Finance* 8:1–54.
- Goldstein, I., and L. Yang. 2015. Information diversity and complementarities in trading and information acquisition. *Journal of Finance* 70:1723–65.
- . 2017. Information disclosure in financial markets. *Annual Review of Financial Economics* 9:101–25.
- . 2019. Good disclosure, bad disclosure. *Journal of Financial Economics* 131:118–38.
- Griffin, J. M., and J. Xu. 2009. How smart are the smart guys? a unique view from hedge fund stock holdings. *Review of Financial Studies* 22:2531–70.
- Grossman, S. J., and J. E. Stiglitz. 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70:393–408.
- Han, B. 2018. Dynamic information acquisition and asset prices. Working Paper, University of Maryland.
- Hellwig, M. F. 1980. On the aggregation of information in competitive markets. *Journal of Economic Theory* 22:477–98.
- Holden, C. W., and A. Subrahmanyam. 2002. News events, information acquisition, and serial correlation. *Journal of Business* 75:1–32.
- Huang, S., and B. Z. Yueshen. 2018. Speed acquisition. Working Paper, INSEAD.
- Kacperczyk, M., and A. Seru. 2007. Fund manager use of public information: New evidence on managerial skills. *Journal of Finance* 62:485–528.
- Karatzas, I., and S. E. Shreve. 1998. Brownian motion and stochastic calculus, 2nd ed. In *Graduate texts in mathematics*, vol. 113. New York: Springer.
- Kendall, C. 2018. The time cost of information in financial markets. *Journal of Economic Theory* 176:118–57.
- Kirk, M. P., and S. Markov. 2016. Come on over: Analyst/investor days as a disclosure medium. *Accounting Review* 91:1725–50.
- Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* 53:1315–35.
- Lawrence, A., J. P. Ryans, and E. Y. Sun. 2017. Investor demand for sell-side research. *Accounting Review* 92:123–49.
- Loughran, T., and B. McDonald. 2017. The use of edgar filings by investors. *Journal of Behavioral Finance* 18:231–48.
- Protter, P. E. 2003. Stochastic integration and differential equations, 2nd ed. In *Applications of mathematics*, vol. 21. New York: Springer.
- Shreve, S. E. 2004. *Stochastic calculus for finance II: Continuous-time models*. New York: Springer.
- Veldkamp, L. L. 2006. Media frenzies in markets for financial information. *American Economic Review* 96:577–601.
- Veronesi, P. 1999. Stock market overreactions to bad news in good times: a rational expectations equilibrium model. *Review of Financial Studies* 12:975–1007.

Verrecchia, R. E. 1982. Information acquisition in a noisy rational expectations economy. *Econometrica* 50:1415–30.

Wermers, R. 2000. Mutual fund performance: An empirical decomposition into stock-picking talent, style, transactions costs, and expenses. *Journal of Finance* 55:1655–95.

Xiong, Y., and L. Yang. 2020. Secret and overt information acquisition in financial markets. Working Paper, Rotman School of Management.

Xu, D. 2019. Costly information acquisition and investment decisions: Quasi-experimental evidence. Working Paper, University of Texas at Austin.