Dynamic Information Acquisition and Strategic Trading

Snehal Banerjee and Bradyn Breon-Drish*

December 2017

Abstract

Consider a strategic trader who dynamically chooses when to acquire costly information about an asset's payoff, instead of being endowed with this information. Whether the market maker observes acquisition is critical. Without observability, we show that an equilibrium with smooth trading and a pure acquisition strategy cannot exist. We also rule out the existence of a natural class of mixed-strategy equilibria. With observability, however, there exists an equilibrium in which optimal acquisition follows a pure strategy and generally exhibits delay. Our results suggest that many strategic trading equilibria considered in the literature are difficult to reconcile with dynamic information acquisition.

JEL: D82, D84, G12, G14

Keywords: Dynamic information acquisition, Strategic trading, Observability

*Banerjee (snehalb@ucsd.edu) and Breon-Drish (bbreondrish@ucsd.edu) are at the University of California, San Diego. All errors are our own. We thank Brett Green for numerous, invaluable discussions during an early stage of this project. We also thank Stathi Avdis, Kerry Back, Darrell Duffie, Joey Engelberg, Mariana Khapko, Igor Makarov, Seymon Malamud, Sophie Moinas, Dmitry Orlov, Christine Parlour, Chris Parsons, Uday Rajan, Allan Timmermann, Dimitri Vayanos, and participants at Collegio Carlo Alberto, UC Davis, LSE, UC Riverside, University of Colorado Boulder, the University of Minnesota Junior Finance Faculty Conference, the FIRS 2017 Meeting, Barcelona GSE Summer Forum, and the Western Finance Association 2017 Meeting for helpful suggestions.

1 Introduction

An investor's incentive to acquire private information changes over time and with economic conditions. For instance, rising oil prices may trigger traders to research whether airlines are hedged against fuel price fluctuations. A falling real estate market may lead investors to acquire loan-level data on their mortgage-backed securities in order to revalue their positions. A consolidation wave in a particular industry may lead market participants to investigate remaining firms as potential targets. Following Grossman and Stiglitz (1980), a large literature has studied how investors choose to acquire information, and what their decisions imply for financial markets. However, despite the inherently dynamic nature of the information acquisition decision, the existing literature treats it as a static problem by requiring that investors make their information choices before the start of trading.

We study the dynamic information acquisition decision of a strategic trader. In contrast to prior work, we allow her to choose the timing of information acquisition in response to the evolution of a public signal. We find that whether or not the market maker observes the trader's acquisition decision plays a crucial role.[1] When the trader's acquisition decision is not observable, we show that there cannot exist equilibria in which information acquisition follows a pure strategy. Moreover, we rule out existence of mixed strategy equilibria under a standard set of regularity conditions. In sharp contrast, when the trader's acquisition decision is observable, there is an equilibrium in which optimal acquisition follows a pure strategy. We show that the optimal decision exhibits delay beyond what would be predicted by a naive "NPV" rule. Furthermore, implications for the likelihood of information acquisition, price dynamics, and announcement effects are qualitatively different from those in settings in which the strategic trader is either endowed with private information or can acquire it only before the financial market opens. Taken together, our analysis suggests that the standard equilibrium in strategic trading models does not arise naturally in a setting with dynamic, costly information acquisition.

We begin with a continuous-time Kyle (1985) framework that builds on Back and Baruch (2004) and Caldentey and Stacchetti (2010). There is a single risky asset, traded by a risk-neutral, strategic trader and a mass of noise traders. We introduce a set of publicly observable signals, which affect the expected value of the asset and evolve stochastically over time. A risk-neutral market maker competitively sets the asset's price, conditional on the public signals and aggregate order flow. The asset payoff is publicly revealed at a random time.[2] The trader and the market maker share a common prior about the random variables

[1] Note that observability refers to whether the market maker can observe the fact that the trader has acquired information, *not* whether the market maker can observe the trader's information itself.

[2] The assumption of a random horizon is largely for tractability and is not qualitatively important for

in the economy. In contrast to much of the literature, we do not assume the strategic trader is endowed with private information or is required to commit to her information acquisition decision before trading begins. Instead, we allow her to pay a cost at any point in time to determine the asset value, and her optimal decision reflects the time-varying value of this information.

For concreteness, consider the following example. Suppose an airline's fuel hedging positions are not publicly known, but will be revealed at some random future time. The strategic trader can pay a cost to investigate whether the airline is exposed to oil price fluctuations. Importantly, the value of this information depends on the evolution of oil prices (a public signal). When oil prices are relatively stable, whether or not the airline is hedged does not have a large effect on the value of the firm. However, when oil prices have changed dramatically, the impact on firm value is much larger. Moreover, acquiring information immediately need not be optimal; instead, the strategic trader might prefer to wait until uncertainty is sufficiently high.

Our analysis reveals that whether or not the market maker observes the trader's information acquisition decision plays a key role in determining the nature, and indeed the very existence, of equilibrium. Section 3 considers the case where the strategic trader's decision to acquire information is not observable by the market maker. First, we explore whether there exist equilibria in which information acquisition follows a pure strategy. One can immediately rule out any equilibrium in which the strategic trader delays acquisition (or does not acquire at all). In any such equilibrium, the market maker should rationally set the price impact to zero before the anticipated acquisition time. But since acquisition is not observable, the uninformed strategic trader can deviate by acquiring information, trading against the unresponsive market maker, and making unbounded profits.

Next, we rule out pure-strategy equilibria in which the strategic trader acquires information immediately. In this case, the trader has another profitable deviation: instead of acquiring information immediately, the strategic trader can wait for a short interval of time, during which she does not trade, and then acquire information. Since future periods are discounted (due to the random horizon) the trader benefits by delaying the cost of acquisition; however, she forgoes trading gains over the interval. We establish that while the cost benefits are of order equal to the delay interval, the trading losses are of a smaller order of magnitude, so the deviation is strictly profitable for the trader.[3]

our primary results. What is key is that a random horizon induces the trader to discount future profits. We expect our results to carry over to settings with fixed horizon that feature discounting for other reasons (e.g., if the trader has a subjective discount factor or the risk-free rate is nonzero).

[3]As we discuss further in Section 3, the deviation is not a pathological feature of continuous time. We show in Appendix B that a similar deviation is profitable in an analogous discrete time setting when the

We show that analogous arguments rule out a natural class of mixed strategy equilibria. These include equilibria that involve "discrete" mixing in which the trader acquires at a countable collection of times and "continuous" mixing in which the trader acquires information with a given intensity over some interval of time. Intuitively, without observability, the strategic trader cannot commit to the equilibrium strategy because she always finds it profitable to deviate (by delaying acquisition or pre-empting it). While our results do not necessarily rule out all conceivable equilibria, our analysis implies that common strategic trading equilibria do not naturally arise as a consequence of dynamic, endogenous information acquisition, when the acquisition decision is unobservable and future payoffs are discounted.[4]

Note that unobservability of acquisition is key for the above arguments: the proposed deviations are no longer profitable when the market maker observes the acquisition decision, and can respond accordingly. Motivated by this, we consider a tractable specification of our general framework in Section 4 in which the market maker can observe the acquisition decision.[5] In this case, we show there exists an equilibrium where information acquisition is uniquely pinned down and follows a pure strategy. Appealing to standard results on optimal stopping, we characterize the trader's optimal strategy and show that it follows a cutoff rule: she chooses to acquire information only when the public uncertainty reaches a threshold. Intuitively, the ability to decide *when* to acquire information endows the trader with a call option on the expected profits from being privately informed, and she chooses to exercise the option only when the uncertainty about the asset payoff is sufficiently high (and therefore the expected profits from being informed are sufficiently large). Moreover, we show that optimal information acquisition exhibits delay — the strategic trader chooses to wait beyond the threshold that would be prescribed by an "NPV" rule. As such, the standard assumption that the trader can only choose to acquire information at the initial date is restrictive if she can condition her acquisition decision on the evolution of public news.

Consistent with the intuition from real option decisions, we show that the benefit from waiting to acquire information increases in the cost of information and the volatility of the public signal, but decreases in the prior uncertainty about the payoff. We show that the

time between trading rounds is sufficiently small.

[4]In particular, our results do not necessarily rule out equilibria in which the trading strategy is non-smooth. Back (1992) establishes that non-smooth trading strategies are not optimal when the strategic trader is exogenously informed. As such, if they were to exist in our setting, such equilibria would be qualitatively different from those considered in existing work.

[5]This is consistent with the standard framework, where the market maker knows with certainty whether there is an informed trader in the market. Entry into new markets or asset classes, addition of star analysts or traders, and regulatory position and capital reporting requirements, can serve as public, albeit noisy, signals of whether or not an institutional investor has acquired information.

likelihood of information acquisition need not always increase with volatility of the public signal. While higher signal volatility increases the likelihood that the option to acquire information ends up "in the money," it also increases the value of waiting. In fact, we show that when acquisition costs are sufficiently low, the likelihood of acquisition decreases with news volatility.

We also find that the likelihood of information acquisition need not be higher when the trading horizon is longer. When the payoff is expected to be revealed quickly, the value from being informed is very low since there is little time over which to profit at the expense of noise traders, and so the acquisition threshold is high. However, as the expected trading horizon increases, there are two offsetting effects. On the one hand, the value from being informed increases with the horizon since the trader expects her information advantage to last longer. On the other hand, the cost of waiting decreases with the horizon, since the likelihood that the payoff is revealed before acquisition is low. We find that initially the first effect dominates, while eventually the second one does. As a result, the trader is less likely to acquire information when the trading horizon is very long or very short.

The dynamic nature of the trader's information acquisition decision leads to novel implications on price dynamics and announcement returns. For instance, we show that information acquisition triggers a jump in instantaneous volatility and price impact, and following acquisition, both evolve stochastically. Notably, these jumps are not driven by jumps in fundamentals (or noise trading), but arise endogenously due to the trader's acquisition decision and the market maker's learning problem.[6] Next, we show that the announcement effect, defined as the average absolute price change at the time the asset payoff is publicly announced, need not be smaller when the strategic trader is informed. One might find this surprising, since the price is more informative about the asset payoff when the strategic trader is informed and intuition would suggest it should therefore be closer to the asset value.[7] However, when information acquisition is endogenous, there is an offsetting effect: the strategic trader only chooses to acquire information when uncertainty is sufficiently high. As a result, when the cost of information acquisition is high, the public signal volatility is high, or the expected trading horizon is extreme (i.e., sufficiently short or sufficiently long),

[6] Although not the focus of their analysis, the model in Back and Baruch (2004) also features stochastic volatility and price impact, but not jumps in volatility. However, our results are distinct from Collin-Dufresne and Fos (2016), where stochastic volatility and price impact are driven by stochastic volatility in noise trading.

[7] For instance, as Back (1992) establishes, the corresponding announcement effect must be zero in the analogous, finite horizon model where the announcement is perfectly anticipated and the trader is endowed with information. When the announcement is stochastic and the trader is endowed with information, as in Back and Baruch (2004), the announcement effect is smaller when the strategic trader is present than when she is not.

we show that the expected announcement effect is larger with information acquisition than without.

Our paper relates to the large literature on asymmetric information models with endogenous information acquisition that was initiated by Grossman and Stiglitz (1980). While a number of papers extend this basic setting to allow for dynamic trading (e.g., Mendelson and Tunca (2004), Avdis (2016)), to allow traders to condition their information acquisition decision on a public signal (e.g., Foster and Viswanathan (1993)), to allow traders to pre-commit to receiving signals at particular dates (e.g., Back and Pedersen (1998), Holden and Subrahmanyam (2002)), to incorporate a time-cost (as opposed to a monetary cost) of information (e.g., Kendall (2017)), or to incorporate a sequence of one-period information acquisition decisions (e.g., Veldkamp (2006)), the information acquisition decision remains essentially static — investors make their information acquisition decision before the start of trade.[8] The unobservable acquisition case of our model is related to a recent literature that studies markets in which some participants face uncertainty about the existence or informedness of others (e.g., Li (2013), Banerjee and Green (2015), Back, Crotty, and Li (2016), Wang and Yang (2016)). To the best of our knowledge, however, our model is the first to allow for dynamic information acquisition in that the strategic trader can choose to become privately informed at any point in time. Our analysis implies that allowing for dynamic information acquisition has economically important consequences and highlights the fact that observability of the acquisition decision plays a critical role.

2 Model

Our framework is based on the continuous-time Kyle (1985) model with random horizon in Back and Baruch (2004) and Caldentey and Stacchetti (2010). Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which is defined an $(n+1)$-dimensional standard Brownian motion $\bar{W} = (W_1, \ldots, W_n, W_Z)$ with filtration $\mathcal{F}_t^W$, independent random variables ξ and T and independent m-dimensional random vector N_0. Let $\mathcal{F}_t$ denote the augmentation of the filtration $\sigma(N_0, \{\bar{W}_s\}_{\{0 \le s \le t\}})$. Suppose that the random variable T is exponentially distributed with rate r, and that ξ and N_0 have finite second moments. Finally, let $W = (W_1, \ldots, W_n)$ denote the first n elements of $\bar{W}$.

There is an $m \geq 0$ dimensional vector of publicly-observable Markovian signals $N_t =$

[8] While Kendall (2017) studies whether or not investors wait for better quality information when there is no explicit monetary cost, the information acquisition decision is implicitly (i) publicly observable, and (ii) made prior to trading.

$(N_{1t}, \ldots, N_{mt})$ with initial value N_0 and which follows

$$dN = \mu(t, N)\, dt + \Sigma(t, N)\, dW_t,$$

where $\mu(t, N) = (\mu_1(t, N), \ldots, \mu_m(t, N))$ and $\Sigma(t, N) = (\Sigma_1(t, N)', \ldots, \Sigma_m(t, N)')$ denote the vector of drifts and matrix of diffusion coefficients. Suppose that $\mu(\cdot)$ and $\Sigma(\cdot)$ are such that there exists a unique strong solution to this set of stochastic differential equations (SDEs).[9]

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 at random time T. We assume that, given knowledge of ξ and the history of N_t, the conditional expected value v_t of this payoff as of time t is

$$v_t = f(t, \xi, N_t)$$

for some function f. There is a single, risk-neutral strategic trader who can pay a fixed cost c at any time τ to observe ξ.[10]

Let X_t denote the cumulative holdings of the trader, and suppose the initial position $X_0 = 0$. Further, suppose X_t is absolutely continuous and let $\theta(\cdot)$ be the trading rate (so $dX_t = \theta(\cdot)dt$).[11] There are noise traders who hold Z_t shares of the asset at time t, where

$$dZ_t = \sigma_Z dW_{Zt}, \tag{1}$$

with $\sigma_Z > 0$ a constant.

Finally, there is a competitive, risk neutral market maker who 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 we describe formally below for the observable and unobservable cases. In either case, the public information set always includes at least the aggregate order flow process $Y_t = X_t + Z_t$ and the public news N_t. The price at time $t < T$ is thus given by

$$P_t = \mathbb{E}\left[v_t \middle| \mathcal{F}_t^P\right]. \tag{2}$$

Let $\mathcal{T}$ denote the set of $\mathcal{F}_t^P$ stopping times. We require that the trader's information acquisition time $\tau \in \mathcal{T}$. That is, we require acquisition to depend only on public information

[9]See for instance Theorem 5.2.9 in Karatzas and Shreve (1998) who present Lipschitz and growth conditions on the coefficients that are sufficient to deliver this result.

[10]Because all market participants are risk-neutral, it is without loss of generality, economically, that v_t represents the conditional expected value of the asset rather than the value itself.

[11]Back (1992) shows that it is optimal for the trader to follow strategies of this form in a model in which she is exogenously informed.

up to that point. Let $\mathcal{F}_t^I$ denote the augmentation of the filtration $\sigma(\mathcal{F}_t^P \cup \sigma(\xi))$. Thus, $\mathcal{F}_t^I$ represents the trader's information set, post-information acquisition. We require the trader's pre-acquisition trading strategy to be adapted to $\mathcal{F}_t^P$ and her post-acquisition strategy to be adapted to $\mathcal{F}_t^I$.

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

Definition 1. *An equilibrium with pure strategy information acquisition is an acquisition time $\tau \in \mathcal{T}$ and admissible trading strategy θ for the trader and a price process P_t such that, given the trader's strategy the price process satisfies* (2) *and, given the price process, the trading strategy and acquisition time maximize the expected profit*

$$\mathbb{E}\left[\int_0^T \theta(f(u,\xi,N_u) - P_u)\,du\right].$$

For now, we focus on pure acquisition strategies. We go into further detail on mixed strategy acquisition when we consider unobservable acquisition below. Also, implicit in the definition of equilibrium is that the trader's expected profit is well-defined. In general, given a particular payoff structure, if an equilibrium exists one must introduce admissibility conditions on trading strategies to rule out strategies that first incur infinite losses by driving the price away from v_t and then reap infinite profits, leaving the expected profit undefined. We return to this point when we construct equilibria with observable information acquisition in Section 4.

Finally, note that our model nests a number of existing models of the literature. For example, Back and Baruch (2004) consider the case in which $\xi \in \{0,1\}$ has a binomial distribution and there are no publicly observable signals, so $v_t = f(t,\xi,N_t) = \xi$. Similarly, the special case of Caldentey and Stacchetti (2010) in which time is continuous and there is no ongoing flow of private information, $\xi \sim N(0,\Sigma_0)$ and $v_t = f(t,\xi,N_t) = \xi$.[12] In contrast to these earlier models, in which the insider is endowed with private information about the asset value, our focus is on allowing her to acquire information at a time of her choosing.

[12] Note that because the strategic trader receives only a lump of private information, our model is not subject to the Caldentey and Stacchetti (2010) critique that the continuous-time equilibrium is not the limit of the corresponding discrete time equilibria. However, all of the results below easily extend to time-varying ξ_t when "acquiring information" entails paying c to perfectly observe ξ_t from time τ forward.

3 Unobservable information acquisition

To begin, we study the case in which the strategic trader's decision to acquire information is not observable by the market maker. First, we show in a very general sense that there cannot exist pure strategy equilibria in which acquisition occurs after the beginning of the game. Intuitively, given a proposed equilibrium acquisition date, the strategic trader finds it optimal to deviate by pre-emptively acquiring information earlier and trading against an unresponsive pricing rule. Second, under typical regularity conditions on the trader's value function and market maker's beliefs, we also rule out pure strategy equilibrium in which acquisition occurs at the beginning of the game. We show that instead of acquiring information at the start, the strategic trader can profitably deviate by waiting — while her trading gains are unaffected, she benefits by delaying the cost of acquisition. Given the non-existence of pure strategy equilibria, we then entertain the possibility of equilibrium in which acquisition follows a mixed strategy. However, we also demonstrate that an economically important class of mixed strategy equilibria cannot exist when information acquisition is unobservable. The intuition for this result is essentially the same as that which rules out pure strategy acquisition at the beginning of the game.

Before presenting the results, we briefly describe the public information environment. When information acquisition is unobservable, the public information filtration $\mathcal{F}_t^P$ is the augmentation of the filtration $\sigma(\{N_t, Y_t\})$. That is, the mere fact that the trader has (or has not) acquired information is not directly observable; rather, as part of updating her beliefs about the asset value, the market maker must also use the public signals and the order flow to update about whether the trader has, in fact, acquired information.

3.1 Pure strategy equilibria

We begin by considering equilibria in which information acquisition follows a pure strategy. Our first observation is immediate: never acquiring information cannot be an equilibrium.

Lemma 1. *There does not exist an equilibrium in which the trader follows a pure acquisition strategy that, with probability 1, never acquires information That is, in which* $\mathbb{P}(\tau = \infty) = 1$.

Proof. In an equilibrium in which the strategic trader never acquires information, the price is insensitive to order flow at all times. But this allows the trader to deviate from the conjectured equilibrium strategy by acquiring information, trading at an arbitrarily large rate with zero price impact, and generating unboundedly large profits. Since acquisition is unobservable by the market maker, she cannot respond to the deviation by adjusting the price impact. □

A similar argument rules out pure strategy equilibria in which information is acquired with some delay.

Lemma 2. *There does not exist an equilibrium in which the trader follows a pure acquisition strategy that acquires information after time* $t = 0$ *with positive probability. That is, in which* $\tau \geq 0$ *with* $\mathbb{P}(\tau > 0) > 0$.

Proof. Again, suppose that there does exist such an equilibrium. Then in equilibrium, the order flow is completely uninformative about ξ prior to time τ and therefore the pricing rule is insensitive to order flow before τ. But in the event $\{\tau > 0\}$ (which occurs with strictly positive probability), the strategic trader can once again profitably deviate by unobservably acquiring information prior to τ and trading at an arbitrarily large rate with zero price impact, thereby generating unbounded profits. □

Intuitively, the previous two results follow from the fact that when acquisition cannot be detected, the strategic trader cannot commit to acquiring information at a future date: she always finds it profitable to deviate by pre-empting herself and acquiring information earlier. The inability to commit to the equilibrium strategy is reminiscent of the durable-good monopolist's inability to commit to high prices in the future (see Coase (1972) and Gul, Sonnenschein, and Wilson (1986)). However, in our setting, the lack of commitment leads to non-existence of pure-strategy equilibria with delay in information acquisition. Moreover, this incentive to pre-emptively acquire information is likely to apply more generally, e.g., in settings with multiple investors, in discrete time, and in settings with a fixed terminal date.

The above results imply that with unobservable acquisition, the only remaining candidate for a pure strategy equilibrium is one in which the trader acquires information immediately. To analyze this scenario, we need to introduce some additional structure and notation. Denote the types of informed strategic trader by $i \in \text{Support}(\xi)$, corresponding to a trader informed of $\xi = i$. Following the literature, we will restrict attention to equilibria in which the asset price is a function of the exogenous public signals N_t, as well as a finite number ℓ of endogenous state variables p_t that follow a Markovian diffusion and which keep track of the market maker's beliefs about ξ. We formalize this in the following.

Assumption 1. *The asset price takes the form* $P_t = g(t, N_t, p_t)$ *for some function* g *that is continuously differentiable in* t *and twice continuously differentiable in* (N, p)*. There are* $\ell > 0$ *endogenous state variables* p_t *with dynamics*

$$dp = \mu_p(t, N, p)\, dt + \Lambda_p(t, N, p) dY, \tag{3}$$

where μ_p *and* Λ_p *are* ℓ*-dimensional vector functions such that there exists a strong solution to this SDE when* $dY = \theta\, dt + \sigma_Z dW_{Zt}$ *and* θ *takes its equilibrium value.*

As suggested by the notation Λ_p is a vector of price-impact coefficients.[13] We emphasize that the function g and the coefficients μ_p and Λ_p are equilibrium objects that, given an equilibrium trading strategy, are pinned down by the rationality of the pricing rule.

We require conditions on the informed trader types' value functions.

Assumption 2. *For $i \in support(\xi)$, the value function $J^i(t,N,p)$ is continuously differentiable in t, twice continuously differentiable in (N,p), and satisfies the HJB equation:*

$$0 = \sup_{\theta} \left\{ \begin{array}{c} -rJ^i + J^i_t + J^i_N \cdot \mu + J^i_p \cdot (\mu_p + \Lambda_p \theta) + \frac{1}{2} tr\left(\Sigma\Sigma' J^i_{NN}\right) + \frac{1}{2} tr\left(\Lambda_p \Lambda_p' J^i_{pp}\right) \\ +\theta\left(f\left(t,i,N_t\right) - P_t\right) \end{array} \right\}. \quad (4)$$

Note that Assumptions 1 and 2 apply to the pricing rule and value function in existing models in the literature (e.g., Back and Baruch (2004), the continuous-time case of Caldentey and Stacchetti (2010)). Under these assumptions, we show that there cannot exist an equilibrium with immediate acquisition.

Lemma 3. *Suppose that Assumptions 1 and 2 hold. There does not exist an equilibrium in which the trader follows a smooth trading strategy and a pure acquisition strategy where she acquires immediately, $\mathbb{P}(\tau = 0) = 1$.*

The argument relies on the following deviation: instead of acquiring information immediately, the strategic trader can *wait* for a short amount of time (Δ), during which she does not trade, and then acquire information. Given that future periods are discounted (due to the stochastic horizon T), she benefits from delaying the cost of acquisition, but forgoes the trading gains over Δ. As we show, while the discounted trading costs are of order Δ, the loss in trading gains are smaller, and so the deviation leaves the trader strictly better off.

Proof. Suppose that there exists an equilibrium in which the trader acquires with probability one at $t = 0$. Let θ^i denote the equilibrium trading rate of a trader who has observed $\xi = i$, and let

$$J^i\left(t, N_t, p_t\right) = \mathbb{E}\left[\int_t^\infty e^{-r(u-t)} \theta^i(\cdot)\left(f(u,\xi,N_u) - P_u\right) du | \mathcal{F}^I_t\right]$$

denote the expected profit from time t onward for a type i trader. By assumption, the J^i are all solutions of the HJB equation (4). Because this equation is linear in θ, and θ is unconstrained, it follows that the sum of the coefficients on θ must be identically zero and

[13]Note that there is no loss in excluding a dW term in (3). Under independence of ξ and N_t, only the order flow is informative about ξ, not N_t directly. However, as the trader's strategy can in general depend on N we cannot similarly suppress N from the drift and price impact coefficients.

therefore the sum of the remaining terms must also equal zero i.e.,

$$-rJ^i + J_t^i + J_N^i \cdot \mu + J_p^i \cdot \mu_p + \frac{1}{2}\mathrm{tr}(\Sigma\Sigma' J_{NN}^i) + \frac{1}{2}\mathrm{tr}(\Lambda_p \Lambda_p' J_{pp}^i) = 0. \tag{5}$$

As first noted by Back (1992), the above reveals a key feature of continuous time Kyle models: over any finite interval of time, the trader is indifferent between playing her equilibrium trading strategy or refraining from trade over that interval and then trading optimally from that time forward. Economically, this result arises because an equilibrium pricing rule must be such that the trader does not perceive any predictability in the price level or price impact if she refrains from trading. Otherwise, she would have a profitable deviation from her conjectured equilibrium *trading* strategy.

Let $\bar{J}(t, N_t, p_t)$ denote the gross expected profit from acquiring information as of time t given that one has not acquired information previously

$$\bar{J}(t, N_t, p_t) = \mathbb{E}\left[J^{\xi}(t, N_t, p_t)\,|\mathcal{F}_t^P\right] = \mathbb{E}\left[\int_t^{\infty} e^{-r(u-t)}\theta^{\xi}(\cdot)(f(u,\xi,N_u) - P_u)\,du|\mathcal{F}_t^P\right],$$

Consider the following deviation by the trader: do not acquire information at $t = 0$, do not trade or acquire over the next small interval $[0, \Delta)$ with $\Delta > 0$, and then acquire at $t = \Delta$ and follow the conjectured equilibrium trading strategy from that point forward. The expected profit from this deviation is

$$\begin{aligned} \bar{\Pi}_{d0} &\equiv e^{-r\Delta}\mathbb{E}[\bar{J}(\Delta, N_\Delta, p_\Delta) - c] - (\bar{J}(0, N_0, p_0) - c) \\ &= (1 - e^{-r\Delta})c + \mathbb{E}[e^{-r\Delta}\bar{J}(\Delta, N_\Delta, p_\Delta) - \bar{J}(0, N_0, p_0)]. \end{aligned} \tag{6}$$

Hence, by deviating the trader benefits by paying the cost later, which is valuable due to discounting, but she forgoes trading profits in the interim.

Applying Ito's Lemma to $e^{-rt}\bar{J}$ under the assumption that $\theta = 0$ on $[0, \Delta)$, integrating, and taking the expectation implies

$$\begin{aligned} &\mathbb{E}\left[e^{-\Delta r}\bar{J}(\Delta, N_\Delta, p_\Delta) - \bar{J}(0, N_0, p_0)\right] \\ &\quad = \mathbb{E}\left[\int_0^{\Delta}\left(-r\bar{J} + \bar{J}_t + \bar{J}_N \cdot \mu + \bar{J}_p \cdot \mu_p + \frac{1}{2}\mathrm{tr}(\Sigma\Sigma'\bar{J}_{NN}) + \frac{1}{2}\mathrm{tr}(\Lambda_p\Lambda_p'\bar{J}_{pp})\right)du\right]. \end{aligned} \tag{7}$$

Suppose that we can interchange the order of integration and differentiation when calculating derivatives of $\bar{J}$ with respect to its arguments.[14] Then, eq. (5) also holds for $\bar{J}$ since $\bar{J}$ is

[14] A simple sufficient condition on primitives would be that ξ is bounded. Technically, this would prevent consideration of the case of normally distributed ξ and no news processes as in Caldentey and Stacchetti (2010); however, one can use the explicit solutions for the value functions from their paper to show that this

simply a linear combination of the J^i. Hence, (7) reduces to zero, and substituting back into eq. (6) yields

$$\bar{\Pi}_{d0} = (1 - e^{-r\Delta})c > 0,$$

which shows that the deviation is strictly profitable. $\square$

While the above argument is particularly transparent in the present setting, it is important to note that the key ingredients for the deviation argument are (i) unobservable information acquisition, (ii) discounting of future profits, and (iii) the fact that forgone trading profits are of order strictly smaller than Δ. Importantly, the result is not a pathology of the continuous-time setting, and the deviation does not necessarily rely on the assumption of smooth trading. For instance, we show in Appendix B that a similar deviation rules out pure-strategy information acquisition in the discrete-time model of Caldentey and Stacchetti (2010) when the time between trading rounds is sufficiently short. In discrete time, the expected loss from not trading over Δ is not zero, but it is of order smaller than Δ. Since the benefit from delaying the cost of acquisition is of order Δ, deviating over a sufficiently small Δ is strictly profitable.

3.2 Mixed strategy equilibria

Next, we entertain the possibility that information acquisition follows a mixed strategy. A mixed information acquisition strategy is a probability distribution over stopping times in $\mathcal{T}$. That is, at the beginning of the game, the trader randomly chooses a (pure) stopping time according to some probability distribution, and follows the realized strategy for the duration of the game.[15] Importantly, note that such a strategy can involve both "continuous" mixing in which the trader acquires information with a given intensity over an interval of times, as well as "discrete" mixing in which the trader acquires at a countable collection of times. Because the trader must be indifferent between any stopping time τ over which she mixes, each such τ must also achieve the maximum in her optimization problem

$$\max_{\{\theta^i(\cdot)\}_{i\in\{U\}\cup\text{Support}(\xi)},\tau\in\mathcal{T}} \mathbb{E}\left[\int_0^\tau \theta^U(f(s,\xi,N_s)-P_s)\,ds + \int_\tau^T \theta^\xi(f(s,\xi,N_s)-P_s)\,ds\right],$$

claim still holds in that case.

[15] There are multiple, equivalent ways of defining randomization over stopping times. Aumann (1964) introduced the notion of randomizing by choosing a stopping time according to some probability distribution at the start of the game. Touzi and Vieille (2002) treat randomization by identifying the stopping strategy with an adapted, non-decreasing, right-continuous processes on $[0, 1]$ that represents the cdf of the time that stopping occurs. Shmaya and Solan (2014) show, under weak conditions, that these definitions are equivalent.

where U indexes a strategic trader who has not yet acquired information (i.e., "uninformed"). Moreover, in any non-degenerate mixed strategy equilibrium there must be information acquisition at date zero with positive probability. If not, the price sensitivity to order flow at $t = 0$ is zero, but this implies the uninformed strategic trader can deviate by acquiring information preemptively, as argued above in Lemma 2.

The following Lemma establishes that, under the same regularity conditions as Lemma 3, there does not exist an equilibrium in which the trader follows a mixed strategy.

Lemma 4. *Suppose that Assumptions 1 and 2 hold. There does not exist an equilibrium in which the trader follows a smooth trading strategy and a mixed acquisition strategy.*

Proof. The intuition and proof for the result is essentially identical to that behind Lemma 3. Instead of acquiring with positive probability at $t = 0$, the trader benefits by instead waiting for a short amount of time Δ, not trading, and then acquiring information.

The fact that any nondegenerate mixed strategy equilibrium must involve a strictly positive probability of acquisition at $t = 0$ implies

$$J^U(0, N_0, p_0) = \bar{J}(0, N_0, p_0) - c, \tag{8}$$

since the trader only mixes when she is indifferent between remaining uninformed and acquiring information immediately.

Suppose that instead of acquiring with positive probability at $t = 0$, the trader deviates by not acquiring information over a short interval Δ and then acquires at that time. The expected profit from this deviation is

$$\begin{aligned}
\bar{\Pi}_{d0} &= e^{-r\Delta}\mathbb{E}[\bar{J}(\Delta, N_\Delta, p_\Delta) - c] - J^U(0, N_0, p_0) \\
&= e^{-r\Delta}\mathbb{E}[\bar{J}(\Delta, N_\Delta, p_\Delta) - c] - (\bar{J}(0, N_0, p_0) - c) \\
&= (1 - e^{-r\Delta})c > 0,
\end{aligned}$$

where the second line follows from eq. (8) and the final line follows from the assumption that J^i and hence $\bar{J}$ satisfies the HJB equation (4). □

The following result summarizes the previous Lemmas.

Proposition 1.
1. *There does not exist an equilibrium in which the trader follows a pure information acquisition strategy in which information is acquired after $t = 0$ with positive probability.*

2. *Suppose that Assumptions 1 and 2 hold. Then,*

(a) there does not exist an equilibrium in which the trader follows a smooth trading strategy and a pure information acquisition strategy in which information is acquired at $t = 0$ with probability 1,

(b) there does not exist an equilibrium in which the trader follows a smooth trading strategy and a mixed acquisition strategy

The results in this section rule out the existence of equilibrium in the case of unobservable information acquisition, under standard regularity conditions. The deviation arguments apply generally to a large class of models that feature discounting (e.g., Back and Baruch (2004), Chau and Vayanos (2008), Caldentey and Stacchetti (2010)), which implies that the trading equilibria in these models do not naturally arise as a consequence of costly dynamic information acquisition. Moreover, as we discuss in Appendix C, a number of the non-existence results carry over to the standard continuous-time Kyle (1985) model with a fixed horizon and no discounting. The arguments behind Lemmas 1 and 2 imply that pure-strategy information acquisition cannot exhibit delay in such a setting. Further, we show that when the cost of information is sufficiently high there are also no pure strategy equilibria with acquisition at time zero, which rules out the existence of any pure strategy equilibria in that case.

While standard strategic trading models provide useful intuition for how exogenous (or costless), private information gets incorporated into prices, our analysis recommends caution when considering settings with endogenous information acquisition. In the next section we consider the case of observable information acquisition, in which case such deviations are no longer profitable, and show that there emerge equilibria with nontrivial dynamic, endogenous information acquisition.

4 Observable information acquisition

The analysis in the previous section highlights the key role that observability plays when the strategic trader can dynamically trade and choose the time at which she acquires information. Sustaining equilibria with costly information acquisition is difficult when information acquisition is unobservable because the strategic trader has incentives to deviate by pre-empting or delaying acquisition. However, such deviations are not feasible when the information acquisition is observable, since the market maker could immediately respond to such decisions by appropriately adjusting the price impact of order flow.

Based on this observation, we consider a tractable specification of the general model from Section 2 in which the strategic trader's decision to acquire information is observable by the market maker. This is a natural benchmark. Since the market maker in standard Kyle (1985) models knows with certainty that there is an informed trader in the market, these models effectively feature observable acquisition. Moreover, in practice, decisions to acquire expertise or information are often publicly detectable. For instance, entry into new markets and asset classes by large institutional investors is usually scrutinized by other market participants. The addition of star traders, portfolio managers, and executives garners significant media attention. Finally, many institutional investors are subject to regulatory reporting requirements, and disclosures about trading positions and capital adequacy can provide noisy information about an investor's trading strategies and private information.

With observability, the information acquisition decision by the strategic trader resembles the exercise of a real option. The standard assumption that the trader makes a one-shot information acquisition decision when the financial market opens is restrictive: we show that, generally, optimal acquisition can exhibit delay. Moreover, we find that allowing for dynamic, endogenous information acquisition has qualitatively novel implications for the likelihood of information acquisition and price dynamics.

4.1 A tractable specification

Suppose the random variable $\xi \in \{0, 1\}$ is binomial with probability $\alpha = \Pr(\xi = 1)$. The risky asset pays off v at time T, where

$$v = \xi N_T. \tag{9}$$

The public news process N_t is a geometric Brownian motion

$$\frac{dN_t}{N_t} = \sigma_N \, dW_{Nt}, \tag{10}$$

where $\sigma_N > 0$ and the initial value $N_0 > 0$ is constant. Hence, the time-t conditional expected value for an informed trader is $v_t = \xi N_t$.

The assumptions that the public signal is perfectly informative about N_t and that N_t has zero drift are without loss of generality. More generally, one could replace N_t with $\hat{N}_t = \mathbb{E}[N_T | \mathcal{F}_t^P]$ in the pricing rule and trading strategy without qualitatively affecting the rest of the analysis. It is also straightforward to generalize to a general continuous, positive martingale for the news process, but at the expense of closed-form solutions to the optimal acquisition problem in most cases.

We interpret ξ as the payoff-relevance of the news process. In particular, the news process is only informative about the payoff of the risky asset if $\xi = 1$. To fix ideas, consider the example from the introduction. Suppose the market faces uncertainty about whether an airline is hedged against fuel price increases. The strategic trader must pay a cost (i.e., c) to investigate whether the airline is exposed (i.e., $\xi \in \{0, 1\}$), and can optimally choose when to do so. When the price of oil is stable, the incremental impact of hedging on firm value (i.e., v) is low. In contrast, when the price of oil moves dramatically, the airline's hedging decision has a larger effect on firm value. As such, one expects the value of learning about exposure varies over time with the publicly observable news about fuel prices (i.e., changes in N_t).[16]

A relabeling of variables offers another natural interpretation of the model. Suppose $v \in \{H_T, L_T\}$ where H_t and L_t are publicly observable processes, independent of ξ, and where $\alpha = \Pr(v = H_T)$. Then, one can express $v = L_T + \xi N_T$, where $N_T \equiv H_T - L_T$. Under this interpretation, acquiring information about ξ reveals whether the asset value is H_T or L_T, and this information is more valuable when the difference between the two possible values is larger.

More generally, the specification of the public news process allows us to introduce stochastic volatility in a parsimonious and tractable manner. Without variation in public news ($N_t \equiv 1$), the above setting reduces to the one analyzed by Back and Baruch (2004) but with endogenous information acquisition. In this case, however, the trader's acquisition decision is effectively static since the value of information is constant over time. With a stochastic news process, the value of information evolves over time, which introduces dynamic considerations to the acquisition decision. We expect alternative specifications that generate time-variation in uncertainty about fundamentals would generate similar predictions, although at the expense of tractability or a less natural economic interpretation.[17]

Let $I_t = \mathbf{1}_{\{\tau \le t\}}$ denote an indicator for whether the strategic trader has acquired information at time t or before. Because the market maker observes the public signal and order flow processes, and the acquisition status of the strategic trader, the public information filtration $\mathcal{F}_t^P$ is the augmentation of the filtration $\sigma(\{N_t, Y_t, I_t\})$.[18]

[16] Note, in this example the level of N_t is interpreted at the deviation of the price of oil from a baseline level (normalized to zero), not the price level itself.

[17] A perhaps more standard specification of the model would be one in which the value v is normally distributed with stochastic volatility (e.g., variance Σ_t). In order for this volatility to impact the acquisition decision, it must be publicly observable. However, this poses a difficulty: how does one interpret a setting in which the value of an asset is unobservable, but exhibits observable stochastic volatility? An alternative specification, in which there is a public signal with an error that exhibits stochastic volatility (e.g., $N_t = v + \epsilon_t$, where ϵ_t exhibits stochastic volatility σ_t), necessitates the introduction of two state variables (i.e., the signal N_t and the conditional variance of v under the public information set, $\Sigma_{P,t}$), which limits tractability.

[18] To reduce clutter, we abuse notation somewhat by using $\mathcal{F}_t^P$ to denote both the market maker's information set, which includes the acquisition indicator I_t in this case, as well as the trader's pre-acquisition (public) information set, which includes only the news process and order flow variables, and defines the admissible class of stopping times for acquisition.

Finally, let p_t denote the market maker's conditional probability that $\xi = 1$. Note that zero and one are absorbing states for p_t. As such, following Back and Baruch (2004), we must rule out trading strategies that first drive the stock price to zero or N_t, incurring infinite losses, and then yield infinite profits by trading against a pricing rule that is unresponsive to order flows. To do so, we add to the existing smoothness and measurability restrictions on trading strategies an additional admissibility condition which requires that the trading strategy for a trader informed of $\xi = 1$ satisfies

$$\mathbb{E}\left[\int_\tau^T N_u(1-p_u)\theta_u^- du\right] < \infty,$$

and analogously for a trader informed of $\xi = 0$,

$$\mathbb{E}\left[\int_\tau^T N_u p_u \theta_u^+ du\right] < \infty.$$

4.2 Financial market equilibrium

We can construct equilibrium by working backwards. We begin by characterizing the financial market equilibrium, conditional on an arbitrary acquisition time, and then find the optimal acquisition time given the financial market equilibrium.

Proposition 2. *Fix an information acquisition time $\tau \in \mathcal{T}$. There exists an equilibrium in the trading game in which the price of the risky asset is given by $P_t = N_t p_t$, where*

$$p_t \equiv \mathbb{E}\left[\xi \mid \mathcal{F}_t^P\right] = \begin{cases} \alpha & 0 \le t < \tau \\ \Phi\left(\Phi^{-1}(\alpha)\, e^{r(t-\tau)} + \sqrt{\frac{2r}{\sigma_Z^2}} \int_\tau^t e^{r(t-s)} dY_s\right) & \tau \le t < T \end{cases}. \tag{11}$$

Prior to information acquisition, the trader does not trade (i.e., $\theta^U \equiv 0$), and conditional on information acquisition, her strategy depends only on p and is given by

$$\theta^1(p) = \frac{\sigma_Z^2 \lambda(p)}{p}, \text{ and } \quad \theta^0(p) = -\frac{\sigma_Z^2 \lambda(p)}{1-p},$$

where θ^i, $i \in \{U, 1, 0\}$, denotes the trading strategy corresponding to the uninformed, informed of $\xi = 1$, and informed of $\xi = 0$ types. In this equilibrium, conditional on becoming informed, the trader's value function is given by

$$J^1(p_t, N_t) = N_t \int_{p_t}^1 \frac{1-a}{\lambda(a)} da, \text{ and } \quad J^0(p_t, N_t) = N_t \int_0^{p_t} \frac{a}{\lambda(a)} da, \tag{12}$$

where $\lambda(p) = \sqrt{\frac{2r}{\sigma_Z^2}}\phi\left(\Phi^{-1}(1-p)\right)$.

Our equilibrium characterization naturally extends the equilibrium in Back and Baruch (2004) to (i) accommodate the news process N_t and (ii) account for the possibility that the strategic trader is uninformed before τ. Before information acquisition, the strategic trader does not trade,[19] and consequently, the order-flow is uninformative and the market-maker does not update his beliefs about ξ. As a result, before τ the price $P_t = \alpha N_t$ evolves linearly with N_t. Conditional on information acquisition, the strategic trader optimally trades according to θ^ξ characterized in the proposition. Since $\theta^1 \neq \theta^0$, the order flow provides a noisy signal about ξ to the market maker. The market maker's conditional expectation about ξ, given by p_t, depends on the cumulative (weighted) order-flow since the acquisition date (i.e., $\int_\tau^t e^{r(t-s)}dY_s$), and consequently, so does the price P_t.

4.3 Optimal information acquisition

Given the value function in Proposition 2, we can characterize the optimal information acquisition decision.

Proposition 3. *Given the financial market equilibrium in Proposition 2, there is a unique optimal acquisition strategy: the strategic trader optimally acquires information the first time N_t hits the optimal acquisition boundary $N^* = \frac{\beta}{\beta-1}\frac{c}{K}$ from below, where*

$$K = \sqrt{\frac{\sigma_Z^2}{2r}}\phi\left(\Phi^{-1}(1-\alpha)\right), \text{ and } \beta = \frac{1+\sqrt{1+8r/\sigma_N^2}}{2}. \tag{13}$$

Moreover, the optimal acquisition boundary N^* increases in c and σ_N, decreases in σ_Z, is U-shaped in α (minimized at $\alpha = 0.5$), and is U-shaped in r.

As we show in the proof of the above, the expected profit immediately prior to acquiring information at any date t (i.e., the value function the instant before ξ is observed) is given by

$$\bar{J}(N_t) \equiv \mathbb{E}_t\left[\alpha J^1(\alpha, N_t) + (1-\alpha) J^0(\alpha, N_t)\right] = KN_t. \tag{14}$$

Note that the value function given information acquisition at date t is higher when there is more noise in the order flow (i.e., higher σ_Z), when there is more prior uncertainty about

[19] Under the posited price function, the pre-acquisition trading strategy is indeterminate. Any strategy that uses only public information earns zero expected profit under the public information set. Given such a trading strategy, it also remains optimal for the market maker to set $P_t = N_t\alpha$. Without loss of generality, we focus on the case in which the trader does not trade before time τ. In the presence of transaction costs, this would be the optimal strategy.

whether N_t is informative (i.e., when α is closer to 0.5), and when the information advantage is expected to be longer lived (i.e., when r is smaller).

The standard approach in the literature restricts the strategic trader to make her information choices before trading begins. In this case, she follows a naive "NPV" rule — she only acquires information if the value from becoming informed is higher than the cost i.e., $\bar{J}(N_0) \geq c$. As the following corollary highlights, the resulting information acquisition decision is effectively a static one.

Corollary 1. *If the strategic trader is restricted to acquiring information at $t = 0$, she optimally acquires information if and only if $N_0 \geq N_0^*$, where $N_0^* = \frac{c}{K}$. Moreover, the optimal acquisition boundary N_0^* increases in c, decreases in σ_Z, is U-shaped in α (minimized at $\alpha = 0.5$), and increases in r.*

With dynamic information acquisition, the optimal time to acquire information is characterized by the following problem:

$$J^U(n) \equiv \sup_{\tau \in \mathcal{T}} \mathbb{E}\left[1_{\{\tau < T\}}(\bar{J}(N_\tau) - c)\middle| N_t = n\right] = \sup_{\tau \in \mathcal{T}} \mathbb{E}\left[e^{-r\tau}(KN_\tau - c)^+ | N_t = n\right]. \tag{15}$$

This problem is analogous to characterizing the optimal exercise time for a perpetual American call option.[20] Notably, the optimal information acquisition decision exhibits delay: information is not acquired when $KN_t = c$, as would be implied by the static NPV rule. The intuition for this effect is analogous to that for investment delay in a real options problem. At any point in time, the trader can exercise her "option" to acquire information and use that information to profit at the expense of the noise traders. However, by waiting and observing the news process she learns additional information about the asset payoff (and therefore her ultimate profits) on which she can condition her decision. Since acquiring information irreversibly sacrifices the ability to wait, it is optimal to acquire only when doing so is sufficiently profitable to overcome this opportunity cost. Moreover, the option to wait is more valuable (and hence N^* is higher) when the volatility of the news process (i.e., σ_N) is higher.

A key difference between the static acquisition boundary of Corollary 1 and the dynamic acquisition boundary of Proposition 3 is how they respond to the expected trading horizon. In the static case, the exercise boundary is increasing in r. Recall that increasing r increases the likelihood that the payoff is revealed sooner i.e., it decreases the expected trading horizon. This naturally decreases the value from acquiring information, since the trader has a shorter window over which to exploit her informational advantage.

[20] Hence, appealing to standard results, we establish that the optimal stopping time is a first hitting time for the N_t process and show that the given N^* is a solution to this problem.

Figure 1: Exercise Boundary N^* versus trading horizon
Unless otherwise specified, parameters are set to $\sigma_Z = \sigma_N = 1$, $c = 0.25$ and $\alpha = 0.5$.

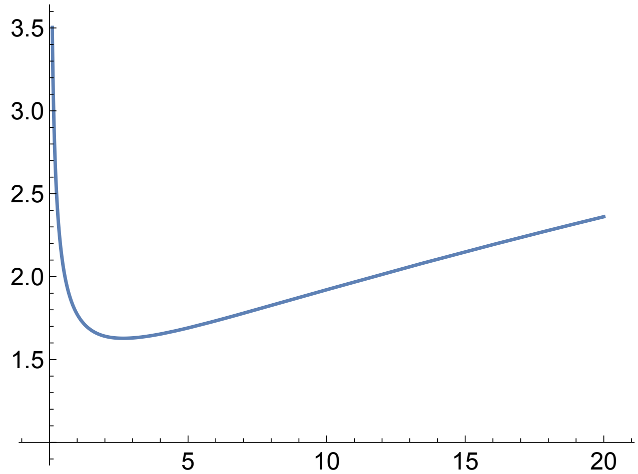

N^* versus trading horizon (i.e., $1/r$)

With dynamic information acquisition, the trader also accounts for the cost of waiting to acquire information. Specifically, as the trading horizon increases (i.e., r decreases), the expected value from acquiring information at any date (i.e., $J^U(N_t)$) increases. However, she is also willing to wait longer to acquire this information, since the cost of waiting (the probability the value will be revealed before she acquires information) also decreases. Initially, the first effect dominates, which leads the exercise boundary to decrease as the trading horizon increases. Eventually, however, the second effect dominates, and the exercise boundary increases with the horizon. As Figure 1 illustrates, this implies that the exercise boundary is non-monotonic in the trading horizon ($1/r$): the trader is less likely to acquire information when the asset payoff is expected to be revealed too quickly or too slowly.

4.4 Likelihood of information acquisition

The likelihood of information acquisition depends on two forces. First, the cost of information may be too high relative to the value of acquiring it: given c, the trader might never find it optimal to acquire the information. Second, even if the (relative) cost of acquisition is not too high, the asset payoff may be revealed before the strategic trader chooses to acquire information. The following results characterize how these effects interact to determine the likelihood of information acquisition.

In what follows, it is useful to define T_N as the first time $N_t \geq N^*$. Then, the time at which information is acquired can be expressed as

$$\tau = T_N 1_{\{T_N \leq T\}} + \infty \times 1_{\{T_N > T\}}, \tag{16}$$

where, as before, $\tau = \infty$ corresponds to no information acquisition. To avoid the trivial case, assume $N_0 < N^*$. We begin with the following observation.

Lemma 5. *Suppose $N_0 < N^*$. For $0 \leq t < \infty$, the probability that $T_N \in [t, t+dt]$ is given by*

$$\Pr\left(T_N \in [t, t+dt]\right) = \frac{\left(\log\left(\frac{N^*}{N_0}\right)\right)}{\sigma_N \sqrt{2\pi t^3}} \exp\left\{-\frac{\left(\frac{1}{\sigma_N}\log\left(\frac{N^*}{N_0}\right) + \frac{1}{2}\sigma_N^2 t\right)^2}{2t}\right\} dt. \tag{17}$$

The probability that T_N is not finite is given by $\Pr\left(T_N = \infty\right) = 1 - \frac{N_0}{N^*}$.

The result follows from applying results on the first hitting time of a Brownian motion with drift. Note that there is a positive probability that the boundary is never hit, even if $T \equiv \infty$. Also, the probability that $T_N = \infty$ is increasing in the N^*. As a result, the probability that the news process reaches N^* decreases in the cost c and volatility σ_N , increases in volatility of noise trading σ_Z and uncertainty about ξ (i.e., is hump-shaped in α), and is hump-shaped in the trading horizon (i.e., $1/r$).

The next result accounts for the possibility that the payoff is revealed before the trader acquires information (i.e, $T_N > T$).

Proposition 4. *Suppose $N_0 < N^*$. The probability that information is acquired is* $\Pr\left(\tau < \infty\right) = \left(\frac{N_0}{N^*}\right)^\beta$. *The probability is decreasing in c, increasing in N_0 and σ_Z, hump-shaped in α (around $\frac{1}{2}$), and hump-shaped in r. When $c \leq N_0 K$, the probability is decreasing in σ_N; when $c > N_0 K$, it is hump-shaped in σ_N.*

Not surprisingly, accounting for the possibility that the payoff is revealed before N_t hits N^* reduces the likelihood of information acquisition (i.e., $\Pr\left(\tau < \infty\right) < \Pr\left(T_N < \infty\right)$, since $N_0 < N^*$ and $\beta > 1$). More interestingly, it reveals novel comparative statics relative to those suggested by an NPV rule.

First, the effect of changes in expected trading horizon (changes in $1/r$) is inherited from the effect of such changes on the acquisition boundary N^*. When the expected trading horizon is short the probability of acquisition is low because, conditional on acquiring, the trader has little time to profit from her informational advantage. On the other hand, when the expected trading horizon is long, the probability of acquisition is also low because in this case the cost of waiting is sufficiently low to offset the longer trading horizon conditional on acquiring.

Second, incorporating the possibility that the payoff is revealed before the trader acquires information also changes the effect of the volatility σ_N on the likelihood of acquisition. Increasing σ_N has two effects on the probability of acquisition: (i) it increases the acquisition boundary (i.e., N^* increases in σ_N), which tends to reduce the probability of acquisition, and (ii) fixing the boundary, it increases the likelihood that N_t will hit the boundary by any given time (i.e., N_t is more volatile), which tends to increase the probability of acquisition. The overall effect of σ_N therefore depends on the relative strength of these two forces.

To gain some intuition for the dependence of the σ_N comparative statics on the initial value N_0, note that the asset value is either $v = 0$ or

$$v = N_T = N_0 e^{-\frac{1}{2}\sigma_N^2 T + \sigma_N W_T}.$$

As a result, the uncertainty about v, and consequently, the benefit of acquiring information depends on both N_0 and σ_N.[21] When N_0 is sufficiently high, uncertainty about v is already high and information acquisition is valuable. Appealing to the analogy with an American call option, the option to acquire information starts in the money. In this case, initial uncertainty is already sufficiently high that the effect of increasing σ_N on N^* dominates, and the probability of acquisition decreases in σ_N. On the other hand, when N_0 is low, uncertainty about v is low and information acquisition is not valuable, i.e., the option starts out of the money. In this case, when σ_N is low, the effect of an increase in σ_N on the volatility of N_t initially dominates, and increases the probability of acquisition. However, as σ_N continues to increase the effect of increasing the boundary N^* begins to dominate, which reduces the acquisition probability.

Figure 2 presents an example of this non-monotonic effect of σ_N on the probability of information acquisition. In panel (a), N_0 is sufficiently high so that $N_0 K \geq c$, and so the probability of information acquisition is decreasing in σ_N. In panel (b), N_0 is low enough so that the probability of information acquisition initially increases and then decreases in σ_N.

Finally, using the distribution of τ derived in the proof of Proposition 4 it is straightforward to characterize the expected time of information acquisition.

Corollary 2. *Suppose $N_0 < N^*$. The expected time of acquisition, conditional on acquisition occurring, is*

$$\mathbb{E}[\tau | \tau < \infty] = \frac{2}{\sigma_N^2} \log(N^*/N_0) \left(1 + \frac{8r}{\sigma_N^2}\right)^{-1/2}.$$

Moreover, all unconditional moments of τ are infinite.

[21] We thank Kerry Back for this intuition.

Figure 2: Probability that information is acquired $\Pr(\tau < \infty)$ versus σ_N.
Unless otherwise specified, parameters are set to $\sigma_Z = 1$, $c = 0.25$ $r = 1.5$, $\alpha = 0.5$.

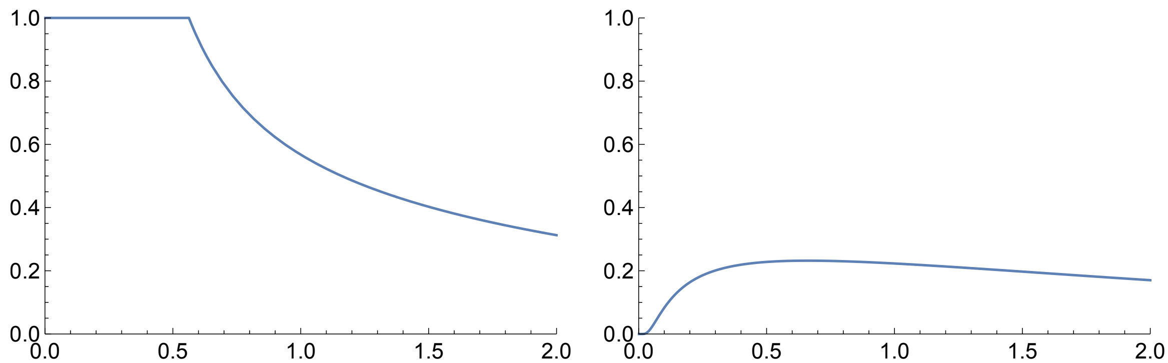

(a) $N_0 K > c$ ($N_0 = 1.5$) (b) $N_0 K < c$ ($N_0 = 1$)

4.5 Implications: Price dynamics and Announcement Effects

In this subsection, we explore some additional properties of the equilibrium with observability. First, the dynamic nature of the trader's information acquisition decision leads to novel price dynamics: information acquisition triggers a jump in instantaneous volatility and price impact, and following acquisition, both evolve stochastically. Notably, these results are not driven by stochastic volatility of fundamentals or noise trading, but arise endogenously due to the trader's acquisition decision and the market maker's learning problem.[22]

Second, we characterize the average absolute price change at the time the asset payoff is publicly announced. Intuitively, one might expect that this announcement effect is smaller when the strategic trader is informed, since the order flow is more informative about the asset payoff.[23] We show that this need not be the case when the timing of information acquisition is endogenous, because the strategic trader only chooses to acquire information when uncertainty is sufficiently high. In fact, when the cost of information acquisition is sufficiently high, the public signal volatility is sufficiently high, or the expected trading horizon is sufficiently extreme (i.e., sufficiently short or sufficiently long), the expected announcement effect is larger when there is information acquisition.

[22] Although not the focus of their analysis, a similar result on stochastic volatility and price impact arises in Back and Baruch (2004). However, our result differs from Collin-Dufresne and Fos (2016), where stochastic volatility and price impact are driven by stochastic volatility in noise trading. Furthermore, neither of these models generate jumps in volatility or price impact.

[23] For instance, as Back (1992) establishes, conditional on the strategic trader being informed the announcement effect must be zero in the analogous, finite horizon model where the announcement is perfectly anticipated. When the announcement date is stochastic, but the strategic trader is exogenously endowed with information, as in Back and Baruch (2004), the announcement effect is smaller on average when the strategic trader is informed.

4.5.1 Price dynamics

The expression for the price in Proposition 2 immediately implies that price impact of order flow before information acquisition is zero, but jumps to $\lambda(p_\tau)$ when information is acquired. Moreover, price impact evolves stochastically post-acquisition, since it is driven by the evolution of the market maker's beliefs p_t.

The following result characterizes return volatility in our model.

Proposition 5. *The instantaneous variance of returns is*

$$\nu_t \equiv \begin{cases} \sigma_N^2 & 0 \leq t < \tau \\ \sigma_N^2 + \left(\frac{\lambda^*(p_t)}{p_t}\right)^2 \sigma_Z^2 & \tau \leq t < T \end{cases}$$

Conditional on information acquisition, volatility is stochastic and exhibits the "leverage" effect i.e., the instantaneous covariance between returns and variance of returns is negative ($cov\left(\nu_t, \frac{dp_t}{p_t}\right) \leq 0$).

The above result highlights that return volatility is higher conditional on information acquisition. Conditional on no acquisition, price changes are driven purely by changes in the news process. However, conditional on the strategic trader being informed, the market maker also conditions on order flow to update his beliefs about the asset payoff, and as a result, return volatility is driven by two sources of variation.

In contrast to the standard Kyle (1985) model, our model generates stochastic return volatility and price impact, even though fundamentals (i.e., N_t) and noise trading (i.e., Z_t) are homoskedastic. This is a consequence of the non-linearity in the filtering problem of the market maker, and is in contrast to models where the (conditionally linear) filtering problem amplifies stochastic volatility in an underlying process (e.g., in Collin-Dufresne and Fos (2016), return volatility amplifies stochastic volatility in noise trading).[24] Moreover, conditional on information acquisition, return volatility also exhibits the "leverage effect" (see Black (1976) and the subsequent literature) — the instantaneous variance increases when returns are negative, and vice versa — even though there is no leverage (debt) in the underlying risky asset.

Despite the large empirical literature documenting the importance of stochastic volatility and jumps in volatility, there are relatively few theoretical explanations for how these patterns arise. Our model provides an explanation for both, but it does not rely on jumps or stochastic volatility in fundamentals. Instead, volatility jumps (and becomes stochastic)

[24] Similar results obtain in the continuous-time models of Back and Baruch (2004), Li (2013), Back et al. (2016), and the discrete-time model of Banerjee and Green (2015).

when the public news process triggers private information acquisition by the strategic trader. Our analysis suggests that further understanding the interaction between public news and private information can provide new insights into what drives empirically observed patterns in volatility.

4.5.2 Announcement effects

Next, we turn to the absolute price change at the time the payoff of the risky asset is announced. In finite horizon models where the announcement is perfectly anticipated (e.g., Back (1992)), the informed trader's optimal strategy ensures that the price change at announcement is zero. While this is no longer the case with a stochastic announcement date, the intuition from these models would suggest that the announcement effect is smaller on average if information is acquired than if it is not. However, as the next result highlights, this is not always the case.

Proposition 6. *The expected absolute price jump on announcement, conditional on information acquisition is*

$$\mathbb{E}\left[\left|\xi N_T - P_{T-}\right| \middle| \tau < \infty\right] = 2N^* h\left(\alpha\right), \tag{18}$$

where $h(\alpha)$ is characterized in the Appendix, and fully illustrated by the plot in Figure 3. The expected absolute price jump on announcement, conditional on no information acquisition is

$$\mathbb{E}\left[\left|\xi N_T - P_{T-}\right| \middle| \tau = \infty\right] = 2\alpha\left(1-\alpha\right) N^* \frac{\frac{N_0}{N^*} - \left(\frac{N_0}{N^*}\right)^{\beta}}{1 - \left(\frac{N_0}{N^*}\right)^{\beta}}.$$

Fixing $\alpha \in (0,1)$ and the other parameters, the announcement effect is larger with information acquisition when: N_0 is sufficiently small, c is sufficiently high, σ_N^2 is sufficiently high, σ_Z^2 is sufficiently low, or r is sufficiently extreme (i.e., sufficiently low, or sufficiently high).

The proposition characterizes conditions under which a potentially surprising result holds: the announcement effect is larger with information acquisition than without. In a setting where the strategic trader is exogenously endowed with information, the standard intuition holds — the announcement effect conditional on an informed trading is smaller than the announcement effect conditional on no informed trading. To see why, note that in this case, the announcement effect can be expressed as

$$\mathbb{E}\left[\left|\xi N_T - P_{T-}\right|\right] = N_0 \mathbb{E}\left[\left|\xi - \pi_T\right|\right] = 2N_0 \mathbb{E}\left[\pi_T\left(1-\pi_T\right)\right], \tag{19}$$

where $\pi_t = \mathbb{E}[\xi | \mathcal{F}_t^P]$. When the strategic trader is not informed, $\pi_T = \alpha$. When the strategic trader is informed, $\pi_T = p_T$, and so Jensen's inequality implies that $\mathbb{E}\left[p_T\left(1-p_T\right)\right] \leq$

Figure 3: $h(\alpha)$ and $\alpha(1-\alpha)$
The figure plots $h(\alpha)$ (solid) and $\alpha(1-\alpha)$ (dashed) as a function of α.

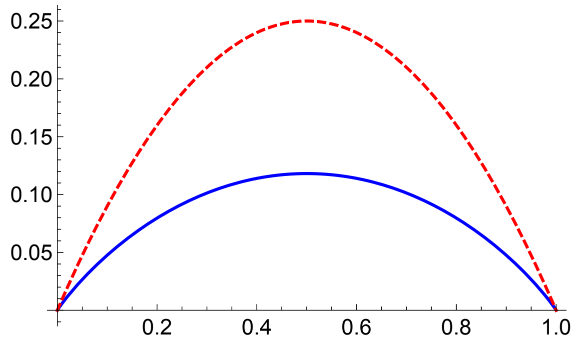

$h(\alpha)$ (solid), $\alpha(1-\alpha)$ (dashed) as a function of α

$\alpha(1-\alpha)$. Intuitively, the market-maker's posterior beliefs are more precise when the strategic trader is informed, and as a result, the price reflects the asset payoff more accurately.

When information acquisition in endogenous, however, there is an offsetting effect at work. Recall that the strategic trader only acquires information when the news process is sufficiently high ($N_t \geq N^*$). This implies that the expected level of N_T, conditional on information acquisition, is higher since $\mathbb{E}[N_T|\tau < \infty] = N^* \geq N_0$. Intuitively, the strategic trader only chooses to acquire information when the prior uncertainty about fundamentals is sufficiently high. This offsetting effect dominates when the initial news level N_0 is sufficiently small or the optimal exercise boundary N^* is sufficiently large, and as a result, the announcement effect conditional on information acquisition is higher in these cases.

5 Conclusions

The canonical Kyle-type framework, in which a market maker sets prices in response to strategic trading by an informed trader, provides an important benchmark for understanding how markets incorporate private information. A key limitation of the standard setup is that the strategic trader is endowed with private information before trading begins, instead of acquiring it endogenously at a time of her choosing. To explore the implications of endogenous information acquisition, we consider a strategic trading model in which the trader can choose *when* to acquire information about the asset payoff in response to the evolution of a public signal.

We show that the existence and nature of equilibrium depends crucially on whether the information acquisition decision is observable by the market maker. When acquisition is not observable, we find there cannot exist pure strategy equilibria. Moreover, mixed strategy equilibria satisfying standard regularity condition are also ruled out. In contrast, when acquisition is observable, we show in a tractable benchmark setting that there exists a unique equilibrium in pure (acquisition) strategies. Moreover, the equilibrium features delay in information acquisition.

Our analysis suggests that key features of the standard, strategic trading framework may be difficult to reconcile with costly dynamic information acquisition. Exploring the robustness of these results to different information acquisition technologies (e.g., costs that depend on the precision of information) and competition among traders are natural next steps. It would also be interesting to study how our analysis changes when public news is endogenous (e.g., in the form of strategic disclosure by firms or regulators).

References

Aumann, R., 1964. Mixed and behavior strategies in infinite extensive games. In: Dresher, M., Shapley, L., Tucker, A. (Eds.), Advances in Game Theory. No. 52. Princeton, pp. 627–650.

Avdis, E., 2016. Information tradeoffs in dynamic financial markets. *forthcoming in* Journal of Financial Economics.

Back, K., 1992. Insider trading in continuous time. Review of financial Studies 5 (3), 387–409.

Back, K., Baruch, S., 2004. Information in securities markets: Kyle meets glosten and milgrom. Econometrica 72 (2), 433–465.

Back, K., Crotty, K., Li, T., 2016. Identifying information asymmetry in securities markets. Available at SSRN 2565216.

Back, K., Pedersen, H., 1998. Long-lived information and intraday patterns. Journal of Financial Markets 1 (3), 385–402.

Banerjee, S., Green, B., 2015. Signal or noise? uncertainty and learning about whether other traders are informed. Journal of Financial Economics 117 (2), 398–423.

Black, F., 1976. Studies of stock price volatility changes. Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section, 177–181.

Caldentey, R., Stacchetti, E., January 2010. Insider trading with a random deadline. Econometrica 78 (1), 245–283.

Chau, M., Vayanos, D., September 2008. Strong-form efficiency with monopolistic insiders. Review of Financial Studies 21 (5), 2275–2306.

Coase, R. H., 1972. Durability and monopoly. Journal of Law and Economics 15 (1), 143–149.

Collin-Dufresne, P., Fos, V., 2016. Insider trading, stochastic liquidity and equilibrium prices. Econometrica, Forthcoming.

Foster, F. D., Viswanathan, S., 1993. The effect of public information and competition on trading volume and price volatility. Review of Financial Studies 6 (1), 23–56.

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

Gul, F., Sonnenschein, H., Wilson, R., 1986. Foundations of dynamic monopoly and the coase conjecture. Journal of Economic Theory 39 (1), 155–190.

Holden, C. W., Subrahmanyam, A., January 2002. News events, information acquisition, and serial correlation. Journal of Business 75 (1), 1–32.

Karatzas, I., Shreve, S. E., 1998. Brownian Motion and Stochastic Calculus, 2nd Edition. Vol. 113 of Graduate Texts in Mathematics. Springer, New York.

Kendall, C., 2017. The time cost of information in financial markets. working paper.

Kyle, A. S., 1985. Continuous auctions and insider trading. Econometrica, 1315–1335.

Li, T., 2013. Insider trading with uncertain informed trading. Available at SSRN 946324.

Mendelson, H., Tunca, T. I., 2004. Strategic trading, liquidity, and information acquisition. Review of Financial Studies 17 (2), 295–337.

Owen, D., 1980. A table of normal integrals: A table. Communications in Statistics-Simulation and Computation 9 (4), 389–419.

Peskir, G., Shiryaev, A., 2006. Optimal Stopping and Free-Boundary Problems. Lectures in Mathematics, ETH Zürich. Birkhäuser, Boston.

Protter, P. E., 2003. Stochastic Integration and Differential Equations, 2nd Edition. Vol. 21 of Applications of Mathematics. Springer, New York.

Shmaya, E., Solan, E., 2014. Equivalence between random stopping times in continuous time, working paper.

Touzi, N., Vieille, N., 2002. Continuous-time dynkin games with mixed strategies. SIAM Journal on Control and Optimization 41 (4), 1073–1088.

Veldkamp, L. L., June 2006. Media frenzies in markets for financial information. American Economic Review 96 (3), 577–601.

Wang, Y., Yang, M., 2016. Insider trading when there may not be an insider. Working paper (https://ssrn.com/abstract=2720736).

A Proofs

Proof of Proposition 2. To establish the equilibrium in the Proposition, we need to show: (i) the proposed price function is rational, and (ii) the informed trader's strategy is optimal. Fix any $\tau \in \mathcal{T}$.

Rationality of pricing function

Consider the set $\{t : t < \tau\}$ on which the trader has not acquired information. Then, because $\{N_t\}$, $\{Z_t\}$ and ξ are independent, and under the proposed trading strategy $Y_t = Z_t$ for $t < \tau$, it is immediate that

$$\mathbb{E}[\xi N_T|\mathcal{F}_t^P] = \mathbb{E}[\xi|\mathcal{F}_t^P]\mathbb{E}[N_T|\mathcal{F}_t^P] = \alpha\mathbb{E}[N_T|\mathcal{F}_t^P].$$

Since T is almost surely finite and is independent of the process N_t we have $\mathbb{E}[N_T|\mathcal{F}_t^P] = N_t$, and so $\mathbb{E}[\xi N_T|\mathcal{F}_t^P] = \alpha N_t$.

Now, consider the set $\{t : \tau \leq t < T\}$ on which the trader has acquired information and the asset payoff has not yet occurred. Up to the addition of the news process, the problem now resembles that considered in Back and Baruch (2004), and we can adapt the proof offered there. Specifically, consider the pricing rule from Back and Baruch (2004), adapted for the fact that information is acquired at time τ,

$$dp_t = \lambda(p)dY_t, \quad p_\tau = \alpha,$$

where $\lambda(p)$ is given in the statement of the Proposition. (Later we will show that this pricing rule can be written in the explicit form in eq. (11).) Note that the proposed trading strategy depends only on ξ and p, the process p depends only on the order flow, and $\{N_t\}$ is independent of ξ and $\{Z_t\}$, so $(\xi, \{p_t\})$ is conditionally independent of $\{N_t\}$, and therefore

$$\mathbb{E}[\xi N_T|\mathcal{F}_t^P] = \mathbb{E}[\xi|\mathcal{F}_t^P]\mathbb{E}[N_T|\mathcal{F}_t^P] = \mathbb{E}[\xi|\{Y_s\}_{s\leq t}]N_t,$$

where the final equality follows since $\mathbb{E}[N_T|\mathcal{F}_t^P] = N_t$. Furthermore, since $Y_t = Z_t$ for $t < \tau$ under the proposed trading strategy and ξ is independent of $\{Z_t\}$ it follows that $\mathbb{E}[\xi|\{Y_s\}_{s\leq t}] = \mathbb{E}[\xi|\{Y_s\}_{\tau\leq s\leq t}]$.

Recall that as of time τ, the informed trader begins trading according to the strategy $\theta^\xi(p)$ and the order flow becomes informative. The market maker's conditional expectation is simply equal to her prior α since before this time only noise traders have been active. It follows that starting at time τ the market maker's filtering problem becomes identical to that of the market maker in Back and Baruch (2004). Hence, their Theorem 1 implies that for $t \geq \tau$ the pricing rule

$$dp_t = \lambda(p)dY_t, \qquad p_\tau = \alpha,$$

satisfies $p_t = \mathbb{E}[\xi|\{Y_s\}_{s\geq\tau}]$.

To complete the proof of the rationality of the proposed price, it suffices to show that the explicit form of $p(\cdot)$ for $\tau \leq t < T$ in eq. (11) satisfies $dp_t = \lambda(p)dY_t$. Applying Ito's Lemma to the function $f(p) = \sqrt{\frac{\sigma_Z^2}{2r}}\Phi^{-1}(p)$ to the above process for p_t gives

$$\begin{aligned} df(p_t) &= \frac{1}{2}\sigma_Z^2\lambda^2(p_t)\frac{\frac{2r}{\sigma_Z^2}f(p_t)}{\lambda^2(p_t)}\,dt + \frac{1}{\lambda(p_t)}\lambda(p_t)dY_t \\ &= rf(p_t)\,dt + dY_t. \end{aligned}$$

Now applying Ito's lemma to the function $e^{-rt}f(p_t)$ and integrating allows one to express

$$f(p_t) = f(p_\tau)e^{rt} + \int_\tau^t e^{r(t-s)}dY_s.$$

Note that $f(p_\tau) = \sqrt{\frac{\sigma_Z^2}{2r}}\Phi^{-1}(\alpha)$, so returning to the explicit form of the function $f(p)$ and inverting it follows that

$$p_t = \Phi\left(\Phi^{-1}\left(\alpha\right)e^{r(t-\tau)} + \sqrt{\frac{2r}{\sigma_Z^2}}\int_\tau^t e^{r(t-s)}dY_s\right).$$

Optimality of trading strategy

Next, we demonstrate the optimality of the proposed trading strategy, taking as given the acquisition time τ. This analysis closely follows the proof in Back and Baruch (2004). Define $V(p) \equiv \int_p^1 \frac{1-a}{\lambda(a)}da$ and consider the proposed post-acquisition value function for the case $\xi = 1$ (the case for $\xi = 0$ is analogous)

$$J^1\left(p_t, N_t\right) = N_t V(p_t).$$

We begin by showing that the given J characterizes the value function for $t \geq \tau$. Consider $\{t : \tau \leq t < T\}$ and suppose $\xi = 1$. Direct calculation on the function V yields

$$V' = \frac{p-1}{\lambda} \tag{20}$$

$$rV = \frac{1}{2}\sigma_Z^2\lambda^2 V'', \tag{21}$$

which coincides with eq. (1.15) and (1.16) in Back and Baruch (2004).

Let θ_t denote an arbitrary admissible trading strategy. Following Back and Baruch (2004), let $\hat{p}_t$ denote the process defined by $\hat{p}_s = \alpha$ for $s \leq \tau$ and $d\hat{p}_t = \lambda(\hat{p})dY_t$ for $t > \tau$ and $0 < \hat{p}_t < 1$, with Y_t generated when the trader follows the given arbitrary trading strategy. In order to condense notation, in this section, we denote $\mathbb{E}[\cdot|\mathcal{F}_t^P] = \mathbb{E}_t[\cdot]$. Since θ is admissible, we know that

$$\mathbb{E}_\tau\left[\int_\tau^T N_u(1-p_u)\theta_u^- du\right] = \mathbb{E}_\tau\left[\int_\tau^\infty e^{-r(u-\tau)}N_u(1-\hat{p}_u)\theta_u^- du\right] < \infty,$$

from which it follows that

$$\int_\tau^\infty e^{-r(u-\tau)}N_u(1-\hat{p}_u)\theta_u^- du < \infty$$

almost surely, and therefore that the integral

$$\int_\tau^\infty e^{-r(u-\tau)}N_u(1-\hat{p}_u)\theta_u du$$

is well-defined, though is possibly infinite.

Let $\hat{T} = \inf\{t \geq \tau : \hat{p} \in \{0,1\}\}$. Applying Ito's lemma to $e^{-r(t-\tau)}J$ yields

$$\begin{aligned}
&e^{-r(t\wedge\hat{T}-\tau)}J^1(\hat{p}_{t\wedge\hat{T}}, N_{t\wedge\hat{T}}) - J^1(\hat{p}_\tau, N_\tau)\\
&\quad= \int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)}N\left(-rV(\hat{p}_u) + \lambda\theta V'(\hat{p}_u) + \frac{1}{2}\sigma_Z^2\lambda^2 V''\right)du\\
&\qquad\qquad + \sigma_Z\int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)}N\lambda V'(\hat{p}_u)dW_{Zu} + \sigma_N\int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)}NV(\hat{p}_u)dW_{Nu}\\
&\quad= -\int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)}N_u\theta_u(1-\hat{p}_u)du - \sigma_Z\int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)}N_u(1-\hat{p}_u)dW_{Zu}\\
&\qquad\qquad + \sigma_N\int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)}N_uV(\hat{p}_u)dW_{Nu}
\end{aligned} \tag{22}$$

where the last equality uses eq. (20) and (21). Since $V \geq 0$, the above implies

$$\int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)}N_u\theta_u(1-\hat{p}_u)du \leq N_\tau V(\alpha) + x(t), \tag{23}$$

where we define $x(t) = \sigma_N \int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)} N_u V(\hat{p}_u) dW_{Nu} - \sigma_Z \int_\tau^{t\wedge\hat{T}} e^{-r(u-\tau)} N_u (1-\hat{p}_u) dW_{Zu}$. The integrands in the stochastic integrals are locally bounded and hence the integrals are local martingales (Thm. 29, Ch. 4, Protter (2003)). It follows that $x(t)$ is itself a local martingale (Thm. 48, Ch. 1, Protter (2003)).

Let $\hat{\tau}_n$ be a localizing sequence of stopping times for $x(t)$. That is, $\hat{\tau}_{n+1} \geq \hat{\tau}_n$, $\hat{\tau}_n \to \infty$, and $x(t \wedge \hat{\tau}_n)$ is a martingale for each n. Because $x(t)$ is a local martingale such a sequence exists (e.g., because $x(t)$ is continuous we can take $\hat{\tau}_n = \inf\{t : |x(t)| \geq n\}$). Further considering the sequence $n \wedge \hat{\tau}_n$, eq. (23) implies

$$\int_\tau^{n\wedge\hat{\tau}_n\wedge\hat{T}} e^{-r(u-\tau)} N_u \theta_u (1-\hat{p}_u) du \leq N_\tau V(\alpha) + x(n \wedge \hat{\tau}_n).$$

Applying Fatou's lemma,[25] along with this inequality, yields

$$\begin{aligned}
\mathbb{E}_\tau \left[\int_\tau^{\hat{T}} e^{-r(u-\tau)} N_u \theta_u (1-\hat{p}_u) du \right] &\leq \liminf_{n\to\infty} \mathbb{E}_\tau \left[\int_\tau^{n\wedge\hat{\tau}_n\wedge\hat{T}} e^{-r(u-\tau)} N_u \theta_u (1-\hat{p}_u) du \right] \\
&\leq N_\tau V(\alpha) + \liminf_{n\to\infty} \mathbb{E}_\tau \left[x(n \wedge \hat{\tau}_n) \right] \\
&\leq N_\tau V(\alpha).
\end{aligned}$$

Note that for $\hat{T} < \infty$ we have $\hat{p}_{\hat{T}} = 1$ since $\hat{p}_{\hat{T}} = 0$ would imply a violation of the admissibility condition. To establish this, note that eq. (22) implies

$$-\mathbb{E}_\tau \left[\int_\tau^{n\wedge\hat{\tau}_n\wedge\hat{T}} e^{-r(u-\tau)} N_u \theta_u (1-\hat{p}_u) du \right] = \mathbb{E}_\tau \left[e^{-r(t\wedge\hat{T}-\tau)} N_{t\wedge\hat{T}} V(\hat{p}_{t\wedge\hat{T}}) - N_\tau V(\alpha) \right] - J^1(\hat{p}_\tau, N_\tau),$$

and therefore

$$\begin{aligned}
& - \mathbb{E}_\tau \left[\int_\tau^{\hat{T}} e^{-r(u-\tau)} N_u \theta_u (1-\hat{p}_u) du \right] \\
& \geq \limsup_{n\to\infty} \mathbb{E}_\tau \left[- \int_\tau^{n\wedge\hat{\tau}_n\wedge\hat{T}} e^{-r(u-\tau)} N_u \theta_u (1-\hat{p}_u) du \right] \\
& = \limsup_{n\to\infty} \mathbb{E}_\tau \left[e^{-r(n\wedge\hat{\tau}_n\wedge\hat{T}-\tau)} N_{n\wedge\hat{\tau}_n\wedge\hat{T}} V(\hat{p}_{n\wedge\hat{\tau}_n\wedge\hat{T}}) - N_\tau V(\alpha) \right] - J^1(\hat{p}_\tau, N_\tau) \\
& \geq \mathbb{E}_\tau \left[e^{-r(\hat{T}-\tau)} N_{\hat{T}} V(\hat{p}_{\hat{T}}) \right] - J^1(\hat{p}_\tau, N_\tau) \\
& = \infty,
\end{aligned}$$

[25] The typical formulation of Fatou's Lemma requires that the integrands f_n be weakly positive. However, if f_n^- is bounded above by an integrable function g, considering $f_n + g$ in Fatou's lemma delivers the result. Here, due to the admissibility condition we can take $g = N_u(1-p_u)\theta_u^-$.

where the first line applies the 'reverse' Fatou's Lemma, the second line uses the equality in the previous displayed equation, the third line applies Fatou's Lemma and the final line follows because $V(0) = \infty$. Furthermore, $\hat{p}_u = \hat{p}_{\hat{T}} = 1$ for all $u \geq \hat{T}$ since 1 is an absorbing state. It follows that

$$\mathbb{E}_\tau \left[\int_\tau^\infty e^{-r(u-\tau)} N_u \theta_u (1 - \hat{p}_u) du \right] = \mathbb{E}_\tau \left[\int_\tau^{\hat{T}} e^{-r(u-\tau)} N_u \theta_u (1 - \hat{p}_u) du \right] \leq N_\tau V(\alpha). \tag{24}$$

Furthermore, this inequality is trivially true for $\hat{T} = \infty$, so it holds regardless of the behavior of $\hat{T}$. It follows that

$$N_\tau V(\alpha) \geq \mathbb{E}_\tau \left[\int_\tau^\infty e^{-r(u-\tau)} N_u \theta_u (1 - \hat{p}_u) du \right] = \mathbb{E}_\tau \left[\int_\tau^T N_u \theta_u (1 - p_u) du \right],$$

since $\hat{p} = p$ for $t \leq T$. Hence $N_\tau V(\alpha)$ is an upper bound on the post-acquisition value function.

To establish the optimality of the trader's post-acquisition strategy and the expression for the value function, it remains to show that the expected profits generated by the strategy attain the bound $N_\tau V(\alpha)$. (We show below that the trader's overall trading strategy is admissible.) Compute the trader's expected profit at time τ. We have

$$\begin{aligned}
\mathbb{E}_\tau \left[\int_\tau^T \theta^1(p_u) N_u (1 - p_u)\, du \right] &= \int_\tau^\infty \mathbb{E}_\tau \left[\mathbf{1}_{\{t \leq T\}} \theta^1(p_u) N_u (1 - p_u) \right] du \\
&= \int_\tau^\infty \mathbb{E}_\tau [N_u] \mathbb{E}_\tau \left[\mathbf{1}_{\{t \leq T\}} \theta^1(p_u)(1 - p_u) \right] du \\
&= N_\tau \int_\tau^\infty \mathbb{E}_\tau \left[\mathbf{1}_{\{t \leq T\}} \theta^1(p_u)(1 - p_u) \right] du \\
&= N_\tau \mathbb{E}_\tau \left[\int_\tau^T \theta^1(p_u)(1 - p_u)\, du \right],
\end{aligned}$$

where the first equality applies Fubini's theorem which is permissible because the integrand is positive, the second equality uses the fact that N is independent of T and $\{p_u\}$, the next-to-last equality follows because N is a martingale, and the final equality applies Fubini's theorem again. Back and Baruch (2004) establish that under the given trading strategy and pricing rule, $V(\alpha) = \mathbb{E}_\tau \left[\int_\tau^T \theta^1(p_u)(1 - p_u)\, du \right]$. Hence,

$$N_\tau V(\alpha) = \mathbb{E}_\tau \left[\int_\tau^T \theta^1(p_u) N_u (1 - p_u)\, du \right],$$

which establishes the optimality of the post-acquisition trading strategy.

Let $J^U(N)$ denote the pre-acquisition value function (i.e., the value function for an uninformed trader). Note that because $p \equiv 0$ for $t < \tau$, J^U effectively depends only on the news process in this case. We need to characterize this function and establish that the overall posited trading strategy, involving no trade prior to acquisition, is optimal. Under the given trading strategy, we have

$$\begin{aligned} J^U(N) &= \mathbb{E}\left[\mathbf{1}_{\{\tau<T\}} \int_\tau^T \theta^\xi(p_u) N_u(\xi - p_u)\, du\right] \\ &= \mathbb{E}\left[\mathbf{1}_{\{\tau<T\}} J^\xi(p_\tau, N_\tau)\right] \end{aligned}$$

Let $\check{\theta}$ be any admissible trading strategy that is adapted to $\mathcal{F}_t^P$ and $\hat{\theta}$ any admissible strategy that is adapted to $\mathcal{F}_t^I$. Then $\theta = \mathbf{1}_{\{t<\tau\}}\check{\theta} + \mathbf{1}_{\{t\geq\tau\}}\hat{\theta}$ is an arbitrary admissible strategy that obeys the restriction that the trader does not observe ξ until time τ. The expected profits from following this strategy are

$$\begin{aligned} &\mathbb{E}_0\left[\mathbf{1}_{\{\tau<T\}} \int_0^\tau \check{\theta}_u N_u(\xi - \alpha)\, du + \mathbf{1}_{\{\tau<T\}} \int_\tau^T \hat{\theta}_u N_u(\xi - p_u)\, du + \mathbf{1}_{\{\tau\geq T\}} \int_0^T \check{\theta}_u N_u(\xi - \alpha)\, du\right] \\ &= \mathbb{E}_0\left[\mathbf{1}_{\{\tau<T\}} \int_\tau^T \hat{\theta}_u N_u(\xi - p_u)\, du\right] \\ &= \mathbb{E}_0\left[\mathbf{1}_{\{\tau<T\}} \mathbb{E}\left[\int_\tau^T \hat{\theta}_u N_u(\xi - p_u)\, du | \mathcal{F}_\tau^I\right]\right] \\ &\leq \mathbb{E}_0\left[\mathbf{1}_{\{\tau<T\}} J^\xi(p_\tau, N_\tau)\right] \\ &= J^U(N), \end{aligned}$$

where the first equality takes expectations over ξ, the second equality uses the law of iterated expectations, and the inequality follows since it was shown above that as of time τ, our posited trading strategy achieves higher expected profit than any other admissible strategy. □

Proof of Proposition 3. Let $\bar{J}(N_t)$ denote the value of acquiring information when the news process is equal to N_t. Using the expression for the post-acquisition value function in Proposition 2, we have

$$\bar{J}\left(N_t\right) = N_t \left(\alpha \int_\alpha^1 \frac{1-a}{\lambda(a)}\, da + (1-\alpha) \int_0^\alpha \frac{a}{\lambda(a)}\, da\right) \equiv N_t K.$$

Make the change of variables $x = \Phi^{-1}(1-a)$ in the integrals in the expression for $J^U(N_t)$

$$K = \alpha \sqrt{\tfrac{\sigma_Z^2}{2r}} \int_\alpha^1 (1-a) \frac{1}{\phi(\Phi^{-1}(1-a))}\, da + (1-\alpha) \sqrt{\tfrac{\sigma_Z^2}{2r}} \int_0^\alpha a \frac{1}{\phi(\Phi^{-1}(1-a))}\, da$$

$$
\begin{aligned}
&= -\alpha\sqrt{\tfrac{\sigma_Z^2}{2r}}\int_{\Phi^{-1}(1-\alpha)}^{-\infty}\Phi(x)dx - (1-\alpha)\sqrt{\tfrac{\sigma_Z^2}{2r}}\int_{\infty}^{\Phi^{-1}(1-\alpha)}(1-\Phi(x))\,dx \\
&= \alpha\sqrt{\tfrac{\sigma_Z^2}{2r}}\int_{-\infty}^{\Phi^{-1}(1-\alpha)}\Phi(x)dx + (1-\alpha)\sqrt{\tfrac{\sigma_Z^2}{2r}}\int_{\Phi^{-1}(1-\alpha)}^{\infty}(1-\Phi(x))\,dx.
\end{aligned}
$$

Now integrate by parts

$$
\begin{aligned}
K &= \alpha\sqrt{\tfrac{\sigma_Z^2}{2r}}\int_{-\infty}^{\Phi^{-1}(1-\alpha)}\Phi(x)dx + (1-\alpha)\sqrt{\tfrac{\sigma_Z^2}{2r}}\int_{\Phi^{-1}(1-\alpha)}^{\infty}(1-\Phi(x))\,dx \\
&= \alpha\sqrt{\tfrac{\sigma_Z^2}{2r}}\left(-\int_{-\infty}^{\Phi^{-1}(1-\alpha)}x\phi(x)\,dx + x\Phi(x)\Big|_{-\infty}^{\Phi^{-1}(1-\alpha)}\right) \\
&\qquad + (1-\alpha)\sqrt{\tfrac{\sigma_Z^2}{2r}}\left(\int_{\Phi^{-1}(1-\alpha)}^{\infty}x\phi(x)\,dx + x(1-\Phi(x))\Big|_{\Phi^{-1}(1-\alpha)}^{\infty}\right) \\
&= \alpha\sqrt{\tfrac{\sigma_Z^2}{2r}}\left(-\int_{-\infty}^{\Phi^{-1}(1-\alpha)}x\phi(x)\,dx + (1-\alpha)\Phi^{-1}(1-\alpha)\right) \\
&\qquad + (1-\alpha)\sqrt{\tfrac{\sigma_Z^2}{2r}}\left(\int_{\Phi^{-1}(1-\alpha)}^{\infty}x\phi(x)\,dx - \alpha\Phi^{-1}(1-\alpha)\right) \\
&= \sqrt{\tfrac{\sigma_Z^2}{2r}}\int_{-\infty}^{\Phi^{-1}(1-\alpha)}-x\phi(x)\,dx = \sqrt{\tfrac{\sigma_Z^2}{2r}}\phi(\Phi^{-1}(1-\alpha)),
\end{aligned}
$$

since $\int -x\phi(x)dx = \int \phi'(x)dx = \phi(x)$.

The pre-acquisition value function under optimal stopping is

$$
J^U(n) \equiv \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[\mathbf{1}_{\{\tau<T\}}(KN_\tau - c) \mid N_t = n\right] = \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[e^{-r\tau}(KN_\tau - c)^+ \mid N_t = n\right],
$$

where the second equality follows because T is independently exponentially distributed and it suffices to consider only the positive part of $KN_\tau - c$ since the trader can always guarantee herself zero profit by not acquiring. Note that this problem is similar to pricing a perpetual American call option on an asset with price process KN_t that follows a geometric Brownian motion and with strike price c. Hence, standard results (Peskir and Shiryaev (2006), Chapter 4) imply that there is a uniquely optimal stopping time and this time is is a first hitting time of the N_t process,

$$
T_N = \inf\{t > 0 : N_t \geq N^*\},
$$

where $N^* > 0$ is a constant to be determined.

The value function and optimal N^* solve the following free boundary problem

$$
\begin{aligned}
rJ^U &= \tfrac{1}{2}\sigma_N^2 N_t^2 J_{NN}^U && \text{for } n < N^* \\
J^U(N^*) &= KN^* - c && \text{`value matching'} \\
J_N^U(N^*) &= K && \text{`smooth pasting'} \\
J^U(n) &> (n-c)^+ && \text{for } n < N^* \\
J^U(n) &= (n-c)^+ && \text{for } n > N^* \\
J^U(0) &= 0.
\end{aligned}
$$

To determine the solution in the continuation region $n < N^*$, consider a trial solution of the form $J^U(n) = An^\beta$. Substituting and matching terms in the differential equation yields

$$
r = \tfrac{1}{2}\sigma_N^2 \beta(\beta - 1), \quad \beta = \tfrac{1}{2} \pm \tfrac{1}{2}\sqrt{1 + \tfrac{8r}{\sigma_N^2}}
$$

and the boundary condition at $N = 0$ requires that one take the positive root

$$
\beta = \tfrac{1}{2} + \tfrac{1}{2}\sqrt{1 + \tfrac{8r}{\sigma_N^2}}.
$$

Applying the above conjecture to the value-matching and smooth pasting conditions implies:

$$
N^* = \frac{\beta}{\beta - 1}\frac{c}{K}, \qquad A = \frac{K}{\beta}\left(\frac{\beta}{\beta - 1}\frac{c}{K}\right)^{1-\beta} = \frac{c}{\beta - 1}\frac{1}{(N^*)^\beta},
$$

and the resulting function satisfies $J^U(n) > n - c$ in the continuation region, which establishes the result. The comparative statics with respect to c, σ_N, σ_Z, and α are immediate from the explicit expression for N^*. Moreover, since

$$
\frac{\partial}{\partial r} N^* = \frac{c}{\sigma_Z^2 \phi\left(\Phi^{-1}(1-\alpha)\right)} \frac{4\sqrt{2}\left(\sqrt{r} - 2\sqrt{\frac{r}{\frac{8r}{\sigma_N^2} + 1}}\right)}{\left(\sigma_N - \sqrt{\sigma_N^2 + 8r}\right)^2} \tag{25}
$$

we know that N^* is decreasing in r when $r < \frac{3}{8}\sigma_N^2$, but increasing otherwise. $\square$

Proof of Lemma 5. Note that

$$
\begin{aligned}
N_t \geq N^* &\iff \log(N_t) \geq \log(N^*) \\
&\iff -\frac{1}{2}\sigma_N t + W_{Nt} \geq \frac{1}{\sigma_N}(\log(N^*/N_0)),
\end{aligned}
$$

so that the first time that N_t hits N^* is the first time that a Brownian motion with drift $-\frac{1}{2}\sigma_N$ hits $\frac{1}{\sigma_N}(\log(\frac{N^*}{N_0}))$. It follows from Karatzas and Shreve (1998) (Chapter 3.5, Part C, p.196-197) that for $N_0 < N^*$ the density of T_N is

$$\Pr\left(T_N \in [t, t+dt]\right) = \frac{\left(\log\left(\frac{N^*}{N_0}\right)\right)}{\sigma_N\sqrt{2\pi t^3}} \exp\left\{-\frac{\left(\frac{1}{\sigma_N}\log\left(\frac{N^*}{N_0}\right) + \frac{1}{2}\sigma_N t\right)^2}{2t}\right\} dt.$$

Moreover, since $\frac{1}{\sigma_N}(\log(\frac{N^*}{N_0})) > 0$ but the drift of the Brownian motion is $-\frac{1}{2}\sigma_N < 0$, it follows from Karatzas and Shreve (1998) (p.197) that $\Pr(T_N = \infty) > 0$. Specifically, note that

$$\Pr\left(T_N < \infty\right) = \int_0^\infty \frac{\left(\log\left(\frac{N^*}{N_0}\right)\right)}{\sigma_N\sqrt{2\pi t^3}} \exp\left\{-\frac{\left(\frac{1}{\sigma_N}\log\left(\frac{N^*}{N_0}\right) + \frac{1}{2}\sigma_N t\right)^2}{2t}\right\} dt = \frac{N_0}{N^*}, \tag{26}$$

which implies $\Pr\left(T_N = \infty\right) = 1 - \frac{N_0}{N^*}$. $\square$

Proof of Proposition 4. Given the definition of τ, we have that for $0 \leq t < \infty$,

$$\begin{aligned}
\Pr\left(\tau \in [t, t+dt]\right) = & \begin{array}{c} \Pr\left(\tau \in [t, t+dt] \,\middle|\, T_N \leq T\right) \Pr\left(T_N \leq T\right) \\ +\Pr\left(\tau \in [t, t+dt] \,\middle|\, T_N > T\right) \Pr\left(T_N > T\right) \end{array} && (27) \\
&= \Pr\left(T_N \in [t, t+dt] \,\middle|\, T_N \leq T\right) \Pr\left(T_N \leq T\right) && (28) \\
&= \Pr\left(T_N \in [t, t+dt]\right) \Pr\left(T \geq t\right) && (29) \\
&= e^{-rt} \Pr\left(T_N \in [t, t+dt]\right). && (30)
\end{aligned}$$

Integrating gives us

$$\begin{aligned}
\Pr\left(\tau < \infty\right) &= \int_0^\infty e^{-rt} \frac{\left(\log\left(\frac{N^*}{N_0}\right)\right)}{\sigma_N\sqrt{2\pi t^3}} \exp\left\{-\frac{\left(\frac{1}{\sigma_N}\log\left(\frac{N^*}{N_0}\right) + \frac{1}{2}\sigma_N t\right)^2}{2t}\right\} dt && (31) \\
&= \frac{e^{-\frac{\log(N^*/N_0)\sqrt{8r+\sigma_N^2}}{2\sigma_N}}}{\sqrt{N^*/N_0}} = \left(\frac{N_0}{N^*}\right)^\beta && (32)
\end{aligned}$$

The comparative statics for c, N_0, σ_Z and α follow from plugging in the expressions for N^* and β. To establish the comparative statics for σ_N, first note that since $\lim_{\sigma_N \to 0} \beta = \infty$,

$\lim_{\sigma_N \to \infty} \beta = 1$, and $N^* = \frac{\beta}{\beta - 1} \frac{c}{K}$,

$$\lim_{\sigma_N \to \infty} \Pr(\tau < \infty) = 0 \tag{33}$$

$$\lim_{\sigma_N \to 0} \Pr(\tau < \infty) = \begin{cases} 0 & \text{if } c > N_0 K \\ 1 & \text{if } c \leq N_0 K \end{cases}. \tag{34}$$

Let

$$\zeta \equiv \frac{\partial}{\partial \beta} \left(\log \left(\Pr(\tau < \infty) \right) \right) = \frac{\partial}{\partial \beta} \left(\frac{N_0}{N^*} \right)^{\beta} = \log \left(\frac{N_0}{N^*} \right) + \frac{1}{\beta - 1} \tag{35}$$

which implies $\lim_{\sigma_N \to 0} \zeta = \lim_{\beta \to \infty} \zeta = \log \left(\frac{N_0 K}{c} \right)$, $\lim_{\sigma_N \to \infty} \zeta = \lim_{\beta \to 1} \zeta = \infty$, and

$$\frac{\partial}{\partial \sigma_N} \zeta = \frac{\partial \zeta}{\partial \beta} \frac{\partial \beta}{\partial \sigma_N} = -\frac{1}{\beta (1-\beta)^2} \frac{\partial \beta}{\partial \sigma_N} > 0. \tag{36}$$

Since $\frac{\partial}{\partial \sigma_N} \log \left(\Pr(\tau < \infty) \right) = \zeta \frac{\partial \beta}{\partial \sigma_N}$, we have the following results:

- When $c \leq N_0 K$, since $\zeta \geq 0$ for $\sigma_N \to 0$ and $\frac{\partial}{\partial \sigma_N} \zeta > 0$ we have $\zeta > 0$ for all σ_N, which in turn implies $\frac{\partial}{\partial \sigma_N} \log \left(\Pr(\tau < \infty) \right) < 0$ for all σ_N.
- When $c > N_0 K$, ζ crosses zero once, from below, as σ_N increases, which implies $\frac{\partial}{\partial \sigma_N} \log \left(\Pr(\tau < \infty) \right) = 0$ at exactly this one point. In this case, $\Pr(\tau < \infty)$ is hump-shaped.

Similarly, for r, $\frac{\partial}{\partial r} \log(\Pr(\tau < \infty)) = \zeta \frac{\partial}{\partial r} \beta - \frac{\beta}{2r}$. We have $\frac{\partial}{\partial r} \zeta = -\frac{1}{\beta(\beta-1)^2} \frac{\partial}{\partial r} \beta - \frac{1}{2r} < 0$. Since $\frac{\partial}{\partial r} \beta = \frac{1}{\sigma_N^2 \left(\beta - \frac{1}{2} \right)} > 0$ this implies $\frac{\partial}{\partial r} \log(\Pr(\tau < \infty))$ crosses zero as most once as r increases and from above if it does so. Consider the limit as r tends to zero,

$$\begin{aligned} \lim_{r \to 0} \frac{\partial}{\partial r} \log(\Pr(\tau < \infty)) &= \lim_{r \to 0} \left(\zeta \frac{\partial}{\partial r} \beta - \frac{\beta}{2r} \right) \\ &= \lim_{r \to 0} \frac{2r\zeta - \sigma_N^2 \beta \left(\beta - \frac{1}{2} \right)}{2\sigma_N^2 r \left(\beta - \frac{1}{2} \right)}. \end{aligned} \tag{37}$$

If it can be shown that the numerator in eq. (37) has a finite, positive limit it will follow that the overall limit is ∞. Considering the numerator, we have

$$\begin{aligned} \lim_{r \to 0} \left(2r\zeta - \sigma_N^2 \beta \left(\beta - \frac{1}{2} \right) \right) &= 2 \lim_{r \to 0} r \left(\frac{1}{\beta - 1} - \log \frac{\beta}{\beta - 1} - \log \sqrt{2r} \right) - \frac{1}{2} \sigma_N^2 \\ &= \sigma_N^2 - 2 \lim_{r \to 0} \frac{\frac{1}{\beta(\beta-1)}}{\frac{1}{r^2}} - \frac{1}{2} \sigma_N^2 \\ &= \frac{1}{2} \sigma_N^2 - 2 \lim_{r \to 0} \frac{2r}{(2\beta - 1) \frac{\partial}{\partial r} \beta} = \frac{1}{2} \sigma_N^2 \end{aligned}$$

where the second equality applies l'Hôspital's rule to the three different terms and uses the fact $\frac{\partial}{\partial r}\beta \to \frac{2}{\sigma_N^2}$ as $\beta \to 1$. The third equality rearranges the expression in the remaining limit to place r^2 in the numerator and uses l'Hôspital's rule again. Returning to eq. (37), this implies $\lim_{r\to 0} \frac{\partial}{\partial r} \log(\mathbb{P}(\tau < \infty)) = \infty$.

Now, consider $\lim_{r\to\infty} \frac{\partial}{\partial r} \log(\Pr(\tau < \infty))$. We have

$$\lim_{r\to\infty} \zeta = \lim_{r\to\infty} \left(\tfrac{1}{\beta-1} - \log\tfrac{\beta}{\beta-1} - \log\sqrt{2r}\right) = \lim_{\beta\to\infty} \left(\tfrac{1}{\beta-1} - \log\tfrac{\beta}{\beta-1}\right) - \lim_{r\to\infty} \log\sqrt{2r} = -\infty.$$

Because $\frac{\partial}{\partial r}\beta > 0$, it follows that $\lim_{r\to\infty} \frac{\partial}{\partial r} \log(\Pr(\tau < \infty)) = -\infty$, which completes the proof. $\square$

Proof of Proposition 5. Using the expression for the asset price in Proposition 2,

$$dP_t = \begin{cases} \alpha\sigma_N N_t dW_{Nt} & 0 \le t < \tau \\ \sigma_N N_t p(Y_t) dW_{Nt} + N_t \lambda^*(p_t)\sigma_Z\, dW_{Yt} & \tau \le t < T \end{cases},$$

where $W_{Yt} \equiv Y_t/\sigma_Z$ is a standard Brownian motion under the public filtration and is independent of W_{Nt}. Hence,

$$\frac{dP_t}{P_t} = \begin{cases} \sigma_N dW_{Nt} & 0 \le t < \tau \\ \sigma_N dW_{Nt} + \frac{\lambda^*(p_t)}{p_t}\sigma_Z\, dW_{Yt} & \tau \le t < T \end{cases}.$$

Letting ν_t denote the instantaneous variance of the return process gives:

$$\nu_t \equiv \begin{cases} \sigma_N^2 & 0 \le t < \tau \\ \sigma_N^2 + \left(\frac{\lambda(p_t)}{p_t}\right)^2 \sigma_Z^2 = \sigma_N^2 + 2r\left(\frac{\phi\left(\Phi^{-1}(p_t)\right)}{p_t}\right)^2 & \tau \le t \end{cases} \tag{38}$$

Let $f(p) \equiv \phi\left(\Phi^{-1}(p)\right)$, and note that $f_p = -\Phi^{-1}(p)$ and $f_{pp} = -\frac{1}{f}$. Conditional on information acquisition, note that by Ito's Lemma, we have:

$$d\nu_t = \nu_p dp_t + \tfrac{1}{2}\nu_{pp}\left(\lambda\left(p_t\right)\right)^2 \sigma_Z^2 dt = \nu_p dp_t + r f(p)^2 \nu_{pp} dt, \tag{39}$$

where $\nu_p = 4r\left(\frac{f}{p}\right)\left(\frac{f_p p - f}{p^2}\right) < 0$, and

$$\nu_{pp} = 4r\left(\tfrac{f_p p - f}{p^2}\right)^2 + 4r\left(\tfrac{f}{p}\right)\left(\tfrac{p^2(f_{pp}p + f_p - f_p) - 2p(f_p p - f)}{p^4}\right). \tag{40}$$

Since $\nu_p < 0$, the above implies that conditional on information acquisition, instantaneous return variance ν_t and returns are negatively related i.e., $\text{cov}\left(\frac{dp_t}{p_t}, d\nu_t\right) < 0$. $\square$

Proof of Proposition 6. For the no acquisition case,

$$\mathbb{E}\left[\left|\xi N_T - P_{T^-}\right| \middle| T_N > T\right] = \mathbb{E}\left[N_T \middle| \xi - \alpha\right| \left| T_N > T\right] = 2\alpha\left(1-\alpha\right)\mathbb{E}\left[N_T \middle| T_N > T\right] \tag{41}$$

Next, note that

$$\mathbb{E}\left[N_T\right] = \Pr\left(T_N < T\right)\mathbb{E}\left[N_T \middle| T_N < T\right] + \Pr\left(T_N \geq T\right)\mathbb{E}\left[N_T \middle| T_N \geq T\right] \tag{42}$$

$$\Rightarrow \mathbb{E}\left[N_T \middle| T_N > T\right] = \frac{N_0 - \Pr\left(T_N < T\right) N^*}{\Pr\left(T_N \geq T\right)} \tag{43}$$

$$= \frac{N_0 - \left(\frac{N_0}{N^*}\right)^{\beta} N^*}{1 - \left(\frac{N_0}{N^*}\right)^{\beta}} \tag{44}$$

since $\mathbb{E}\left[N_T\right] = N_0$, $\mathbb{E}\left[N_T \middle| T_N < T\right] = N^*$ and $\Pr\left(T_N < T\right) = \left(\frac{N_0}{N^*}\right)^{\beta}$. This produces the desired expression.

Conditional on information acquisition, the expected announcement effect is

$$\mathbb{E}\left[\left|\xi N_T - P_{T^-}\right| \middle| \tau < \infty\right] = \mathbb{E}\left[N_T \middle| \xi - p\left(Y_T\right)\right| \left| T_N < T\right] \tag{45}$$

$$= 2\mathbb{E}\left[N_T p\left(Y_T\right)\left(1 - p\left(Y_T\right)\right) \middle| \middle| T_N < T\right] \tag{46}$$

$$= 2\mathbb{E}\left[\mathbb{E}_{T_N}\left[N_T p\left(Y_T\right)\left(1 - p\left(Y_T\right)\right) \middle| T_N < T\right] \middle| \middle| T_N < T\right] \tag{47}$$

$$= 2\mathbb{E}\left[N_{T_N}\mathbb{E}_{T_N}\left[p\left(Y_T\right)\left(1 - p\left(Y_T\right)\right) \middle| T_N < T\right] \middle| \middle| T_N < T\right] \tag{48}$$

$$= 2N^*\mathbb{E}\left[p\left(Y_T\right)\left(1 - p\left(Y_T\right)\right) \middle| T_N < T\right] \tag{49}$$

the first and second equalities use the law of iterated expectations, the third equality uses the fact that conditional on $\sigma\left(\mathcal{F}^P_{T_N} \cup \{T_N < T\}\right)$, $N_T - N_{T_N}$ and Y_T are independent, the fourth equality uses the fact that N is a martingale, and the final equality uses $N_{T_N} = N^*$.

Suppose $\tau \in [t, t + dt]$. Given the characterization of p_t in Proposition 2, we can express p_s for $s \geq t$ as $p_s = \Phi\left(\frac{\sqrt{2r}}{\sigma_Z} z_s\right)$, where

$$z_s \middle| \{\tau \in [t, t + dt]\} \sim N\left(\Phi^{-1}\left(\alpha\right) e^{r(s-t)}, \frac{\sigma_Z^2}{2r}\left(e^{2r(s-t)} - 1\right)\right). \tag{50}$$

Next, note that for $w \sim N(0, 1)$, we have

$$\mathbb{E}\left[\Phi\left(a + bw\right)\left[1 - \Phi\left(a + bw\right)\right]\right] = \Phi\left(\frac{a}{\sqrt{1+b^2}}\right) - \left[\Phi\left(\frac{a}{\sqrt{1+b^2}}\right) - 2T^o\left(\frac{a}{\sqrt{1+b^2}}, \frac{1}{\sqrt{1+2b^2}}\right)\right] \tag{51}$$

from Owen (1980) 10,010.8 and 20,010.4, where $T^o(a,b)$ is the Owen T function. Let $\tilde{z}_s \equiv \frac{z_s - e^{r(s-t)}z_0}{\sqrt{\frac{\sigma_Z^2}{2r}\left(e^{2r(s-t)}-1\right)}} \sim N(0,1)$, and note that $p(z_s) = \Phi(a + b\tilde{z}_s)$. This implies

$$G(t,s) \equiv \mathbb{E}_t\left[p_s(1-p_s)\,|\tau \in [t,t+dt], s>t\right] = 2T^o\left(\Phi^{-1}(\alpha), \tfrac{1}{\sqrt{2e^{2r(s-t)}-1}}\right). \tag{52}$$

Since the stopping time T is exponentially distributed, we have

$$\begin{aligned}
&\mathbb{E}_t\left[p(Y_T)(1-p(Y_T))\,|T>t, \tau\in[t,t+dt]\right] \\
&= e^{-rt}\int_{s=t}^{\infty} re^{-r(s-t)}\mathbb{E}_t\left[p(Y_s)(1-p(Y_s))\,\middle|\tau\in[t,t+dt]\right]ds && (53)\\
&= \int_0^{\infty} re^{-rs}G(0,s)\,ds && (54)\\
&= 2\int_0^{\infty} re^{-rs}T^o\left(\Phi^{-1}(\alpha), \tfrac{1}{\sqrt{2e^{2rs}-1}}\right)ds && (55)\\
&= 2\int_0^{\infty} e^{-x}T^o\left(\Phi^{-1}(\alpha), \tfrac{1}{\sqrt{2e^{2x}-1}}\right)dx, \text{ where } x = rs && (56)\\
&\equiv h(\alpha) && (57)
\end{aligned}$$

This implies that

$$\begin{aligned}
&\mathbb{E}\left[p(Y_T)(1-p(Y_T))\,\middle|\tau<T\right] \\
&= \int_0^{\infty}\mathbb{E}_t\left[p(Y_T)(1-p(Y_T))\,|T>t,\tau\in[t,t+dt]\right]\Pr\left(\tau\in[t,t+dt]\,|T>\tau\right)dt && (58)\\
&= h(\alpha)\int_0^{\infty}\Pr\left(\tau\in[t,t+dt]\,|T>\tau\right)dt = h(\alpha) && (59)
\end{aligned}$$

which implies $\mathbb{E}\left[|\xi N_T - P_{T-}|\,\middle|\tau<T\right] = 2N^*h(\alpha)$.

Note that the announcement effect is bigger conditional on no acquisition if and only if:

$$\begin{aligned}
2N^*h(\alpha) < 2\alpha(1-\alpha)N^*\frac{\frac{N_0}{N^*}-\left(\frac{N_0}{N^*}\right)^{\beta}}{1-\left(\frac{N_0}{N^*}\right)^{\beta}} &\Leftrightarrow \frac{h(\alpha)}{\alpha(1-\alpha)} < \frac{\frac{N_0}{N^*}-\left(\frac{N_0}{N^*}\right)^{\beta}}{1-\left(\frac{N_0}{N^*}\right)^{\beta}} && (60)\\
&\Leftrightarrow \tfrac{h(\alpha)}{\alpha(1-\alpha)}\left(1-\left(\tfrac{N_0}{N^*}\right)^{\beta}\right) < \tfrac{N_0}{N^*}-\left(\tfrac{N_0}{N^*}\right)^{\beta} && (61)\\
&\Leftrightarrow \tfrac{h(\alpha)}{\alpha(1-\alpha)} < \tfrac{N_0}{N^*}-\left(\tfrac{N_0}{N^*}\right)^{\beta}\left(1-\tfrac{h(\alpha)}{\alpha(1-\alpha)}\right) && (62)
\end{aligned}$$

For a fixed α, since

$$\frac{N_0}{N^*} = N_0 \frac{\beta-1}{\beta c} K = \frac{N_0}{c} \frac{\frac{1}{2}\left(1+\sqrt{1+8\frac{r}{\sigma_N^2}}\right)-1}{\frac{1}{2}\left(1+\sqrt{1+8\frac{r}{\sigma_N^2}}\right)} \frac{\sigma_Z}{\sqrt{2r}} \phi\left(\Phi^{-1}(\alpha)\right), \tag{63}$$

implies that $\frac{N_0}{N^*} \to 0$ when $r \to 0$, $r \to \infty$, $\sigma_N \to \infty$,$c \to \infty$ or $\sigma_Z \to 0$. Moreover, since $\beta > 1$ and $\frac{N_0}{N^*} < 1$, we have $\left(\frac{N_0}{N^*}\right)^\beta \to 0$ when $\left(\frac{N_0}{N^*}\right) \to 0$. Now, fix α and pick a δ such that $0 < \delta < \frac{h(\alpha)}{\alpha(1-\alpha)}$. Then, the above implies that for sufficiently extreme r, sufficiently large σ_N, sufficiently large c or sufficiently small σ_Z, $\frac{N_0}{N^*} - \left(\frac{N_0}{N^*}\right)^\beta \left(1 - \frac{h(\alpha)}{\alpha(1-\alpha)}\right) < \delta$, and so the announcement effect is bigger conditional on acquisition. $\square$

B Unobservable acquisition in discrete time

In this appendix, we establish that there does not exist an equilibrium with pure-strategy information acquisition in the discrete-time model of Caldentey and Stacchetti (2010) when the time between trading rounds, Δ, is sufficiently small. Note that Lemmas 1 and 2 apply to this setting: never acquiring information, or acquiring it with a delay, cannot be an equilibrium. In either case, the strategic trader can unobservably deviate by acquiring information earlier, and trade profitably against an insensitive pricing rule. The rest of the appendix establishes that information acquisition at date zero is not an equilibrium when the length between trading rounds is sufficiently small. The argument follows that of Lemma 3: instead of acquiring immediately, the strategic trader can wait for a period and re-evaluate her decision. The expected gain from delaying acquisition is of order Δ, but the expected loss from not trading in the first period is of order smaller than Δ. As a result, when Δ is sufficiently small, the deviation is strictly profitable.

B.1 Setup

Consider the setting in Caldentey and Stacchetti (2010). Time is discrete and trade takes place at dates $t_n = n\Delta$ for $n \geq 0$ and $\Delta > 0$. There is a risky asset that pays off $V \sim N(0, \Sigma_0)$ immediately after trading round T, where T is random. Specifically, $T = \eta\Delta$, where η is geometrically distributed with failure probability $\rho = e^{-r\Delta}$.[26] There is a risk-neutral strategic trader who observes V. Let x_n denote her trade at date t_n. There are noise traders who

[26] There are at least two different distributions that are often referred to as "the geometric distribution". The one we use here is supported on the nonnegative integers $n \in \{0, 1, 2, \dots\}$ and has probability mass function $f_n = \rho^n(1-\rho)$ and cdf $F_n = 1 - \rho^{n+1}$.

submit iid trades $z_t \sim N(0, \Sigma_z)$ with $\Sigma_z = \sigma_z^2 \Delta$. Let $y_n = x_n + z_n$ denote the time n order flow. Competitive risk-neutral market makers set the price p_n in each trading round equal to the conditional expected value. Following Caldentey and Stacchetti (2010), we focus on linear, Markovian equilibria in which the time t_n price depends only on p_{n-1} and y_n. Let $\bar{V}_n$ and Σ_n denote the market maker's conditional expectation and variance, immediately *before* the time t_n trading round. So, $\bar{V}_n = p_{n-1}$. Finally, set $p_{-1} = \mathbb{E}[V] = 0$.

Caldentey and Stacchetti (2010) show that there exists an equilibrium in which the asset price and trading strategy are given by

$$\begin{aligned} p_n(\bar{V}_n, y_n) &= \bar{V}_n + \lambda_n y_n \\ x_n(V, \bar{V}_n) &= \beta_n (V - \bar{V}_n), \end{aligned}$$

the trader's expected profit is

$$\Pi_n(p_{n-1}, V) = \alpha_n (V - p_{n-1})^2 + \gamma_n,$$

and the constants are characterized by the difference equations

$$\begin{aligned} \Sigma_{n+1} &= \frac{\Sigma_n \Sigma_z}{\beta_n^2 \Sigma_n + \Sigma_z} \\ \beta_{n+1}\Sigma_{n+1} &= \rho \beta_n \Sigma_n \left(\frac{\Sigma_z^2}{\Sigma_z^2 - \beta_n^4 \Sigma_n^2} \right) \\ \lambda_n &= \frac{\beta_n \Sigma_n}{\beta_n^2 \Sigma_n + \Sigma_z} \\ \alpha_n &= \frac{1 - \lambda_n \beta_n}{2\lambda_n} \\ \rho \gamma_{n+1} &= \gamma_n - \frac{1 - 2\lambda_n \beta_n}{2\lambda_n (1 - \lambda_n \beta_n)} \lambda_n^2 \Sigma_z \\ \gamma_0 &= \sum_{k=0}^{\infty} \rho^k \left(\frac{1 - 2\lambda_k \beta_k}{2\lambda_k (1 - \lambda_k \beta_k)} \right) \lambda_k^2 \Sigma_z \end{aligned}$$

B.2 Nonexistence of equilibrium with unobservable acquisition

In this section, we show that for sufficiently small $\Delta > 0$ there does not exist an equilibrium in which information acquisition follows a pure strategy. First, note that the same preemption argument used in the tex rules out a pure strategy equilibrium in which the trader acquires information after round $n = 0$. Hence, we need only to search for equilibria in which the trader acquires with probability 1 at the beginning of round $n = 0$.

Suppose that there is such an equilibrium. The ex-ante expected profit from acquiring information immediately before the $t = 0$ trading round is

$$\begin{aligned}\bar{\Pi}_0 \equiv \mathbb{E}[\Pi_0(p_{-1}, V)] - c &= \alpha_0 \mathbb{E}[(V - p_{-1})^2] + \gamma_0 - c \\ &= \alpha_0 \Sigma_0 + \gamma_0 - c.\end{aligned}$$

We would like to compare this to the expected profit if the trader deviates by remaining uninformed for the $n = 0$ trading round and then acquiring immediately before round $n = 1$. Supposing that she does so, trades x units at time zero, and then follows the prescribed equilibrium trading strategy in the following rounds, the expected profit is

$$\begin{aligned}&\mathbb{E}[x(V - p_0(\bar{V}_0, x + z_0))] + \mathbb{E}\left[\sum_{n=1}^{\infty} \rho^n (V - p_n) x_n - \rho c\right] \\ &= \mathbb{E}[x(V - \lambda_0(x + z_0))] + \rho \mathbb{E}\left[\sum_{n=1}^{\infty} \rho^n (V - p_n) x_n - c\right] \\ &= \mathbb{E}[x(V - \lambda_0(x + z_0))] + \rho \mathbb{E}[\Pi_1(p_0, V) - c] \\ &= \mathbb{E}[x(V - \lambda_0(x + z_0))] + \rho \mathbb{E}[\alpha_1 (V - p_0)^2 + \gamma_1 - c] \\ &= \mathbb{E}[x(V - \lambda_0(x + z_0))] + \rho \mathbb{E}[\alpha_1 (V - \lambda_0(x + z_0))^2 + \gamma_1 - c] \\ &= -x^2 \lambda_0 + \rho \left(\alpha_1(\Sigma_0 + \lambda_0^2 \Sigma_z + \lambda_0^2 x^2) + \gamma_1 - c\right),\end{aligned}$$

Take $x = 0$. This yields ex-ante deviation profits

$$\bar{\Pi}_{d0} = \rho \left(\alpha_1(\Sigma_0 + \lambda_0^2 \Sigma_z) + \gamma_1 - c\right).$$

This deviation is profitable if and only if

$$\begin{aligned}&\bar{\Pi}_{d0} - \bar{\Pi}_0 > 0 \\ &\iff \rho \left(\alpha_1(\Sigma_0 + \lambda_0^2 \Sigma_z) + \gamma_1 - c\right) - (\alpha_0 \Sigma_0 + \gamma_0 - c) > 0 \\ &\iff (\rho \alpha_1 - \alpha_0)\Sigma_0 + \rho \gamma_1 - \gamma_0 + \rho \alpha_1 \lambda_0^2 \Sigma_z + (1 - \rho) c > 0.\end{aligned}$$

We have

$$\begin{aligned}\rho \alpha_1 - \alpha_0 &= \rho \frac{1 - \lambda_1 \beta_1}{2\lambda_1} - \frac{1 - \lambda_0 \beta_0}{2\lambda_0} \\ &= \rho \frac{1 - \frac{\beta_1^2 \Sigma_1}{\beta_1^2 \Sigma_1 + \Sigma_z}}{2\frac{\beta_1 \Sigma_1}{\beta_1^2 \Sigma_1 + \Sigma_z}} - \frac{1 - \frac{\beta_0^2 \Sigma_0}{\beta_0^2 \Sigma_0 + \Sigma_z}}{2\frac{\beta_0 \Sigma_0}{\beta_0^2 \Sigma_0 + \Sigma_z}} \\ &= \frac{\rho}{2} \frac{\Sigma_z}{\beta_1 \Sigma_1} - \frac{1}{2} \frac{\Sigma_z}{\beta_0 \Sigma_0}\end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{2}\Sigma_z \frac{1}{\beta_0\Sigma_0}\left(\frac{\rho\beta_0\Sigma_0}{\beta_1\Sigma_1} - 1\right) \\
&= \frac{1}{2}\Sigma_z \frac{1}{\beta_0\Sigma_0}\left(\frac{\Sigma_z^2 - \beta_0^4\Sigma_0^2}{\Sigma_z^2} - 1\right) \\
&= -\frac{1}{2}\frac{\beta_0^3\Sigma_0}{\Sigma_z},
\end{aligned}
$$

where the first equality substitutes in from the difference equation for α_n, the second equality substitutes from the equation for λ_n, the third and fourth simplify and collect terms, the fifth equality uses the difference equation for $\beta_{n+1}\Sigma_{n+1}$, and the final equality simplifies and collects terms.

Similarly,

$$
\begin{aligned}
\rho\gamma_1 - \gamma_0 &= -\frac{1 - 2\lambda_0\beta_0}{2\lambda_0(1-\lambda_0\beta_0)}\lambda_0^2\Sigma_z \\
&= -\frac{1}{2}\frac{1 - 2\lambda_0\beta_0}{1-\lambda_0\beta_0}\lambda_0\Sigma_z \\
&= -\frac{1}{2}\frac{1 - 2\frac{\beta_0^2\Sigma_0}{\beta_0^2\Sigma_0+\Sigma_z}}{1 - \frac{\beta_0^2\Sigma_0}{\beta_0^2\Sigma_0+\Sigma_z}}\lambda_0\Sigma_z \\
&= -\frac{1}{2}\frac{\Sigma_z - \beta_0^2\Sigma_0}{\Sigma_z}\lambda_0\Sigma_z \\
&= -\frac{1}{2}(\Sigma_z - \beta_0^2\Sigma_0)\lambda_0,
\end{aligned}
$$

where the first equality uses the difference equation for γ_n, the second equality cancels a λ_0, the third equality substitutes for λ_0, and the last two equalities simplify.

Furthermore, recalling from the calculations for $\rho\alpha_1 - \alpha_0$ that $\rho\alpha_1 = \frac{\rho}{2}\frac{\Sigma_z}{\beta_1\Sigma_1}$ we have

$$
\begin{aligned}
\rho\alpha_1\lambda_0^2\Sigma_z &= \frac{\rho}{2}\frac{\Sigma_z}{\beta_1\Sigma_1}\lambda_0^2\Sigma_z \\
&= \frac{1}{2}\frac{\Sigma_z^2}{\beta_0\Sigma_0\left(\frac{\Sigma_z^2}{\Sigma_z^2-\beta_0^4\Sigma_0^2}\right)}\lambda_0^2 \\
&= \frac{1}{2}\frac{1}{\beta_0\Sigma_0}(\Sigma_z^2 - \beta_0^4\Sigma_0^2)\lambda_0^2 \\
&= \frac{1}{2}\frac{1}{\beta_0\Sigma_0}(\Sigma_z^2 - \beta_0^4\Sigma_0^2)\frac{\beta_0\Sigma_0}{\beta_0^2\Sigma_0+\Sigma_z}\lambda_0 \\
&= \frac{1}{2}(\Sigma_z^2 - \beta_0^4\Sigma_0^2)\frac{1}{\beta_0^2\Sigma_0+\Sigma_z}\lambda_0,
\end{aligned}
$$

where the second equality substitutes from the difference equation for $\beta_n\Sigma_n$, the third equality simplifies, the fourth equality substitutes for λ_0, and the final equality simplifies.

Combining the most recent two displayed expressions, we have

$$\begin{aligned}
\rho\gamma_1 - \gamma_0 + \rho\alpha_1\lambda_0^2\Sigma_z &= \frac{1}{2}(\Sigma_z^2 - \beta_0^4\Sigma_0^2)\frac{1}{\beta_0^2\Sigma_0 + \Sigma_z}\lambda_0 - \frac{1}{2}(\Sigma_z - \beta_0^2\Sigma_0)\lambda_0 \\
&= \frac{1}{2}\lambda_0\frac{1}{\beta_0^2\Sigma_0 + \Sigma_z}(\Sigma_z^2 - \beta_0^4\Sigma_0^2 - (\Sigma_z - \beta_0^2\Sigma_0)(\beta_0^2\Sigma_0 + \Sigma_z)) \\
&= \frac{1}{2}\lambda_0\frac{1}{\beta_0^2\Sigma_0 + \Sigma_z}(\Sigma_z^2 - \beta_0^4\Sigma_0^2 - \beta_0^2\Sigma_z\Sigma_0 - \Sigma_z^2 + \beta_0^4\Sigma_0^2 + \beta_0^2\Sigma_z\Sigma_0) \\
&= 0.
\end{aligned}$$

It follows that

$$\begin{aligned}
\bar{\Pi}_{d0} - \bar{\Pi}_0 &= (\rho\alpha_1 - \alpha_0)\Sigma_0 + \rho\gamma_1 - \gamma_0 + \rho\alpha_1\lambda_0^2\Sigma_z + (1-\rho)c \\
&= -\frac{1}{2}\frac{\beta_0^3\Sigma_0^2}{\Sigma_z} + (1-\rho)c. \qquad (64)
\end{aligned}$$

We would like to study the behavior of the above expression as $\Delta \to 0$. To make clear the dependence of the various coefficients h_0 on Δ, we write h_0^Δ as applicable in the following.

Lemma 6. *There exists a strictly increasing function ψ such that*

$$\beta_0 = \frac{\sqrt{\Sigma_Z}}{\Sigma_0}\psi(\Sigma_0).$$

Furthermore, $\psi(0) = 0$ and we have

$$\psi(\Sigma_0) \leq [1-\rho]^{1/4}\sqrt{\Sigma_0}.$$

Proof. This proof leans heavily on the Appendix of Caldentey and Stacchetti (2010) but specialized to the case in which there is no flow of private information ($\Sigma_v \equiv 0$). As such, we point out only the essential differences in the analysis.

Define

$$A_n = \Sigma_n, \quad B_n = \frac{\beta_n\Sigma_n}{\sqrt{\Sigma_z}}$$

Then the difference equations for Σ_n and $\beta_n\Sigma_n$ imply that $(A_{n+1}, B_{n+1}) = F(A_n, B_n)$, where

$$F_A(A_n, B_n) = \frac{A_n^2}{A_n + B_n^2}, \quad F_B(A_n, B_n) = \rho\left[\frac{A_n^2 B_n}{A_n^2 - B_n^4}\right].$$

Note that these are similar to those in Caldentey and Stacchetti (2010), with the exception that there is no $+1$ term in F_A owing to the absence of a flow of private information Σ_v. Further, define

$$G_1(A) = 0, \quad G_2(A) = \sqrt{A}[1-\rho]^{1/4}, \quad G_3(A) = \sqrt{A},$$

where the function G_1 is defined so that $F_A(A, G_1(A)) = A$ and $G_2(A)$ is such that $F_B(A, G_2(A)) = G_2(A)$. (Note that there is a typo in this definition in Caldentey and Stacchetti (2010) which claims that $F_B(A, G_2(A)) = B$. However, this cannot hold in general since B is on only one side of the equation.) Finally, $G_3(A)$ is defined so that a point (A, B) is feasible (i.e., leads to a strictly positive value of $\Sigma_{n+1}\beta_{n+1}$ in its difference equation) if and only if $B < G_3(A)$.

These curves divide $\mathbb{R}^2_+$ into three mutually exclusive regions, the union of which comprises all of $\mathbb{R}^2_+$. First, define the infeasible region $R_5 = \{(A, B) : A \geq 0, B \geq G_3(A)\}$. Second, define region $R_1 = \{(A, B) : A \geq 0, G_2(A) < B < G_3(A)\}$. In this region, $F(A, B)$ is always to the left and higher than (A, B) and the given expression for F implies that starting the iteration $(A_{n+1}, B_{n+1}) = F(A_n, B_n)$ in R_1 will eventually lead the sequence to enter the infeasible region R_5. Finally, define $R_2 = \{(A, B) : A \geq 0, 0 = G_1(A) < B \leq G_2(A)\}$. Note that any sequence (A_n, B_n) always remains feasible, as shown by Caldentey and Stacchetti (2010). Hence, any candidate (A_0, B_0) must lie in R_2.

To put the above more clearly in the setting of Figure 1 in Caldentey and Stacchetti (2010), note that in our case, the function $G_1(A)$ is shifted identically downward to zero. This completely eliminates the regions R_3 and R_4 in their plot. Furthermore, the stationary point $(\hat{A}, \hat{B})$ in our case is defined by the point at which $G_1(A)$ and $G_2(A)$ intersect. This point is precisely $(0, 0)$. That is, with no flow of information to the insider, in the stationary limit the trader perfectly reveals her information and the market maker faces no residual uncertainty.

To complete the proof, we need to find a curve $\mathcal{C} \subset R_2$ such that $(0, 0) \in \mathcal{C}$ and $F(\mathcal{C}) \subset \mathcal{C}$. Note further that because for sequences in R_2, we have $(A_{n+1}, B_{n+1}) = (F_A(A_n, B_n), F_B(A_n, B_n)) < (A_n, B_n)$ we know that such a curve must be strictly increasing. Furthermore, because F is continuous we know that such a curve exists. This curve can be defined by an increasing function $0 \leq \psi(A) \leq G_2(A)$ with $\psi(0) = 0$ so that $\mathcal{C} = \{(A, B) : A \geq 0, B = \psi(A)\}$.

Clearly if we take $B_0 = \psi(A_0)$ then the associated sequence always lies in $\mathcal{C}$ and we have $(A_n, B_n) \downarrow 0$, the stationary point. Hence, returning to the definitions of A_0 and B_0, this implies that we need to set

$$\frac{\beta_0 \Sigma_0}{\sqrt{\Sigma_Z}} = \psi(\Sigma_0)$$

$$\Rightarrow \beta_0 = \Psi(\Sigma_0) \equiv \frac{\sqrt{\Sigma_Z}}{\Sigma_0}\psi(\Sigma_0).$$

The claimed inequality holds because it was shown above that $0 \leq \psi(A) \leq G_2(A)$. $\square$

We will now proceed with analyzing the behavior of the deviation profit in eq. (64) as Δ shrinks. Recall that for a positive function of $\Delta > 0$, h^Δ, we define $h^\Delta \sim O(\Delta^p)$ for $p > 0$ as $\Delta \to 0$ if and only if

$$\limsup_{\Delta\to 0} \frac{h^\Delta}{\Delta^p} < \infty.$$

The following result establishes a property of the limiting behavior of β_0^Δ as $\Delta \to 0$.

Lemma 7. *We have*

$$\beta_0^\Delta \sim O(\Delta^{3/4}), \quad \text{as } \Delta \to 0.$$

Proof. From Lemma 6 we know

$$\begin{aligned}
0 < \beta_0^\Delta &= \frac{\sqrt{\Sigma_Z}}{\Sigma_0}\psi(\Sigma_0) \\
&= \frac{\sigma_z\sqrt{\Delta}}{\Sigma_0}\psi(\Sigma_0) \\
&\leq \frac{\sigma_z\sqrt{\Delta}}{\Sigma_0}[1-\rho]^{1/4}\sqrt{\Sigma_0} \\
&= \frac{\sigma_z}{\sqrt{\Sigma_0}}\sqrt{\Delta}[1-e^{-r\Delta}]^{1/4}.
\end{aligned}$$

We have $[1-e^{-r\Delta}]^{1/4} \sim O(\Delta^{1/4})$ from which it follows that

$$0 \leq \limsup_{\Delta\to 0} \frac{\beta_0^\Delta}{\Delta^{3/4}} \leq \frac{\sigma_z}{\sqrt{\Sigma_0}} \limsup_{\Delta\to 0} \frac{\sqrt{\Delta}[1-e^{-r\Delta}]^{1/4}}{\Delta^{3/4}} < \infty,$$

which establishes the result. $\square$

Proposition 7. *For all $\Delta > 0$ sufficiently small we have*

$$\frac{\bar{\Pi}_{d0} - \bar{\Pi}_0}{\Delta} = \frac{-\frac{1}{2}\frac{(\beta_0^\Delta)^3\Sigma_0^2}{\sigma_z^2\Delta} + (1-e^{-r\Delta})c}{\Delta} > 0.$$

Proof. From Lemma 7 we know $\beta_0^\Delta \sim O(\Delta^{3/4})$. It follows that $(\beta_0^\Delta)^3 \sim O(\Delta^{9/4})$. Hence

$$\liminf_{\Delta\to 0}\left(-\frac{1}{2}\frac{(\beta_0^\Delta)^3\Sigma_0^2}{\sigma_z^2\Delta^2}\right) = -\frac{1}{2}\frac{\Sigma_0^2}{\sigma_z^2}\limsup_{\Delta\to 0}\frac{(\beta_0^\Delta)^3}{\Delta^2}$$

$$= -\frac{1}{2}\frac{\Sigma_0^2}{\sigma_z^2}\limsup_{\Delta\to 0}\frac{(\beta_0^\Delta)^3}{\Delta^{9/4}}\Delta^{1/4}$$
$$= 0.$$

Similarly since $\frac{1-e^{-r\Delta}}{\Delta} \to rc$, we have

$$\liminf_{\Delta\to 0}\frac{(1-e^{-r\Delta})c}{\Delta} = rc.$$

Combining the above two results yields

$$\liminf_{\Delta\to 0}\frac{\bar{\Pi}_{d0} - \bar{\Pi}_0}{\Delta} = rc > 0,$$

which establishes the result. □

C Unobservable acquisition in Kyle (1985)

In this appendix, we study unobservable information acquisition in the continuous-time version of Kyle (1985). As in the case of a random horizon considered in the body of the paper, we first establish that the only candidate pure strategy equilibria involve acquisition at the beginning of the trading game. We then show that when the cost of information is sufficiently high, there does not exist a pure strategy equilibrium with acquisition at time zero, and therefore there does not exist any pure strategy equilibrium. The intuition for this result is a stronger version of the deviation argument in the text – when the cost of information is sufficiently high, it is profitable for the trader to deviate by *never* acquiring information.

Trading takes place on the interval $[0,1]$. That is, following Kyle (1985) and Back (1992) the terminal date T is fixed and normalized to 1. As the end of trading the asset pays off v, where $v \sim N(0,\Sigma_0)$. There is a risk-neutral strategic trader who maximizes her expected trading profit and at any time can choose to pay $c > 0$ to observe v. As above, we assume that she trades smoothly so that her cumulative holdings $X_t = \int_0^t \theta_u\, du$ for trading rate θ_u. There are noise traders whose cumulative holdings Z_t follow a Brownian motion with variance σ_Z^2. Prices are set by a risk neutral market maker who observes the cumulative order flow $Y_t = X_t + Z_t$ and sets $P_t = \mathbb{E}[v|\mathcal{F}_{Pt}]$. Because we are considering unobservable acquisition, public information includes only the order flow so $\mathcal{F}_t^P$ is the augmentation of the filtration $\sigma(\{Y_t\})$. As is standard, we search for equilibria in which the time-t asset price is a smooth function of the cumulative order flow up to that point. Thus, in the notation in the text, the conditional expected payoff is $f(t, \xi, N) = \xi$, where $\xi = v$ and there is no public signal, and the endogenous state variable is simply the cumulative order flow, $p_t = Y_t$.

An identical argument to that in Lemmas 1 and 2 immediately implies that any pure-strategy equilibrium cannot involve acquisition after time zero. Otherwise the strategic trader has a profitable deviation from any conjectured equilibrium since she can preemptively acquire information, trade against an unresponsive pricing rule, and make unboundedly large profits.

Now, suppose that there is a pure strategy equilibrium in which the trader acquires information immediately at $t = 0$. In such an equilibrium, the pricing rule and the trader's post-acquisition value function are those from Kyle (1985) (or the special case of Back (1992) with normally-distributed payoff).

Hence,

$$P(t, y) = \lambda y,$$

where $\lambda = \sqrt{\frac{\Sigma_0}{\sigma_Z^2}}$ and

$$J^v(t, y) = \frac{1}{2}\sqrt{\frac{\sigma_Z^2}{\Sigma_0}}\left(v - \sqrt{\frac{\Sigma_0}{\sigma_Z^2}}y\right)^2 + \frac{1}{2}\sqrt{\Sigma_0 \sigma_Z^2}(1 - t).$$

The ex-ante (gross) expected profit from being informed is therefore

$$\begin{aligned}
\overline{J}(0, 0) &= \mathbb{E}[J^v(0, 0)] \\
&= \frac{1}{2}\sqrt{\frac{\sigma_Z^2}{\Sigma_0}}\Sigma_0 + \frac{1}{2}\sqrt{\Sigma_0 \sigma_Z^2} \\
&= \sqrt{\Sigma_0 \sigma_Z^2}.
\end{aligned}$$

We would like to compare the above to the expected payoff if the trader deviates, remains uninformed for the duration of the trading game, and trades against the posited equilibrium price function. Following the argument of Lemma 1 in Back (1992), it is straightforward to construct the trader's value function under this deviation. It is

$$J^{d,U}(t, y) = J^0(t, y) = \frac{1}{2}\sqrt{\frac{\Sigma_0}{\sigma_Z^2}}y^2 + \frac{1}{2}\sqrt{\Sigma_0 \sigma_Z^2}(1 - t).$$

At time zero, this becomes

$$J^{d,U}(0, 0) = \frac{1}{2}\sqrt{\Sigma_0 \sigma_Z^2},$$

which is half of the informed trader's ex-ante gross profit.

Notice that because the trader is risk-neutral, under the given price process her optimal trading profit depends only on her conditional expectation of the asset value (as well as the cumulative order flow and the calendar time), so her value function is identical to that for a trader who acquires information but for whom the realized signal does not change her prior expectation ($v = 0$). The fact that this trading profit is not zero is a consequence of the fact that in a dynamic model the trader expects profitable trading opportunities to arise in the future when the realized noise trade pushes the price away from zero.

With the above value functions, we arrive immediately at the following.

Proposition 8. *Suppose that the asset price is a smooth function of calendar time and the cumulative order flow, Assumption 2 holds, and* $c > \frac{1}{2}\sqrt{\Sigma_0 \sigma_Z^2}$*. Then, there does not exist an equilibrium in which information acquisition follows a pure strategy.*

Proof. We have already ruled out pure strategy equilibria in which the trader acquires with some delay. Suppose that there is an equilibrium in which the trader acquires information immediately. The analysis of Back (1992) implies that under the stated assumptions the asset price and the informed trader's value functions must be of the form above. Consider the trader's expected profit from unobservably deviating and never acquiring information. The net expected profit from this deviation is

$$J^{d,u}(0,0) - (\bar{J}(0,0) - c) = c - \frac{1}{2}\sqrt{\Sigma_0 \sigma_Z^2} > 0,$$

so the trader is better off by undertaking the deviation. □

Importantly, note that the above result implies that when $c > \frac{1}{2}\sqrt{\Sigma_0 \sigma_Z^2}$, the financial market equilibrium in Kyle (1985) cannot arise as a consequence of endogenous information acquisition.