Strong-form efficiency with monopolistic insiders

Minh Chau and Dimitri Vayanos

Strong-Form Efficiency with Monopolistic Insiders

Minh Chau and Dimitri Vayanos*

November 7, 2005

Abstract

We study market efficiency in an infinite-horizon model with a monopolistic insider. The insider can trade with a competitive market maker and noise traders, and observes privately the expected growth rate of asset dividends. In the absence of the insider, this information would be reflected in prices only after a long series of dividend observations. The insider chooses, however, to reveal the information very quickly, within a time converging to zero as the market approaches continuous trading. Although the market converges to strong-form efficiency, the insider's profits do not converge to zero.

*Chau is from the ESSEC Business School, e-mail chau@essec.fr, and Vayanos is from the London School of Economics, CEPR and NBER, e-mail d.vayanos@lse.ac.uk. We thank Bruno Biais, Denis Gromb, Maureen O'Hara, Anna Pavlova, Jean-Pierre Zigrand, and participants at the EFA meetings for helpful comments. A previous version of this paper was circulated under the title "Positive Profits when Prices are Strongly Efficient."

1 Introduction

How efficient are financial markets in incorporating information? This question has generated a large body of research, both theoretical and empirical. Some studies focus on information known to all market participants, such as earnings and macroeconomic announcements. Others consider private information held, for example, by corporate insiders. Fama (1970) uses the concept of strong-form efficiency to characterize a market where private information is fully reflected in prices.

Understanding how closely markets approximate the strong-form-efficiency ideal requires an analysis of the trading strategies of privately informed agents. If, for example, these agents trade aggressively, then their information should be reflected in prices quickly. In a seminal paper, Kyle (1985) provides the first analysis of strategic informed trading. He considers a monopolistic insider who can trade with competitive market makers in the presence of noise traders. When trading is continuous, the insider reveals her information slowly at a rate which is constant over time. Information is fully reflected in prices only at the end of the trading session, just before the time when it is to be announced publicly.

Kyle and much of the subsequent literature, assume that the insider receives information only once, at the beginning of the trading session. This can be a good description of a corporate insider who knows the content of an earnings announcement. In other cases, however, the assumption that the insider receives information repeatedly might be more appropriate. For example, the insider could be a hedge fund or proprietary-trading desk generating a continuous flow of private information on a stock through their superior research.

In this paper we consider an infinite-horizon, steady-state model where a monopolistic insider receives information in each period. The information concerns the expected growth rate of asset dividends, and in the absence of the insider would be reflected in prices only after a long series of dividend observations. Quite surprisingly, however, in the presence of the insider the information is reflected very quickly, within a time converging to zero as the market approaches continuous trading (i.e., as the time between consecutive transactions goes to zero). Thus, a market with a monopolistic insider can be arbitrarily close to strong-form efficiency, in contrast to Kyle. We also show that the insider's profits do not converge to zero despite the market converging to efficiency. Thus, markets can be almost efficient and yet offer sizeable returns to information acquisition.

While our results are in sharp contrast to previous literature, they are not driven by any peculiar modelling assumptions. We consider an economy with a dividend-paying risky asset and an exogenous riskless rate. As in Kyle, the agents are a risk-neutral insider who can submit a market order in each period, noise traders who submit *i.i.d.* market orders, and a risk-neutral competitive market maker who sets a price to absorb the aggregate order. We depart from Kyle in assuming an infinite horizon and new private information arriving in each period. To model private information we follow Wang (1993), setting the dividend growth rate to the sum of a time-varying drift, observed only by the insider, and *i.i.d.* noise. The drift represents the expected growth rate and follows a random walk.[1]

Our model has a unique linear equilibrium that we compute in closed form when the market approaches continuous trading. To characterize the speed of information revelation, we examine how quickly the price adjusts to an innovation in the drift process. In the absence of the insider, the adjustment occurs slowly (i.e., at a finite rate) even in the continuous-time limit. Intuitively, the market maker can learn about the drift only by observing the dividend. In the continuous-time limit the dividend process becomes a Brownian motion with drift, and it is well known that the drift cannot be fully inferred within any finite time. In the presence of the insider, however, the price adjustment occurs at a rate that converges to infinity, meaning that prices reflect private information almost instantly.

Why does the insider choose to reveal her information quickly? In general, the insider can minimize the price impact of a large order by breaking it into small pieces and "going down" the market maker's demand curve. When the market approaches continuous trading, the small orders can be placed within a short time interval without increasing the price impact. This allows the insider to exploit her information quickly and avoid the costs linked to impatience that are (i) the time-discounting of her profits and (ii) the revelation of her information through the dividend. Impatience cannot, however, provide a full explanation because the insider does not trade quickly in Kyle.[2] The additional element has to do with the time-pattern of information arrival. In Kyle, the insider receives information only once, at the beginning of the trading session. If she trades quickly, then the price impact of her trades will be large early on when her information is being revealed, and small afterwards when information has become symmetric. But this cannot be an equilibrium because the insider would then prefer to wait until the price impact gets small. In our model, by contrast, the price impact is constant over time, whether the insider trades quickly or not, because we are in a steady state where new private information always arrives. Given the constant price impact, impatience induces the insider to trade quickly.

[1]The random-walk assumption is for simplicity; introducing mean-reversion would only burden the notation without changing the results.

[2]Kyle assumes no impatience, but it is easy to introduce impatience in his model and show that the insider still trades slowly. We return to this point in Footnote 11.

Although the market converges to strong-form efficiency, the insider's profits do not converge to zero. Intuitively, the insider's profit margin per share decreases as the market approaches efficiency. This is, however, compensated by the fact that the insider can trade more as trading opportunities become more frequent.

To assess the practical significance of our results, we calibrate the model. We select the noise in the dividend process so that the price-adjustment to new information in the absence of the insider has a half-life of four months. We find that the half-life drops to only four days in the presence of an insider who can trade every ten minutes. Thus, our model implies that markets can be much closer to strong-form efficiency than suggested by previous literature. For example, if the information of Kyle's insider is to be announced publicly in four months, then the insider will take two months to incorporate half of it into the price.

Holden and Subrahmanyam (1992), Foster and Vishwanathan (1996) and Back, Cao and Willard (2000) introduce multiple insiders into Kyle's model (where information is received only once and there is a finite horizon). In Holden and Subrahmanyam all insiders receive the same information, and reveal it almost immediately as the market approaches continuous trading. Thus, the market becomes strong-form efficient but for a different reason than in our model - each insider tries to exploit her information before others do. An additional difference with our model is that the insiders' profits converge to zero as the market approaches efficiency. In Foster and Vishwanathan the insiders receive imperfectly correlated signals, and information revelation slows down because of a "waiting-game" effect, whereby each insider attempts to learn the others' signals. Back, Cao and Willard formulate the problem directly in continuous time. They show, in particular, that when signals are imperfectly correlated, information is not fully reflected in prices until the end of the trading session because of the waiting-game effect.[3]

Back and Pedersen (1998) consider a continuous-time, finite-horizon model where a monopolistic insider receives a flow of private information during the trading session. They show that the insider reveals her information slowly, and thus the market is not strong-form efficient. To ensure existence of equilibrium, they endow the insider with a stock of initial information in addition to the subsequent flow. It is because of this stock that information revelation is slow as in Kyle.

Taub, Bernhardt and Seiler (2005) consider a discrete-time, infinite-horizon model with multiple insiders receiving information in each period. They propose a general method to compute the equilibrium, using functional-analysis techniques. Their main focus is to solve the complicated

[3]See also Back (1992) for a general continuous-time formulation of the single-insider problem.

problem of infinite regress, and they do not characterize the equilibrium close to the continuous-time limit. Also, their model is different than ours in many respects. For example, the insiders' private information concerns a liquidating dividend paid at a stochastic time when the economy ends, and no information about the dividend is revealed publicly beforehand.

The rest of this paper is organized as follows. Section 2 presents the model. Section 3 determines the equilibrium in the general discrete-time case. Section 4 considers the behavior of the equilibrium when the market approaches continuous trading, and establishes our main results. Section 5 calibrates the model and Section 6 concludes. All proofs are in the Appendix.

2 Model

Time is continuous and goes from $-\infty$ to ∞. Trading takes place at a set of discrete times $\{\ell h\}_{\ell \in \mathbb{Z}}$, where h is a positive constant. We refer to time ℓh as period ℓ. There is a consumption good and two financial assets. The first asset is a riskless bond with an exogenous, continuously compounded rate of return r. The return on the bond between two consecutive periods is e^{rh}. The second asset is a risky stock that pays a dividend $d_\ell h$ in period ℓ. The dynamics of the dividend rate d_ℓ are given by:

$$d_\ell = d_{\ell-1} + g_{\ell-1} h + \varepsilon_{d,\ell} \tag{1}$$

$$g_\ell = g_{\ell-1} + \varepsilon_{g,\ell}, \tag{2}$$

where the shocks $\varepsilon_{d,\ell}$ and $\varepsilon_{g,\ell}$ are independent of each other and across periods, and normally distributed with mean zero and variance $\sigma_d^2 h$ and $\sigma_g^2 h$ respectively. The variable g_ℓ is the drift of the dividend rate. Our specification for d_ℓ and g_ℓ ensures that when the time h between consecutive periods goes to zero, g_ℓ converges to an arithmetic Brownian motion and d_ℓ to an arithmetic Brownian motion with stochastic drift.

There are three types of traders: a market maker, an insider, and noise traders. The market maker behaves competitively, while the insider is strategic.[4] Both are risk-neutral, discount the future at rate r, and have the utility function

$$E_\ell \left[\sum_{\ell'=\ell}^{\infty} c_{\ell'}^j e^{-r(\ell'-\ell)h} \middle| \mathcal{F}_\ell^j \right], \tag{3}$$

[4] As in Kyle (1985), the assumption of a competitive market maker can be viewed as a reduced form for multiple market makers competing in a Bertrand fashion.

where c_ℓ^j denotes consumption in period ℓ, $\mathcal{F}_\ell^j$ denotes the information set, and the superscript j is m for the market maker and i for the insider. Under the utility function (3), agents are indifferent as to the timing of consumption, and value a cash-flow stream according to the present value (PV) of expected cash flows discounted at rate r. The insider's private information consists of the dividend drift g_ℓ. In period ℓ, the insider can trade with the market maker via a market (i.e., price-inelastic) order that we denote by x_ℓ. Noise traders also submit a market order that we denote by u_ℓ. The noise traders' order is independent of the dividend process, independent across periods, and normally distributed with mean 0 and variance $\sigma_u^2 h$. We adopt the convention that x_ℓ and u_ℓ are positive if the insider and the noise traders buy. As in Kyle, we assume that the market maker observes only the aggregate order $x_\ell + u_\ell$, and sets a price p_ℓ at which he is willing to take the other side of the trade.

The timing of events in period ℓ is as follows. First, the insider and the noise traders submit their orders. Next, the dividend rate d_ℓ is publicly revealed, and the insider observes the drift g_ℓ. The market maker then sets a price p_ℓ at which he is willing to take the other side of the trade. Finally, the asset pays the dividend $d_\ell h$ and agents consume.[5]

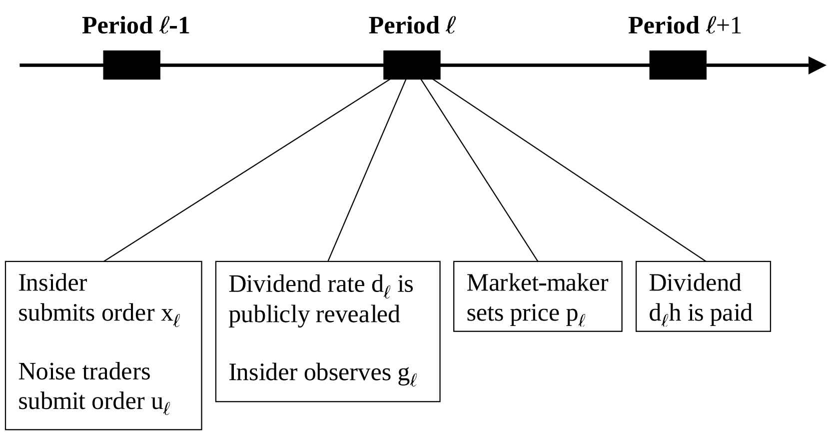

Figure 1: Timing of events in period ℓ.

Competitive behavior ensures that the market maker sets the price equal to his marginal valuation. The latter is equal to the PV of expected dividends conditional on the market maker's information. Therefore,

$$p_\ell = E_\ell \left[\sum_{\ell'=\ell}^{\infty} d_{\ell'} h e^{-r(\ell'-\ell)h} \middle| \mathcal{F}_\ell^m \right].$$

[5] Alternatively, we could assume that d_ℓ and g_ℓ are observed first, and then orders are submitted. This would complicate the notation without changing the results.

From Equations (1) and (2), the price is equal to

$$\begin{aligned}
p_\ell &= \sum_{\ell'=\ell}^{\infty} E_\ell(d_{\ell'}|\mathcal{F}_\ell^m) h e^{-r(\ell'-\ell)h} \\
&= \sum_{\ell'=\ell}^{\infty} \left[d_\ell + E_\ell(g_\ell|\mathcal{F}_\ell^m)(\ell'-\ell)h\right] h e^{-r(\ell'-\ell)h} \\
&= \left[d_\ell + E_\ell(g_\ell|\mathcal{F}_\ell^m)\frac{he^{-rh}}{1-e^{-rh}}\right]\frac{h}{1-e^{-rh}}. \qquad (4)
\end{aligned}$$

From now on, we set $\widehat{g}_\ell \equiv E_\ell(g_\ell|\mathcal{F}_\ell^m)$ to denote the market maker's expectation of g_ℓ in period ℓ. Note that the expectation is evaluated after the market maker observes the dividend rate d_ℓ and order flow $x_\ell + u_\ell$.

The insider's valuation for the asset is equal to the PV of expected dividends conditional on her information. Denoting the valuation in period ℓ by v_ℓ, an analogous calculation as for the market maker implies that

$$v_\ell = \left(d_\ell + g_\ell \frac{he^{-rh}}{1-e^{-rh}}\right)\frac{h}{1-e^{-rh}}, \qquad (5)$$

since the insider observes g_ℓ perfectly. The insider's optimization problem in period ℓ is to choose a sequence of market orders $\{x_{\ell'}\}_{\ell'\geq\ell}$ to maximize the PV of expected profits. Expected profits for the insider in period ℓ are equal to her order x_ℓ times the difference between valuation and price. Therefore, the insider's objective is

$$E\left[\sum_{\ell'=\ell}^{\infty} x_{\ell'}\left(v_{\ell'} - p_{\ell'}\right) e^{-r(\ell'-\ell)h} \middle| F_\ell^i\right]. \qquad (6)$$

3 Equilibrium

3.1 Candidate Strategies

An equilibrium consists of a trading strategy $\{x_\ell\}_{\ell\in\mathbb{Z}}$ for the insider and a pricing strategy $\{p_\ell\}_{\ell\in\mathbb{Z}}$ for the market maker such that

- The insider maximizes the PV of expected profits, given the price process generated by the market maker's strategy.

- The market maker sets prices equal to the PV of expected dividends, where the expectation is conditional on information revealed by the insider's strategy.

We look for an equilibrium in which strategies are linear functions of the state variables. Additionally, we assume that we are in a steady state where these functions are time-independent.[6] The state variables in period ℓ are the dividend rate d_ℓ, the drift g_ℓ, and the market maker's expectation of the drift $\widehat{g}_\ell$. The price quoted by the market maker is given by Equation (4), i.e.,

$$p_\ell = \left(d_\ell + \widehat{g}_\ell \frac{he^{-rh}}{1-e^{-rh}} \right) \frac{h}{1-e^{-rh}}. \tag{7}$$

We conjecture that the expectation $\widehat{g}_\ell$ evolves according to

$$\widehat{g}_\ell = \widehat{g}_{\ell-1} + \lambda_d \left(d_\ell - (d_{\ell-1} + \widehat{g}_{\ell-1} h) \right) + \lambda_x (x_\ell + u_\ell), \tag{8}$$

for two constants λ_d and λ_x. Intuitively, the market maker updates the expectation held in period $\ell - 1$ because of two pieces of information learned in period ℓ: the dividend rate d_ℓ and the order flow $x_\ell + u_\ell$. The latter is informative provided that the insider's order depends on the drift. We conjecture that the insider's order is proportional to the market maker's error in forecasting the drift, i.e.,

$$x_\ell = \beta(g_{\ell-1} - \widehat{g}_{\ell-1}), \tag{9}$$

for a constant β. The forecast error is evaluated as of period $\ell - 1$ because when the insider submits her order she only knows $g_{\ell-1}$ and not g_ℓ.

To solve for the equilibrium, we derive a set of equations linking the coefficients λ_d, λ_x, and β. These equations follow from the market maker's inference problem and the insider's optimization problem.

3.2 Market Maker's Inference

The market maker's inference problem consists in forming a belief about the drift g_ℓ, given the history of dividend rates and order flows up to period ℓ. To solve this problem, we use recursive

[6] In assuming that time goes from $-\infty$ to ∞, we are implicitly assuming convergence to the steady state. To show convergence, we can start the economy at a finite time and endow the market maker with a normal prior on the drift. We can then examine the limit of the coefficients that characterize the linear equilibrium when time goes to ∞. While a comprehensive analysis of convergence is beyond the scope of this paper, we have established numerically local convergence, i.e., when the initial condition (the variance of the market maker's prior) is close to its steady-state value.

(Kalman) filtering. That is, we derive the belief about g_ℓ given (i) the belief about $g_{\ell-1}$ held in period $\ell-1$, and (ii) the new information learned in period ℓ consisting of the dividend rate d_ℓ and order flow $x_\ell + u_\ell$.

Suppose that in period $\ell - 1$ the market maker takes $g_{\ell-1}$ to be normal with mean $\widehat{g}_{\ell-1}$ and variance Σ_g^2. Then, we show in Appendix A that the belief about g_ℓ is also normal. The mean of the normal distribution is given by

$$\widehat{g}_\ell = \widehat{g}_{\ell-1} + \lambda_d \left(d_\ell - (d_{\ell-1} + \widehat{g}_{\ell-1}h)\right) + \lambda_x(x_\ell + u_\ell),$$

i.e., Equation (8), with

$$\lambda_d = \frac{\Sigma_g^2 \sigma_u^2 h}{\Sigma_g^2 \left(\beta^2\sigma_d^2 + \sigma_u^2 h^2\right) + \sigma_d^2\sigma_u^2 h} \tag{10}$$

$$\lambda_x = \frac{\beta\Sigma_g^2 \sigma_d^2}{\Sigma_g^2 \left(\beta^2\sigma_d^2 + \sigma_u^2 h^2\right) + \sigma_d^2\sigma_u^2 h}. \tag{11}$$

Intuitively, the market maker starts with a prior mean for g_ℓ, which is $\widehat{g}_{\ell-1}$ since $g_\ell = g_{\ell-1} + \varepsilon_{g,\ell}$. The prior mean is then adjusted to reflect the information learned from d_ℓ and $x_\ell + u_\ell$. The adjustment is proportional to the surprises in these signals, i.e., the differences between the signals and their prior means. The prior mean of

$$d_\ell = d_{\ell-1} + g_{\ell-1}h + \varepsilon_{d,\ell}$$

is $d_{\ell-1} + \widehat{g}_{\ell-1}h$, while that of

$$x_\ell + u_\ell = \beta(g_{\ell-1} - \widehat{g}_{\ell-1}) + u_\ell$$

is zero. In Appendix A we show that the variance of the market maker's belief about g_ℓ is

$$\mathrm{Var}(g_\ell | F_\ell^m) = \frac{\Sigma_g^2\sigma_d^2\sigma_u^2 h}{\Sigma_g^2 \left(\beta^2\sigma_d^2 + \sigma_u^2 h^2\right) + \sigma_d^2\sigma_u^2 h} + \sigma_g^2 h. \tag{12}$$

In steady state the variance must be constant over time, implying that $\mathrm{Var}(g_\ell | F_\ell^m) = \Sigma_g^2$. This yields the equation

$$\Sigma_g^2 \left(\Sigma_g^2 - \sigma_g^2 h\right)\left(\beta^2\sigma_d^2 + \sigma_u^2 h^2\right) - \sigma_g^2\sigma_d^2\sigma_u^2 h^2 = 0. \tag{13}$$

3.3 Insider's Optimization

The insider maximizes the objective in Equation (6). Using Equation (7), we can simplify this objective to

$$E\left[\sum_{\ell'=\ell}^{\infty} x_{\ell'}\left(g_{\ell'}-\widehat{g}_{\ell'}\right)e^{-r(\ell'-\ell)h}\middle|\,\mathcal{F}_{\ell}^{i}\right]. \tag{14}$$

Thus, a buy order in period ℓ ($x_\ell > 0$) is profitable to the insider if the market maker underestimates the drift ($g_\ell - \widehat{g}_\ell > 0$). When the insider submits her order in period ℓ, she only knows the market maker's forecast error up to period $\ell - 1$. We conjecture that the insider's value function in period ℓ is a quadratic function of the forecast error in period $\ell - 1$, i.e.,

$$V(g_{\ell-1},\widehat{g}_{\ell-1}) = B(g_{\ell-1}-\widehat{g}_{\ell-1})^2 + C,$$

for two constants B and C. The Bellman equation is

$$V(g_{\ell-1},\widehat{g}_{\ell-1}) = \max_{x_\ell} E\left[x_\ell\left(g_\ell-\widehat{g}_\ell\right) + e^{-rh}V(g_\ell,\widehat{g}_\ell)\middle|\,\mathcal{F}_{\ell}^{i}\right].$$

To evaluate the right-hand side, we need to compute the market maker's forecast error as of period ℓ. This is

$$\begin{aligned} g_\ell - \widehat{g}_\ell &= (g_{\ell-1}+\varepsilon_{g,\ell}) - \left[\widehat{g}_{\ell-1} + \lambda_d\left(d_\ell - (d_{\ell-1}+\widehat{g}_{\ell-1}h)\right) + \lambda_x(x_\ell+u_\ell)\right] \\ &= (1-\lambda_d h)(g_{\ell-1}-\widehat{g}_{\ell-1}) - \lambda_d\varepsilon_{d,\ell} - \lambda_x(x_\ell+u_\ell) + \varepsilon_{g,\ell}, \end{aligned} \tag{15}$$

where the first step follows from Equations (2) and (8), and the second from Equation (1). Substituting into the Bellman equation, we find

$$\begin{aligned} &B(g_{\ell-1}-\widehat{g}_{\ell-1})^2 + C = \max_{x_\ell}\Big\{x_\ell\left[(1-\lambda_d h)(g_{\ell-1}-\widehat{g}_{\ell-1}) - \lambda_x x_\ell\right] \\ &+e^{-rh}\left[B\left[\left[(1-\lambda_d h)(g_{\ell-1}-\widehat{g}_{\ell-1}) - \lambda_x x_\ell\right]^2 + \lambda_d^2\sigma_d^2 h + \lambda_x^2\sigma_u^2 h + \sigma_g^2 h\right] + C\right]\Big\}. \end{aligned} \tag{16}$$

The first-order condition yields

$$x_\ell = \beta(g_{\ell-1}-\widehat{g}_{\ell-1}),$$

i.e., Equation (9), with

$$\beta = \frac{(1-\lambda_d h)\left(1-2e^{-rh}B\lambda_x\right)}{2\lambda_x\left(1-e^{-rh}B\lambda_x\right)}. \tag{17}$$

Substituting for x_ℓ in equation (16), we can determine B and C:

$$B = \frac{(1-\lambda_d h)^2}{4\lambda_x \left(1 - e^{-rh} B \lambda_x\right)} \tag{18}$$

$$C = \frac{e^{-rh} B \left(\lambda_d^2 \sigma_d^2 + \lambda_x^2 \sigma_u^2 + \sigma_g^2\right) h}{1 - e^{-rh}}. \tag{19}$$

3.4 Existence and Uniqueness

Our conjectured equilibrium is characterized by the six parameters λ_d, λ_x, Σ_g^2, β, B and C. These are the solution to the system of six equations (10), (11), (13) and (17)-(19). In Appendix B we show that the system has a unique solution, which in addition satisfies the insider's second-order condition. This implies that there exists a unique equilibrium of the conjectured form.

Proposition 1 *There exists a unique equilibrium of the form conjectured in Equations (7)-(9).*

4 Near-Continuous Trading

Our main results concern the behavior of the equilibrium when the market approaches continuous trading, i.e., the time h between consecutive periods goes to zero.[7] To better illustrate the results, we start with the benchmark case where the insider is prevented from trading due to exogenous reasons. The market maker then quotes infinite depth (i.e., price not sensitive to order flow), but still learns about the drift by observing the dividend. Inference is characterized by the parameters λ_d and Σ_g^2, and these are given by Equations (10) and (13) with the insider's trading intensity β set to zero. To distinguish with the case where the insider is trading, we denote λ_d and Σ_g^2 by $\overline{\lambda}_d$ and $\overline{\Sigma}_g^2$, respectively.

Proposition 2 *When the insider is not trading, the asymptotic behavior of the equilibrium is*

[7] Although our main results concern the continuous-time limit, we avoid formulating the model directly in continuous time. Starting with discrete time and then taking the limit has the advantage of illustrating how the equilibrium changes with the trading frequency. Discrete time is also important when we calibrate the model. Finally, in continuous-time formulations (e.g., Kyle (1985) and Back (1992)) insider trading is a flow, i.e., proportional to dt. In our model, by contrast, insider trading is of order larger than dt, and this is central to our strong-form efficiency result.

characterized by

$$\lim_{h\to 0} \overline{\lambda}_d = \frac{\sigma_g}{\sigma_d}$$

$$\lim_{h\to 0} \overline{\Sigma}_g^2 = \sigma_g \sigma_d.$$

Proposition 2 implies that in the absence of the insider, the equilibrium close to the continuous-trading limit is qualitatively similar to that away from the limit. In particular, the market maker's uncertainty about the drift, characterized by $\overline{\Sigma}_g^2$, remains bounded away from zero. The intuition is that when h goes to zero the dividend process converges to a Brownian motion with drift. Because the drift is changing over time, it cannot become known to the market maker.

We next consider the case where the insider is trading, and establish our main results.

Proposition 3 *When the insider is trading, the asymptotic behavior of the equilibrium is characterized by*

$$\lim_{h\to 0} \frac{\lambda_d}{\sqrt{h}} = \frac{\sigma_g^2}{\sigma_d^2 \sqrt{r}} \tag{20}$$

$$\lim_{h\to 0} \lambda_x = \frac{\sigma_g}{\sigma_u} \tag{21}$$

$$\lim_{h\to 0} \frac{\Sigma_g^2}{\sqrt{h}} = \frac{\sigma_g^2}{\sqrt{r}} \tag{22}$$

$$\lim_{h\to 0} \frac{\beta}{\sqrt{h}} = \frac{\sigma_u \sqrt{r}}{\sigma_g} \tag{23}$$

$$\lim_{h\to 0} B = \frac{\sigma_u}{2\sigma_g} \tag{24}$$

$$\lim_{h\to 0} C = \frac{\sigma_g \sigma_u}{r}. \tag{25}$$

Proposition 3 shows that in the presence of the insider, the equilibrium close to the continuous-trading limit and that away from the limit have very different properties. In particular, the parameter Σ_g^2 characterizing the market maker's uncertainty about the drift is approximately $(\sigma_g^2/\sqrt{r})\sqrt{h}$ for small h. Therefore, it converges to zero when h goes to zero, implying that the information asymmetry between the insider and the market maker vanishes. Put differently, a market with a monopolistic insider can become strong-form efficient when the trading frequency is sufficiently large.

An alternative way to characterize strong-form efficiency is through the speed at which private information is incorporated into prices. Suppose that at time zero the insider learns that the drift g_0 is different from the market maker's expectation $\widehat{g}_0$. To measure how quickly the insider's information is incorporated into prices, we can examine the dynamics of the market maker's expectation. Conditional on all information available in the economy at time zero, the drift is expected to remain equal to g_0 because it follows a random walk. The market maker's expectation $\widehat{g}_\ell$ is then expected to converge to g_0 over time. This convergence is in expectation only, conditional on all available time-zero information (which in our model is the insider's information $\mathcal{F}_0^i$), because the drift keeps changing over time. Thus, the convergence concerns the variable $E\left(\widehat{g}_\ell \middle| \mathcal{F}_0^i\right)$. To evaluate this variable when trading is almost continuous, we fix a calendar time t, corresponding to period $\ell = t/h$, and consider the limit $\widehat{G}_t \equiv \lim_{h\to 0} E\left(\widehat{g}_{\frac{t}{h}} \middle| \mathcal{F}_0^i\right)$. Proposition 4 characterizes how $\widehat{G}_t$ varies over time.

Proposition 4 *When the insider is not trading,*

$$\widehat{G}_t = g_0 + e^{-\frac{\sigma_g}{\sigma_d}t}(\widehat{g}_0 - g_0).$$

When the insider is trading, $\widehat{G}_t = g_0$ *for* $t > 0$.

Proposition 4 shows that in the absence of the insider, information about the drift is incorporated into prices slowly (on average). For small h, the market maker's expectation converges to g_0 at the finite rate σ_g/σ_d. By contrast, in the insider's presence, information about the drift is incorporated very quickly. For small h, the market maker's expectation reaches g_0 within any positive time t, and not only when t goes to infinity. Thus, the rate of convergence to g_0 becomes infinite. This is, of course, consistent with Proposition 3: insider trading can result in strong-form efficiency when the trading frequency is large.

To understand the intuition for strong-form efficiency, we consider the insider's trading strategy. Recall that the insider submits an order proportional to the market maker's forecast error, with the proportionality parameter β interpreted as the trading intensity. The parameter β is a key determinant of the speed at which prices incorporate information. In Kyle (1985), β is of order h, and prices incorporate information within a calendar time not converging to zero. When, however, β is larger than order h, the insider's trades reveal more information, and the calendar time converges to zero. Proposition 3 implies that β in our model is of order $\sqrt{h}$ and thus larger than h.[8]

Why does the insider trade quickly in our model? To answer this question, we examine the determinants of the insider's order size. In general, a large order is costly to the insider because it generates an adverse price impact. To minimize the impact, the insider can trade slowly, breaking her order into small pieces and "going down" the market maker's demand curve. Slow trading, however, generates costs linked to impatience: by realizing her profits quickly, the insider can avoid (i) time-discounting and (ii) the public revelation of her information through the dividend.

When the trading frequency is large, the trade-off between price impact and impatience disappears. Indeed, the insider can squeeze all small pieces of a large order into a short time interval. Therefore, she can can execute the order quickly without increasing the price impact and without incurring costs linked to impatience.[9]

That impatience is necessary for the insider to trade quickly can be seen formally as follows. Impatience is eliminated when (i) there is no time-discounting, i.e., the interest rate r is zero, and (ii) no information is revealed publicly through the dividend, i.e., the noise σ_d^2 in the dividend process is infinite. Proposition 3 shows that when $r > 0$, β is of order $\sqrt{h}$, regardless of whether σ_d^2 is infinite or not. Therefore, a positive interest rate induces the insider to trade quickly. In Appendix C we consider the opposite case where $r = 0$.[10] When σ_d^2 is finite, we show that β is of order $h^{\frac{2}{3}}$ which is larger than h. Therefore, the public revelation of information induces the insider to trade quickly, even in the absence of time-discounting. When, however, σ_d^2 is infinite, we show that the insider prefers β to be as close to zero as possible. Therefore, a patient insider prefers to trade slowly.

While impatience is necessary for the insider to trade quickly, it is not sufficient. This can be seen by contrasting our model with Kyle. Kyle assumes no impatience because there is no time-discounting and no public revelation of information until a final period. It is easy, however,

[8] Other properties of β are as in Kyle. For example, the insider trades more aggressively when there is more noise trading (σ_u large), or when the market maker expects her to have less private information (σ_g small).

[9] See, however, Vayanos (2001) where a strategic hedger goes down the demand curve slowly, even in the continuous-time limit. Suppose, for example, that the market expects the hedger to sell 100 shares over ten hours, at a rate of ten shares per hour. If the hedger sells all 100 shares over the first hour, this will exceed the market's expectation of ten shares. Therefore, the market will increase its estimate of the hedger's inventory, expect more future sales from the hedger, and set a lower price for the 100 shares. By contrast, an insider can sell 100 shares over one hour at the same price as over ten hours. The difference with the hedger is that the market expects a zero average order from the insider, both over one and over ten hours. Therefore, the updating generated by the 100-share order is independent of the time it takes to complete the order. See also Spiegel and Subrahmanyam (1995) where the market expects non-zero orders from rational liquidity traders.

[10] For $r = 0$, we define the insider's objective as the long-run average of per-period payoffs. This type of objective is standard in the literature on repeated games with no discounting. See, for example, Fudenberg and Tirole (1991).

to introduce impatience in his model and show that the insider still trades slowly.[11]

The crucial difference with Kyle has to do with the time-pattern of information arrival. Kyle's model is non-stationary in that the insider receives information only once, at the beginning of the trading session. If the insider trades quickly, market depth will be small early on when her information is being revealed, and large afterwards when information has become symmetric. But this cannot be an equilibrium because the insider would prefer to wait until depth increases. Our model, by contrast, is stationary because the insider always receives new private information. Stationarity ensures that market depth is constant over time, whether the insider trades quickly or not. Given the constant depth, the insider trades quickly because of impatience (generated by either time-discounting, or public revelation of information, or both).

Summarizing, our strong-form efficiency result is due to the combination of impatience and stationarity. When the insider is patient, she trades slowly. Likewise, in a non-stationary setting (e.g., Kyle or Back and Pedersen (1998)) trading occurs slowly even with an impatient insider.

We next consider the insider's trading profits. These are

$$\left[B(g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C\right] \frac{h^2 e^{-rh}}{(1-e^{-rh})^2} \equiv B'(g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C', \tag{26}$$

i.e., the value function times a scaling factor that was dropped for simplicity when writing the insider's objective as (14) instead of (6). When h goes to zero, $(g_{\ell-1} - \widehat{g}_{\ell-1})^2$ converges to zero because the market becomes strong-form efficient. From Proposition 3, however, C converges to the positive limit $\sigma_g \sigma_u / r$, implying that C' converges to $\sigma_g \sigma_u / r^3$. Thus, the insider's profits remain positive despite the market converging to strong-form efficiency. In some sense, this is natural: since the insider chooses to incorporate her information quickly into prices, this must guarantee her a larger payoff than trading slowly. At the same time, the result can appear paradoxical: how can the insider realize positive profits when prices reflect almost all of her information?

To address the paradox, we recall that the insider's profits in period ℓ are

$$x_\ell (g_{\ell-1} - \widehat{g}_{\ell-1}) \frac{h^2 e^{-rh}}{(1-e^{-rh})^2}.$$

The term $(g_{\ell-1} - \widehat{g}_{\ell-1})$ corresponds to the profit margin, and converges to zero when h goes to zero. Asymptotically it is of order Σ_g, which is of order $h^{\frac{1}{4}}$ from Proposition 3. The term x_ℓ corresponds

[11] More specifically, we can allow for a positive interest rate and a noisy signal about the asset value revealed in each period. The noise in the signal must be such that information is revealed slowly in the benchmark case where the insider is not trading, even in the continuous-time limit. The analysis is available upon request.

to the trading volume. Since β is of order $\sqrt{h}$, $x_\ell = \beta(g_{\ell-1} - \widehat{g}_{\ell-1})$ is of order $h^{\frac{3}{4}}$. Thus, the volume generated by the insider within a fixed time interval is of order $h^{\frac{3}{4}}/h$, and converges to infinity when h goes to zero. This explains the paradox of positive profits: the insider's profit margin per share goes to zero but the number of shares traded goes to infinity.[12]

Finally, note that the price impact of order flow, given by λ_x, remains finite even in the continuous-time limit. This might appear surprising because the insider's trading volume converges to infinity, and hence order flow is much more informative in our model than in Kyle. Because the market is close to efficient, however, the market maker faces little uncertainty, so a given amount of information has a smaller effect on the price.

5 Calibration

The results of the previous section are asymptotic, i.e., hold to any given degree of approximation by choosing a time h between consecutive periods close enough to zero. In this section we calibrate the model and examine how well the results hold for plausible values of h. We are interested, in particular, in how close the market is to strong-form efficiency.

The exogenous parameters in our model are the interest rate r, the time h between consecutive periods, and the variance parameters σ_d of the dividend process, σ_g of the drift process, and σ_u of the noise trading. We set r to 5%, but values in the interval $[0\%, 10\%]$ would change the times in Table 1 by less than 6%. We allow h to take two possible values: the insider can trade every ten minutes, or every three hours. Assuming 250 trading days per year and ten trading hours per day, the values for h are $10/(60 \times 10 \times 250) = 0.000067$ and $180/(60 \times 10 \times 250) = 0.0012$.

Before turning to the parameters σ_d, σ_g, and σ_u, we define and compute our measure of market efficiency. We measure efficiency through the speed at which private information is incorporated into prices. As in Proposition 4, we assume that at time zero the insider observes a drift g_0 different from the market maker's expectation $\widehat{g}_0$. We then examine how quickly the market maker's expectation converges to g_0. From the proof of Proposition 4, the convergence dynamics for a given h are

$$E\left(\widehat{g}_{\frac{t}{h}} \,\middle|\, \mathcal{F}_0^i\right) = g_0 + (1 - \lambda_d h - \lambda_x \beta)^{\frac{t}{h}} (\widehat{g}_0 - g_0). \tag{27}$$

[12]That trading volume converges to infinity in the continuous-time limit is not pathological. For example, volume is infinite in the basic Merton (1971) model, where a CRRA investor keeps a constant fraction of wealth in a risky asset and needs to rebalance continuously. Mathematically, the investor's volume is infinite because the Brownian motion has infinite variation.

Our measure of market efficiency is the time t_χ by which a given percentage χ of the insider's information is incorporated into prices. This time is defined by

$$E\left(\widehat{g}_{\frac{t_\chi}{h}} \middle| \mathcal{F}_0^i\right) = \chi g_0 + (1-\chi)\widehat{g}_0. \tag{28}$$

Comparing Equations (27) and (28), we find

$$t_\chi = \frac{h\log(1-\chi)}{\log(1-\lambda_d h - \lambda_x \beta)}. \tag{29}$$

We next observe that t_χ is independent of two of the three parameters σ_d, σ_g and σ_u. Indeed, consider the solution $(\lambda_d, \lambda_x, \Sigma_g, \beta, B, C)$ to Equations (10), (11), (13) and (17)-(19). If σ_u in these equations is multiplied by $z > 0$, the new solution is $(\lambda_d, \lambda_x/z, \Sigma_g, z\beta, zB, zC)$. Equation (29) then implies that t_χ stays constant, meaning that t_χ is independent of σ_u. Intuitively, when noise trading increases, prices tend to become less informative, but the insider trades more aggressively restoring the same informativeness. Similarly, if σ_d and σ_g are multiplied by the same $z > 0$, the new solution to Equations (10), (11), (13) and (17)-(19) becomes $(\lambda_d, z\lambda_x, z\Sigma_g, \beta/z, zB, zC)$. Equation (29) then implies that t_χ stays constant, meaning that t_χ depends on σ_d and σ_g only through their ratio. Therefore, our results are robust to any choice of σ_u, and to choices of σ_d and σ_g that generate the same ratio.

To calibrate σ_g/σ_d, we consider the speed of information revelation in the absence of the insider, when the only source of information about the drift is the dividend. The parameter σ_g/σ_d is the signal-to-noise ratio, and measures the extent to which the dividend process is informative. The time $\bar{t}_\chi$ by which χ percent of an innovation in the drift process is incorporated into prices is

$$\bar{t}_\chi = \frac{h\log(1-\chi)}{\log(1-\bar{\lambda}_d h)}. \tag{30}$$

We calibrate σ_g/σ_d through the time $\bar{t}_{0.5}$ by which prices reflect half of the information (i.e., the half-life of the convergence dynamics). We allow σ_g/σ_d to take two possible values: 12 and 1.5. In the first case the half-life $\bar{t}_{0.5}$ is approximately 15 days, and in the second it is approximately four months.[13]

Table 1 compares the half-life $\bar{t}_{0.5}$ in the absence of the insider to the half-life $t_{0.5}$ in her presence. The table shows that the insider speeds information revelation by an order of magnitude.

[13]The half-life depends on h, but the dependence appears after the second decimal digit.

	$\bar{t}_{0.5}$ (days)	$t_{0.5}$ (days)	
		$h = 0.0012$	$h = 0.000067$
$\sigma_g/\sigma_d = 12$	14.44	2.81	1.06
$\sigma_g/\sigma_d = 1.5$	115.52	10.59	3.70

Table 1: Speed of information revelation.

Consider, for example, the case where the price adjustment to new information has a half-life of 115.52 days in the insider's absence. The insider reduces this to 10.59 days (i.e., less than a tenth) if she can trade every three hours, and 3.70 days (i.e., less than a thirtieth) if she can trade every ten minutes. We should emphasize that the insider can choose to reveal her information slowly, stretching the half-life closer to 115.52 days. Our main result, however, is that she prefers to reveal it quickly. When she can trade every ten minutes, for example, a half-life of 3.70 days ensures a minimal price impact, while allowing her to reap the benefits associated to impatience.

Using our calibration, we can also examine how the insider's trading profits depend on h. Recall from Equation (26) that profits are

$$\left[B(g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C\right] \frac{h^2 e^{-rh}}{(1-e^{-rh})^2} \equiv B'(g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C'.$$

To evaluate (B', C'), we must select values for (σ_u, σ_g), and for simplicity set both parameters to one.[14]

	B'		C'	
	$h = 0.0012$	$h = 0.000067$	$h = 0.0012$	$h = 0.000067$
$\sigma_g/\sigma_d = 12$	196.0	199.4	7482	7937
$\sigma_g/\sigma_d = 1.5$	198.9	199.8	8084	8060

Table 2: Insider's trading profits.

Table 2 reports the values of (B', C'). When $\sigma_g/\sigma_d = 12$, both B' and C' increase when h decreases, and thus the insider's profits increase. This reinforces our result that the insider makes positive profits even in the limit when the market becomes strong-form efficient. When $\sigma_g/\sigma_d = 1.5$, C' decreases when h decreases, but it still converges to a positive limit.

[14]The parameters (σ_u, σ_g) can be calibrated through the insider's trading volume and the bid-ask spread. See Chau (2002) for an example of such a calibration.

6 Conclusion

In this paper we consider a discrete-time, infinite-horizon model with a monopolistic insider. The insider observes the expected growth rate of asset dividends in each period, and can trade with competitive market makers in the presence of noise traders. Our main result is that when the market approaches continuous trading, the insider's information is reflected in prices almost immediately. This is especially surprising given that in the absence of the insider, the information would be reflected only after a long series of dividend observations. We also show that although the market converges to strong-form efficiency, the insider's profits do not converge to zero.

Our results have two important implications. First, markets can be close to strong-form efficiency even in the presence of monopolistic insiders. Second, despite being close to efficiency, markets can offer significant returns to information acquisition. These implications are in sharp contrast to previous literature, and are not driven by any peculiar assumptions in our model. Indeed, the main difference with Kyle (1985) is that we assume an infinite horizon, with new private information generated in each period. To model private information in infinite horizon, we adopt the information structure in Wang (1993).

The insider in our model can be best interpreted as a hedge fund or proprietary-trading desk, generating a continuous flow of private information through their superior research. Given that there are multiple such agents, one might question the assumption of a monopolistic insider. The work of Holden and Subrahmanyam (1992), Foster and Vishwanathan (1996) and Back, Cao and Willard (2000) shows, however, that competing insiders generally reveal their information faster than a monopolistic one. Thus, our strong-form efficiency result is likely to carry through with multiple insiders. Of course, this conjecture needs to be verified, and this could be an interesting extension of our research. The main technical difficulty is that combining multiple insiders with repeated information arrival generates an infinite-regress problem when the insiders' signals are imperfectly correlated. Taub, Bernhardt and Seiler (2005) develop a technique for dealing with this problem, however, and one could possibly use it to study the continuous-time limit.

A Market-Maker's Inference

Suppose that conditional on information up to period $\ell - 1$, the market maker believes that $g_{\ell-1}$ is normal with mean $\widehat{g}_{\ell-1} = E(g_{\ell-1}|\mathcal{F}^m_{\ell-1})$ and variance $\Sigma_g^2 = \mathrm{Var}(g_{\ell-1}|\mathcal{F}^m_{\ell-1})$. For notational simplicity, we omit the conditioning set $\mathcal{F}^m_{\ell-1}$ in the rest of this appendix because all moments are conditional. The signals observed by the market maker in period ℓ are

$$d_\ell = d_{\ell-1} + g_{\ell-1}h + \varepsilon_{d,\ell}$$

and

$$x_\ell + u_\ell = \beta(g_{\ell-1} - \widehat{g}_{\ell-1}) + u_\ell.$$

Because all variables are jointly normal, the market maker's posterior about $g_{\ell-1}$ is of the form

$$\begin{aligned} g_{\ell-1} &= E(g_{\ell-1}) + \lambda_d\left(d_\ell - E\left(d_\ell\right)\right) + \lambda_x\left(x_\ell + u_\ell - E\left(x_\ell + u_\ell\right)\right) + \eta_\ell \\ &= \widehat{g}_{\ell-1} + \lambda_d\left(d_\ell - \left(d_{\ell-1} + \widehat{g}_{\ell-1}h\right)\right) + \lambda_x\left(x_\ell + u_\ell\right) + \eta_\ell, \end{aligned} \tag{A.1}$$

where λ_d and λ_x are two constants, and η_ℓ is a normal random variable with mean zero and independent of d_ℓ and $x_\ell + u_\ell$. The posterior about $g_\ell = g_{\ell-1} + \varepsilon_{g,\ell}$ is as in Equation (A.1) with η_ℓ replaced by $\eta_\ell + \varepsilon_{g,\ell}$.

To compute λ_d and λ_x, we take the covariance of both sides of Equation (A.1) with d_ℓ and $x_\ell + u_\ell$:

$$\mathrm{Cov}\left(d_\ell, g_{\ell-1}\right) = \mathrm{Cov}\left(d_\ell, \lambda_d d_\ell + \lambda_x\left(x_\ell + u_\ell\right)\right), \tag{A.2}$$

$$\mathrm{Cov}\left(x_\ell + u_\ell, g_{\ell-1}\right) = \mathrm{Cov}\left(x_\ell + u_\ell, \lambda_d d_\ell + \lambda_x\left(x_\ell + u_\ell\right)\right). \tag{A.3}$$

Since

$$\mathrm{Cov}\left(d_\ell, g_{\ell-1}\right) = \mathrm{Cov}\left(d_{\ell-1} + g_{\ell-1}h + \varepsilon_{d,\ell}, g_{\ell-1}\right) = \mathrm{Var}\left(g_{\ell-1}\right)h = \Sigma_g^2 h, \tag{A.4}$$

$$\mathrm{Cov}\left(x_\ell + u_\ell, g_{\ell-1}\right) = \mathrm{Cov}\left(\beta(g_{\ell-1} - \widehat{g}_{\ell-1}) + u_\ell, g_{\ell-1}\right) = \beta\mathrm{Var}\left(g_{\ell-1}\right) = \beta\Sigma_g^2, \tag{A.5}$$

$$\mathrm{Var}\left(d_\ell\right) = \mathrm{Var}\left(d_{\ell-1} + g_{\ell-1}h + \varepsilon_{d,\ell}\right) = \mathrm{Var}\left(g_{\ell-1}\right)h^2 + \mathrm{Var}\left(\varepsilon_{d,\ell}\right) = \Sigma_g^2 h^2 + \sigma_d^2 h,$$

$$\mathrm{Cov}\left(d_\ell, , x_\ell + u_\ell\right) = \mathrm{Cov}\left(d_{\ell-1} + g_{\ell-1}h + \varepsilon_{d,\ell}, \beta(g_{\ell-1} - \widehat{g}_{\ell-1}) + u_\ell\right) = \beta\mathrm{Var}\left(g_{\ell-1}\right)h = \beta\Sigma_g^2 h,$$

$$\mathrm{Var}\left(x_\ell + u_\ell\right) = \mathrm{Var}\left(\beta(g_{\ell-1} - \widehat{g}_{\ell-1}) + u_\ell\right) = \beta^2\mathrm{Var}\left(g_{\ell-1}\right) + \mathrm{Var}\left(u_\ell\right) = \beta^2\Sigma_g^2 + \sigma_u^2 h,$$

we can write Equations (A.2) and (A.3) as

$$\Sigma_g^2 h = \lambda_d(\Sigma_g^2 h^2 + \sigma_d^2 h) + \lambda_x \beta \Sigma_g^2 h,$$

$$\beta\Sigma_g^2 = \lambda_d \beta \Sigma_g^2 h + \lambda_x(\beta^2\Sigma_g^2 + \sigma_u^2 h).$$

The solution to this linear system is given by Equations (10) and (11). Therefore, the posterior expectation of g_ℓ is as in Equation (8).

The posterior variance of g_ℓ is

$$\mathrm{Var}(g_\ell) = \mathrm{Var}(\eta_\ell + \varepsilon_{g,\ell}) = \mathrm{Var}(\eta_\ell) + \sigma_g^2 h. \tag{A.6}$$

To compute the variance of η_ℓ, we take the variance of both sides of Equation (A.1). Since η_ℓ is independent of d_ℓ and $x_\ell + u_\ell$, we have

$$\begin{aligned}
\mathrm{Var}(\eta_\ell) &= \mathrm{Var}(g_{\ell-1}) - \mathrm{Var}\left(\lambda_d d_\ell + \lambda_x\left(x_\ell + u_\ell\right)\right) \\
&= \mathrm{Var}(g_{\ell-1}) - \lambda_d \mathrm{Cov}\left(d_\ell, \lambda_d d_\ell + \lambda_x\left(x_\ell + u_\ell\right)\right) - \lambda_x \mathrm{Cov}\left(x_\ell + u_\ell, \lambda_d d_\ell + \lambda_x\left(x_\ell + u_\ell\right)\right) \\
&= \mathrm{Var}(g_{\ell-1}) - \lambda_d \mathrm{Cov}\left(d_\ell, g_{\ell-1}\right) - \lambda_x \mathrm{Cov}\left(x_\ell + u_\ell, g_{\ell-1}\right) \\
&= \Sigma_g^2 - \lambda_d \Sigma_g^2 h - \lambda_x \beta \Sigma_g^2,
\end{aligned}$$

where the third step follows from (A.2) and (A.3), and the fourth from (A.4) and (A.5). Using (10) and (11) to substitute for λ_d and λ_x, and plugging back into (A.6), we find (12).

B Proof of Propositions 1-4

Proof of Proposition 1: We will show that the system of Equations (10), (11), (13) and (17)-(19) has a unique solution, which also satisfies the insider's second-order condition. We will reduce the system to a single equation in β. Equation (18) can be written as

$$B^2 - \frac{e^{rh}}{\lambda_x} B + \frac{e^{rh}(1-\lambda_d h)^2}{4\lambda_x^2} = 0.$$

This quadratic equation in B has the two solutions

$$B = \frac{e^{rh}}{2\lambda_x}\left[1 \pm \sqrt{1 - e^{-rh}(1-\lambda_d h)^2}\right].$$

The solution with the plus sign cannot be part of a solution to the overall system. Indeed, Equation (11) implies that $\lambda_x \beta > 0$, which from Equation (17) means that

$$\frac{1 - 2e^{-rh} B \lambda_x}{1 - e^{-rh} B \lambda_x} > 0.$$

This is violated by the solution with the plus sign. Therefore, the only possible solution for B is

$$B = \frac{e^{rh}}{2\lambda_x}\left[1 - \sqrt{1 - e^{-rh}(1-\lambda_d h)^2}\right]. \tag{B.1}$$

Plugging into Equation (17), we find

$$\frac{\lambda_x \beta}{1-\lambda_d h} = \frac{\sqrt{1 - e^{-rh}(1-\lambda_d h)^2}}{1 + \sqrt{1 - e^{-rh}(1-\lambda_d h)^2}}. \tag{B.2}$$

Substituting for λ_d and λ_x using Equations (10) and (11), we can write Equation (B.2) as

$$\frac{\Sigma_g^2\beta^2}{\Sigma_g^2\beta^2 + \sigma_u^2 h} = \frac{\sqrt{1 - e^{-rh}\left[\frac{\Sigma_g^2\beta^2\sigma_d^2 + \sigma_d^2\sigma_u^2 h}{\Sigma_g^2\left(\beta^2\sigma_d^2 + \sigma_u^2 h^2\right) + \sigma_d^2\sigma_u^2 h}\right]^2}}{1 + \sqrt{1 - e^{-rh}\left[\frac{\Sigma_g^2\beta^2\sigma_d^2 + \sigma_d^2\sigma_u^2 h}{\Sigma_g^2\left(\beta^2\sigma_d^2 + \sigma_u^2 h^2\right) + \sigma_d^2\sigma_u^2 h}\right]^2}}. \tag{B.3}$$

We can reduce Equation (B.3) to one in the single unknown β by substituting for Σ_g^2 as a function of β. This can be done using Equation (13), which is quadratic in Σ_g^2 and has a unique positive solution. The solution is a decreasing function of β, converges to $\sigma_g^2 h$ when β goes to ∞, and to a value $\overline{\Sigma}_g^2 > \sigma_g^2 h$ when β goes to zero.

To show that the system of Equations (10), (11), (13) and (17)-(19) has a solution, we note that the left-hand side (LHS) of Equation (B.3) (in which Σ_g^2 is an implicit function of β) converges to one when β goes to ∞, and to zero when β goes to zero. By contrast, the right-hand side (RHS) converges to values strictly between zero and one in both cases. Therefore, Equation (B.3) has a solution $\beta \in (0, \infty)$. From this solution, we can deduce Σ_g^2, λ_d, λ_x, B and C using Equations (13), (10), (11), (B.1) and (19), respectively. The insider's second-order condition is the requirement that the problem (16) is concave, i.e., $1 - e^{-rh}B\lambda_x > 0$. This inequality is satisfied because of Equation (B.1).

To show that the solution is unique, we will show that the LHS of Equation (B.3) is increasing in β, while the RHS is decreasing. The LHS is increasing in β if $\Sigma_g^2\beta^2\sigma_d^2$ is increasing. Equation (13) implies that

$$\Sigma_g^2\beta^2\sigma_d^2 = \frac{\sigma_g^2\sigma_d^2\sigma_u^2 h^2}{\Sigma_g^2 - \sigma_g^2 h} - \Sigma_g^2\sigma_u^2 h^2. \tag{B.4}$$

Since Σ_g^2 is decreasing in β, $\Sigma_g^2\beta^2\sigma_d^2$ is increasing. The RHS of Equation (B.3) is decreasing in β if

$$Z \equiv \frac{\Sigma_g^2\beta^2\sigma_d^2 + \sigma_d^2\sigma_u^2 h}{\Sigma_g^2\left(\beta^2\sigma_d^2 + \sigma_u^2 h^2\right) + \sigma_d^2\sigma_u^2 h}$$

is increasing. Using Equation (B.4) to eliminate β, we find

$$Z = \frac{\sigma_d^2 + \sigma_g^2 h^2 - \Sigma_g^2 h}{\sigma_d^2}.$$

Since Σ_g^2 is decreasing in β, Z is increasing. ∎

Proof of Proposition 2: When $\beta = 0$, Equations (10) and (13) become

$$\overline{\lambda}_d = \frac{\overline{\Sigma}_g^2}{\overline{\Sigma}_g^2 h + \sigma_d^2} \tag{B.5}$$

and

$$\overline{\Sigma}_g^2\left(\overline{\Sigma}_g^2 - \sigma_g^2 h\right) - \sigma_g^2\sigma_d^2 = 0, \tag{B.6}$$

respectively. These equations have a unique solution. The solution for $h = 0$ is $\overline{\Sigma}_g^2 = \sigma_g\sigma_d$ and $\overline{\lambda}_d = \sigma_g/\sigma_d$. By continuity, this is also the limit of the solution when h goes to zero. ∎

Proof of Proposition 3: We first solve the system of Equations (13) and (B.3) in the unknowns Σ_g^2 and β. Setting $\Sigma_g^2 \equiv S_g^2\sqrt{h}$ and $\beta \equiv b\sqrt{h}$, we can write the system as

$$S_g^2\left(S_g^2 - \sigma_g^2\sqrt{h}\right)(b^2\sigma_d^2 + \sigma_u^2 h) - \sigma_g^2\sigma_d^2\sigma_u^2 = 0$$

and

$$\frac{S_g^2 b^2}{S_g^2 b^2\sqrt{h} + \sigma_u^2} = \frac{\sqrt{\frac{1}{h}\left[1 - e^{-rh}\left[1 + \frac{S_g^2\sigma_u^2 h^{\frac{3}{2}}}{\sigma_d^2\sigma_u^2 + S_g^2 b^2\sigma_d^2\sqrt{h}}\right]^{-2}\right]}}{1 + \sqrt{1 - e^{-rh}\left[1 + \frac{S_g^2\sigma_u^2 h^{\frac{3}{2}}}{\sigma_d^2\sigma_u^2 + S_g^2 b^2\sigma_d^2\sqrt{h}}\right]^{-2}}}.$$

For $h = 0$, the system becomes

$$S_g^4 b^2 - \sigma_g^2\sigma_u^2 = 0$$

and

$$\frac{S_g^2 b^2}{\sigma_u^2} = \sqrt{r},$$

and has the solution $S_g^2 = \sigma_g^2/\sqrt{r}$ and $b = \sigma_u\sqrt{r}/\sigma_g$. By continuity, this is also the limit of the solution when h goes to zero. This establishes the limits (22) and (23) since

$$\lim_{h\to 0}\frac{\Sigma_g^2}{\sqrt{h}} = \lim_{h\to 0} S_g^2 = \frac{\sigma_g^2}{\sqrt{r}}$$

and

$$\lim_{h\to 0}\frac{\beta}{\sqrt{h}} = \lim_{h\to 0} b = \frac{\sigma_u\sqrt{r}}{\sigma_g}.$$

To prove the limits (20), (21), (24), and (25), we write Equations (10), (11), (19), and (B.1) in terms of S_g^2 and b, and use the limits of S_g^2 and b. ∎

Proof of Proposition 4: Taking expectations in Equation (8), conditional on the insider's time-zero information, we find

$$E\left(\widehat{g}_\ell|\,\mathcal{F}_0^i\right) = E\left(\widehat{g}_{\ell-1}|\,\mathcal{F}_0^i\right)+\lambda_d\left[E\left(d_\ell|\,\mathcal{F}_0^i\right) - \left(E\left(d_{\ell-1}|\,\mathcal{F}_0^i\right) + E\left(\widehat{g}_{\ell-1}|\,\mathcal{F}_0^i\right)h\right)\right]+\lambda_x E\left(x_\ell|\,\mathcal{F}_0^i\right) \quad \text{(B.7)}$$

for $\ell > 0$. Since

$$E\left(d_\ell|\,\mathcal{F}_0^i\right) = d_0 + g_0\ell h$$

for $\ell \geq 0$, and

$$E\left(x_\ell|\,\mathcal{F}_0^i\right) = E\left(\beta(g_{\ell-1} - \widehat{g}_{\ell-1})|\,\mathcal{F}_0^i\right) = \beta\left(E\left(g_{\ell-1}|\,\mathcal{F}_0^i\right) - E\left(\widehat{g}_{\ell-1}|\,\mathcal{F}_0^i\right)\right) = \beta\left(g_0 - E\left(\widehat{g}_{\ell-1}|\,\mathcal{F}_0^i\right)\right)$$

for $\ell > 0$, we can write Equation (B.7) as

$$E\left(\widehat{g}_\ell|\,\mathcal{F}_0^i\right) = E\left(\widehat{g}_{\ell-1}|\,\mathcal{F}_0^i\right) + (\lambda_d h + \lambda_x\beta)\left[g_0 - E\left(\widehat{g}_{\ell-1}|\,\mathcal{F}_0^i\right)\right].$$

Therefore,

$$E\left(\widehat{g}_\ell|\,\mathcal{F}_0^i\right) = g_0 + (1-\lambda_d h - \lambda_x\beta)^\ell(\widehat{g}_0 - g_0).$$

To prove the proposition, we need to determine the limit of $(1-\lambda_d h-\lambda_x\beta)^{\frac{t}{h}}$ when h goes to zero. When the insider is not trading, the limit is $e^{-\frac{\sigma_g}{\sigma_d}h}$ because λ_d converges to σ_g/σ_d and $\beta = 0$. When the insider is trading, the limit is zero because β is of order $\sqrt{h}$ and λ_x of order 1. ∎

C No Time-Discounting

For $r = 0$, we define the insider's objective as the long-run average of per-period payoffs. The average payoff over the L periods starting from ℓ is

$$\Pi_L \equiv \frac{1}{L} E\left[\sum_{\ell'=\ell}^{\ell+L-1} x_{\ell'} \left(g_{\ell'} - \widehat{g}_{\ell'} \right) \middle| \mathcal{F}_\ell^i \right],$$

and the long-run average is

$$\Pi \equiv \lim_{L \to \infty} \Pi_L.$$

To compute the equilibrium, we assume that the insider follows the linear strategy (9) for some constant β. Then, λ_d, λ_x and Σ_g are given as a function of β by Equations (10), (11) and (13), respectively. Moreover, the insider chooses β to maximize Π, taking λ_d and λ_x as given.

To determine Π, we first compute Π_L. We conjecture that

$$\Pi_L = B_L (g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C_L,$$

for two constants B_L and C_L. These constants satisfy the equation

$$B_L (g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C_L = \frac{1}{L} E\left[x_\ell \left(g_\ell - \widehat{g}_\ell \right) + (L-1) \left[B_{L-1} (g_\ell - \widehat{g}_\ell)^2 + C_{L-1} \right] \middle| \mathcal{F}_\ell^i \right],$$

where $g_\ell - \widehat{g}_\ell$ is given by Equation (15), and $x_\ell = \beta(g_{\ell-1} - \widehat{g}_{\ell-1})$. Substituting for $g_\ell - \widehat{g}_\ell$ and x_ℓ, we find

$$B_L = \frac{\beta \left(1 - \lambda_d h - \lambda_x \beta\right) + (L-1) B_{L-1} \left(1 - \lambda_d h - \lambda_x \beta\right)^2}{L}$$

$$C_L = \frac{(L-1) \left[B_{L-1} \left(\lambda_d^2 \sigma_d^2 + \lambda_x^2 \sigma_u^2 + \sigma_g^2 \right) + C_{L-1} \right]}{L}.$$

It is easy to check by induction, starting from $L = 1$, that

$$B_L = \frac{\beta \sum_{k=0}^{L-1} \left(1 - \lambda_d h - \lambda_x \beta\right)^{2k+1}}{L} \tag{C.1}$$

$$C_L = \frac{\beta \left(\lambda_d^2 \sigma_d^2 + \lambda_x^2 \sigma_u^2 + \sigma_g^2 \right) \sum_{k=0}^{L-2} (L-1-k) \left(1 - \lambda_d h - \lambda_x \beta\right)^{2k+1}}{L}. \tag{C.2}$$

Suppose that $\sigma_d^2 < \infty$. Equations (10) and (11) imply that

$$1 - \lambda_d h - \lambda_x \beta = \frac{\sigma_d^2 \sigma_u^2 h}{\Sigma_g^2 \left(\beta^2 \sigma_d^2 + \sigma_u^2 h^2\right) + \sigma_d^2 \sigma_u^2 h}.$$

Since $\sigma_d^2 < \infty$, we have $0 < 1 - \lambda_d h - \lambda_x \beta < 1$. Equations (C.1) and (C.2) then imply that $\lim_{L\to\infty} B_L = 0$ and

$$\begin{aligned}\Pi = \lim_{L\to\infty} C_L &= \beta \left(\lambda_d^2 \sigma_d^2 + \lambda_x^2 \sigma_u^2 + \sigma_g^2\right) \sum_{k=0}^{\infty} (1 - \lambda_d h - \lambda_x \beta)^{2k+1} \\ &= \frac{\beta (1 - \lambda_d h - \lambda_x \beta)}{1 - (1 - \lambda_d h - \lambda_x \beta)^2} \left(\lambda_d^2 \sigma_d^2 + \lambda_x^2 \sigma_u^2 + \sigma_g^2\right).\end{aligned}$$

Therefore, maximizing Π is equivalent to maximizing

$$\frac{\beta (1 - \lambda_d h - \lambda_x \beta)}{1 - (1 - \lambda_d h - \lambda_x \beta)^2}.$$

The first-order condition is

$$\begin{aligned}& 1 - \lambda_d h - 2\lambda_x \beta - (1 - \lambda_d h)(1 - \lambda_d h - \lambda_x \beta)^2 = 0 \\ \Leftrightarrow\ & (\lambda_x \beta)^2 = \lambda_d h \left[(1 - \lambda_d h - \lambda_x \beta)^2 + 1 - \lambda_d h - 2\lambda_x \beta\right]. \qquad \text{(C.3)}\end{aligned}$$

In equilibrium, λ_d, λ_x, Σ_g and β are the solution to the system of Equations (10), (11), (13) and (C.3). To determine the solution for small h, we set $\lambda_d = l_d h^{\frac{1}{3}}$, $\Sigma_g^2 \equiv S_g^2 h^{\frac{1}{3}}$ and $\beta \equiv b h^{\frac{2}{3}}$. It is then easy to check that the resulting system in l_d, λ_x, S_g and b has the solution

$$\begin{aligned} b &= \frac{2^{\frac{1}{3}} \sigma_u}{\sigma_d^{\frac{2}{3}} \sigma_g^{\frac{1}{3}}} \\ S_g^2 &= \frac{\sigma_g \sigma_u}{b} \\ \lambda_x &= \frac{b S_g^2}{\sigma_u^2} \\ l_d &= \frac{S_g^2}{\sigma_d^2} \end{aligned}$$

for $h = 0$. Continuity then implies that a solution for small h exists, and

$$\lim_{h\to 0} \frac{\beta}{h^{\frac{2}{3}}} = \lim_{h\to 0} b = \frac{2^{\frac{1}{3}} \sigma_u}{\sigma_d^{\frac{2}{3}} \sigma_g^{\frac{1}{3}}}.$$

Therefore, β is of order $h^{\frac{2}{3}}$.

Suppose next that $\sigma_d = \infty$. Equation (10) implies that $\lambda_d = 0$, and Equation (11) implies that

$$1 - \lambda_x \beta = \frac{\sigma_u^2 h}{\Sigma_g^2 \beta^2 + \sigma_u^2 h}.$$

When $\beta > 0$, we have $0 < 1 - \lambda_x \beta < 1$. Equations (C.1) and (C.2) then imply that $\lim_{L\to\infty} B_L = 0$ and

$$\Pi = \lim_{L\to\infty} C_L = \beta \left(\lambda_x^2 \sigma_u^2 + \sigma_g^2\right) \sum_{k=0}^{\infty} (1 - \lambda_x \beta)^{2k+1}$$

$$= \frac{\beta (1 - \lambda_x \beta)}{1 - (1 - \lambda_x \beta)^2} \left(\lambda_x^2 \sigma_u^2 + \sigma_g^2\right).$$

When $\beta = 0$, we have $\Pi = 0$ since Equations (C.1) and (C.2) imply that $B_L = C_L = 0$. Therefore, maximizing Π is equivalent to maximizing a function that is equal to $(1 - \lambda_x \beta)/(2 - \lambda_x \beta)$ if $\beta > 0$ and zero if $\beta = 0$. This function increases as β decreases to zero, and drops discontinuously to zero for $\beta = 0$. Therefore, the insider prefers to set β as close to zero as possible. Intuitively, since the insider is infinitely patient, she chooses to minimize price impact by spreading her trades maximally over time.

References

Back, Kerry, 1992, "Insider Trading in Continuous Time," *Review of Financial Studies,* 5, 387-409.

Back, Kerry, Henry Cao and Gregory Willard, 2000, "Imperfect Competition among Informed Traders," *Journal of Finance,* 55, 2117-2155.

Back, Kerry and Hal Pedersen, 1998, "Long-Lived Information and Intraday Patterns," *Journal of Financial Markets,* 1, 385-402.

Chau, Minh, 2002, "Market-Making when Traders Follow Dynamic Order-Placement Strategies," working paper, ESSEC.

Fama, Eugene, 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance,* 25, 383-417.

Foster, Douglas and S. Viswanathan, 1996, "Strategic Trading when Agents Forecast the Forecasts of Others," *Journal of Finance,* 51, 1437-1478.

Fudenberg, Drew and Jean Tirole, 1991, *Game Theory,* MIT Press.

Holden, Craig and Avanidhar Subrahmanyam, 1992, "Long-Lived Private Information and Imperfect Competition," *Journal of Finance,* 47, 247-270.

Kyle, Albert, 1985, "Continuous Auctions and Insider Trading," *Econometrica,* 53, 1315-1335.

Merton, Robert, 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory,* 3, 373-413.

Spiegel, Matthew and Avanidhar Subrahmanyam, 1995, "On Intraday Risk Premia," *Journal of Finance,* 50, 319-339.

Taub, Bart, Dan Bernardt and Peter Seiler, 2005, "Speculative Dynamics," working paper, University of Illinois.

Vayanos, Dimitri, 2001, "Strategic Trading in a Dynamic Noisy Market," *Journal of Finance,* 56, 131-171.

Wang, Jiang, 1993, "A Model of Intertemporal Asset Prices under Asymmetric Information," *Review of Economic Studies,* 60, 249-282.