Strong-Form Efficiency with Monopolistic Insiders

Minh Chau and Dimitri Vayanos*

August 7, 2006

*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, Andrew Ellul, Denis Gromb, Harrison Hong, Jonathan Lewellen, Maureen O'Hara, Andrew Patton, Anna Pavlova, Michela Verardo, Jiang Wang, Jean-Pierre Zigrand, an anonymous referee, and seminar participants at AFFI, EFA, ESSEC, and HEC for helpful comments and suggestions. A previous version of this paper was circulated under the title "Positive Profits when Prices are Strongly Efficient." Please address correspondence to Dimitri Vayanos, Department of Accounting and Finance, London School of Economics, Houghton Street, London WC2A 2AE, UK, tel 44-20-79556382.

Abstract

We study market efficiency in an infinite-horizon model with a monopolistic insider. The insider can trade with competitive market makers 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. Thus, the insider's information is "long-lived." Surprisingly, however, the monopolistic insider chooses to reveal her 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.

1 Introduction

How efficient are financial markets at 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 proprietary-trading desk, hedge fund, or mutual fund, 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's 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 stock 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), assuming that dividends revert to a time-varying mean observed only by the insider. The mean can be interpreted as the firm's underlying profitability, and reverts to a constant long-run value.

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 a shock in the firm's profitability. 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 profitability only by observing the dividend. In the continuous-time limit the dividend follows an Ito process, and since profitability enters in the drift it cannot be fully inferred within any finite time. In the presence of the insider, however, the 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, (ii) the revelation of her information through the dividend, and (iii) the obsolescence of her information through the mean-reversion in the firm's profitability. Impatience cannot, however, provide a full explanation because the insider does not trade quickly in Kyle.[1] 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 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.

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, offset 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 to parameters corresponding to a large-cap and a small-cap stock. In both cases, the insider speeds information revelation dramatically. As an example, suppose that the noise in the dividend process is such that in the insider's absence, the price adjustment to profitability shocks has a half-life of one year. Then, the insider reduces this to 1.5 day for the large-cap stock and 24 days for the small-cap stock. 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 one year, the insider takes six 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 from that in our model - each insider tries to exploit her information before others do. An additional difference from 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.[2]

Back and Pedersen (1998) consider a continuous-time, finite-horizon model in which 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 the 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 problem of infinite regress, and they do not characterize the equilibrium close to the continuous-time limit. Their work differs from ours in many other 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.

Wang (1993,1994) considers infinite-horizon models where a set of insiders receive information on a firm's dividends over time in steady state. In contrast to our paper (and all others mentioned in this section), Wang's insiders are risk-averse and competitive rather than risk-neutral and strategic. Absent risk-aversion, competitive insiders would reveal their information instantly. Risk-aversion, however, limits the insiders' trading aggressiveness, ensuring that information is revealed slowly.

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 derives empirical implications. Section 6 concludes, and 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 stock that pays a dividend $D_\ell h$ in period ℓ. The dividend rate D_ℓ reverts to a time-varying mean g_ℓ according to the process

$$D_\ell = D_{\ell-1} + \nu h(g_{\ell-1} - D_{\ell-1}) + \varepsilon_{D,\ell}, \tag{1}$$

where ν determines the reversion rate and $\varepsilon_{D,\ell}$ is an *i.i.d.* shock. The parameter g_ℓ can be interpreted as the firm's "true" underlying profitability, and is observable only to the insider. This parameter reverts to a constant value $\overline{g}$ according to the process

$$g_\ell = g_{\ell-1} + \kappa h(\overline{g} - g_{\ell-1}) + \varepsilon_{g,\ell}, \tag{2}$$

where κ determines the reversion rate and $\varepsilon_{g,\ell}$ is an *i.i.d.* shock. The shocks $\varepsilon_{D,\ell}$ and $\varepsilon_{g,\ell}$ are independent of each other and normally distributed with mean zero and variance $\sigma_D^2 h > 0$ and $\sigma_g^2 h > 0$, respectively. We assume that $\nu h \in (0,1)$ and $\kappa h \in [0,1)$, allowing in particular $\kappa = 0$ in which case g_ℓ follows a random walk. Our specification for D_ℓ and g_ℓ ensures that when the time

h between consecutive periods goes to zero, D_ℓ and g_ℓ converge to the Ito processes

$$dD_t = \nu(g_t - D_t)dt + \sigma_D dB_{D,t} \tag{3}$$

and

$$dg_t = \kappa(\overline{g} - g_t)dt + \sigma_g dB_{g,t}, \tag{4}$$

respectively, where the Brownian motions $B_{D,t}$ and $B_{g,t}$ are independent. Equations (3) and (4) are quite standard for modelling private information in infinite horizon, e.g., Wang (1993).

There are three types of traders: a market maker, an insider, and noise traders. The market maker behaves competitively, while the insider is strategic.[3] Both are risk-neutral, discount the future at rate r, and have the utility function

$$E\left[\sum_{\ell'=\ell}^{\infty} c_{\ell'}^j e^{-r(\ell'-\ell)h} \middle| \mathcal{F}_\ell^j\right], \tag{5}$$

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 (5), 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 profitability 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 $\sigma_u^2 h > 0$. 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 profitability 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.[4]

INSERT FIGURE 1 SOMEWHERE HERE

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,

$$p_\ell = E\left[\sum_{\ell'=\ell}^{\infty} D_{\ell'} h e^{-r(\ell'-\ell)h} \middle| \mathcal{F}_\ell^m\right]. \tag{6}$$

To evaluate this expression, we substitute $\{D_{\ell'}\}_{\ell'>\ell}$ from Equations (1) and (2). In Appendix A we show that this yields

$$p_\ell = A_0 D_\ell + A_1 E(g_\ell|\mathcal{F}_\ell^m) + A_2 \overline{g} \tag{7}$$

for three positive constants A_0, A_1 and A_2 that depend on h, r, ν and κ. Thus, the price is a linear and increasing function of the current dividend level D_ℓ, the market maker's expectation of current profitability $E(g_\ell|\mathcal{F}_\ell^m)$, and the long-run mean of profitability $\overline{g}$. From now on, we denote the market maker's expectation by $\widehat{g}_\ell \equiv E(g_\ell|\mathcal{F}_\ell^m)$. Note that this 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. This PV is given by the same equation as for the market maker, with the difference that the expectation $\widehat{g}_\ell$ is replaced by true value g_ℓ. Thus, the valuation v_ℓ in period ℓ is

$$v_\ell = A_0 D_\ell + A_1 g_\ell + A_2 \overline{g}, \tag{8}$$

for the constants A_0, A_1 and A_2 in Equation (7). 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

$$E\left[\sum_{\ell'=\ell}^{\infty} x_{\ell'} \left(v_{\ell'} - p_{\ell'}\right) e^{-r(\ell'-\ell)h} \middle| F_\ell^i\right]. \tag{9}$$

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.[5] The state variables in period ℓ are the dividend rate D_ℓ, the profitability g_ℓ, and the market maker's expectation of profitability $\widehat{g}_\ell$. The price quoted by the market maker is given by Equation (7), i.e.,

$$p_\ell = A_0 D_\ell + A_1 \widehat{g}_\ell + A_2 \overline{g}. \tag{10}$$

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

$$\widehat{g}_\ell = \overline{g} + (1 - \kappa h)(\widehat{g}_{\ell-1} - \overline{g}) + \lambda_D \left(D_\ell - (1 - \nu h) D_{\ell-1} - \nu h \widehat{g}_{\ell-1}\right) + \lambda_x (x_\ell + u_\ell), \tag{11}$$

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 profitability. We conjecture that the insider's order is proportional to the market maker's error in estimating profitability, i.e.,

$$x_\ell = \beta(g_{\ell-1} - \widehat{g}_{\ell-1}), \tag{12}$$

for a constant β. The 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 profitability g_ℓ, given the history of dividend rates and order flows up to period ℓ. To solve this problem, we use recursive (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 B that the belief about g_ℓ is also normal. The mean of the normal distribution is given by

$$\widehat{g}_\ell = (1-\kappa h)\widehat{g}_{\ell-1} + \kappa h\overline{g} + \lambda_D\left(D_\ell - (1-\nu h)D_{\ell-1} - \nu h\widehat{g}_{\ell-1}\right) + \lambda_x(x_\ell+u_\ell),$$

i.e., Equation (11), with

$$\lambda_D = \frac{(1-\kappa h)\Sigma_g^2\nu\sigma_u^2 h}{\Sigma_g^2\left(\beta^2\sigma_D^2+\nu^2\sigma_u^2h^2\right)+\sigma_D^2\sigma_u^2 h}, \tag{13}$$

$$\lambda_x = \frac{(1-\kappa h)\beta\Sigma_g^2\sigma_D^2}{\Sigma_g^2\left(\beta^2\sigma_D^2+\nu^2\sigma_u^2h^2\right)+\sigma_D^2\sigma_u^2 h}. \tag{14}$$

Intuitively, the market maker starts with a prior mean for g_ℓ, which is $(1-\kappa h)\widehat{g}_{\ell-1}+\kappa h\overline{g}$ since

$$g_\ell = (1-\kappa h)g_{\ell-1} + \kappa h\overline{g} + \varepsilon_{g,\ell}.$$

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

$$D_\ell = (1-\nu h)D_{\ell-1} + \nu h g_{\ell-1} + \varepsilon_{D,\ell}$$

is $(1-\nu h)D_{\ell-1}+\nu h\widehat{g}_{\ell-1}$, while that of

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

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

$$\mathrm{Var}(g_\ell|F_\ell^m) = \frac{(1-\kappa h)^2\Sigma_g^2\sigma_D^2\sigma_u^2 h}{\Sigma_g^2\left(\beta^2\sigma_D^2+\nu^2\sigma_u^2h^2\right)+\sigma_D^2\sigma_u^2 h} + \sigma_g^2 h. \tag{15}$$

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

$$\left(\Sigma_g^2-\sigma_g^2h\right)\left[\Sigma_g^2\left(\beta^2\sigma_D^2+\nu^2\sigma_u^2h^2\right)+\sigma_D^2\sigma_u^2h\right] - (1-\kappa h)^2\Sigma_g^2\sigma_D^2\sigma_u^2 h = 0. \tag{16}$$

3.3 Insider's Optimization

The insider maximizes the objective in Equation (9). Using Equation (10), we can simplify this objective to

$$E\left[\sum_{\ell'=\ell}^{\infty} x_{\ell'}\left(g_{\ell'}-\widehat{g}_{\ell'}\right) e^{-r(\ell'-\ell)h} \middle| \mathcal{F}_{\ell}^{i}\right]. \tag{17}$$

Thus, a buy order in period ℓ ($x_\ell > 0$) is advantageous to the insider if the market maker underestimates the firm's profitability ($g_\ell - \widehat{g}_\ell > 0$). When the insider submits her order in period ℓ, she only knows the market maker's estimation error up to period $\ell-1$. We conjecture that the insider's value function in period ℓ, evaluated at the time of order submission, is a quadratic function of the period $\ell - 1$ estimation error, 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 estimation error as of period ℓ. This is

$$\begin{aligned}
g_\ell - \widehat{g}_\ell &= (1-\kappa h) g_{\ell-1} + \kappa h \overline{g} + \varepsilon_{g,\ell} \\
&\quad - \left[(1-\kappa h)\widehat{g}_{\ell-1} + \kappa h \overline{g} + \lambda_D \left(D_\ell - (1-\nu h) D_{\ell-1} - \nu h \widehat{g}_{\ell-1}\right) + \lambda_x (x_\ell + u_\ell)\right] \\
&= \left[1 - (\kappa + \nu \lambda_D) h\right](g_{\ell-1} - \widehat{g}_{\ell-1}) - \lambda_D \varepsilon_{D,\ell} - \lambda_x (x_\ell + u_\ell) + \varepsilon_{g,\ell},
\end{aligned} \tag{18}$$

where the first step follows from Equations (2) and (11), and the second from Equation (1). Substituting into the Bellman equation, we find

$$\begin{aligned}
&B(g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C = \max_{x_\ell} \left\{x_\ell \left[\left[1 - (\kappa + \nu\lambda_D)h\right](g_{\ell-1} - \widehat{g}_{\ell-1}) - \lambda_x x_\ell\right]\right. \\
&\left. + e^{-rh}\left[B\left[\left[\left[1 - (\kappa + \nu\lambda_D)h\right](g_{\ell-1} - \widehat{g}_{\ell-1}) - \lambda_x x_\ell\right]^2 + \lambda_D^2 \sigma_D^2 h + \lambda_x^2 \sigma_u^2 h + \sigma_g^2 h\right] + C\right]\right\}.
\end{aligned} \tag{19}$$

The first-order condition yields

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

i.e., Equation (12), with

$$\beta = \frac{[1-(\kappa+\nu\lambda_D)h]\left(1-2e^{-rh}B\lambda_x\right)}{2\lambda_x\left(1-e^{-rh}B\lambda_x\right)}. \tag{20}$$

Substituting for x_ℓ in equation (19), we can determine B and C:

$$B = \frac{[1-(\kappa+\nu\lambda_D)h]^2}{4\lambda_x\left(1-e^{-rh}B\lambda_x\right)} \tag{21}$$

$$C = \frac{e^{-rh}B\left(\lambda_D^2\sigma_D^2+\lambda_x^2\sigma_u^2+\sigma_g^2\right)h}{1-e^{-rh}}. \tag{22}$$

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 (13), (14), (16) and (20)-(22). In Appendix C 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 (10)-(12).*

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.[6] 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 profitability by observing the dividend. Inference is characterized by the parameters λ_D and Σ_g^2, and these are given by Equations (13) and (16) with the insider's trading intensity β set to zero. To distinguish from 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. We also set $\rho \equiv \sqrt{\kappa^2+\frac{\nu^2\sigma_g^2}{\sigma_D^2}}$.

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

$$\lim_{h\to 0} \overline{\lambda}_D = \frac{\rho-\kappa}{\nu},$$

$$\lim_{h\to 0} \overline{\Sigma}_g^2 = \frac{\sigma_D^2(\rho-\kappa)}{\nu^2}.$$

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 profitability, characterized by $\overline{\Sigma}_g^2$, remains bounded away from zero. The intuition is that when h goes to zero the dividend follows an Ito process. Since profitability enters in the drift, it cannot be fully inferred within any finite time.

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*

$$\lim_{h\to 0} \frac{\lambda_D}{\sqrt{h}} = \frac{\sigma_g^2\nu}{\sigma_D^2\sqrt{r+2\kappa}} \tag{23}$$

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

$$\lim_{h\to 0} \frac{\Sigma_g^2}{\sqrt{h}} = \frac{\sigma_g^2}{\sqrt{r+2\kappa}} \tag{25}$$

$$\lim_{h\to 0} \frac{\beta}{\sqrt{h}} = \frac{\sigma_u\sqrt{r+2\kappa}}{\sigma_g} \tag{26}$$

$$\lim_{h\to 0} B = \frac{\sigma_u}{2\sigma_g} \tag{27}$$

$$\lim_{h\to 0} C = \frac{\sigma_g\sigma_u}{r}. \tag{28}$$

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 profitability is approximately $(\sigma_g^2/\sqrt{r+2\kappa})\sqrt{h}$ 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 profitability 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 consider the deviation between the price p_ℓ and the insider's valuation v_ℓ. This deviation is stochastic because of the new shocks $(\varepsilon_{D,\ell}, \varepsilon_{g,\ell}, u_\ell)$ subsequent to time zero. To isolate the effect of the time-zero shock, we set the subsequent shocks to their expected value of zero. Since our model is linear, this amounts to computing the expected deviation conditional on the insider's time-zero information $\mathcal{F}_0^i$. From Equations (8) and (10), $E\left(v_\ell - p_\ell | \mathcal{F}_0^i\right)$ is proportional to the market maker's expected estimation error $E\left(g_\ell - \widehat{g}_\ell | \mathcal{F}_0^i\right)$. To evaluate the latter when trading is almost continuous, we fix a calendar time t corresponding to period $\ell = t/h$, and consider the limit $e_t \equiv \lim_{h\to 0} E\left(g_{\frac{t}{h}} - \widehat{g}_{\frac{t}{h}} \middle| \mathcal{F}_0^i\right)$. Proposition 4 characterizes how e_t varies over time.

Proposition 4 *When the insider is not trading,*

$$e_t = e^{-\rho t}(\widehat{g}_0 - g_0).$$

When the insider is trading, $e_t = 0$ *for* $t > 0$.

Proposition 4 shows that in the absence of the insider, information about profitability is incorporated into prices slowly. For small h, the market maker's expected estimation error converges to zero at the finite rate ρ. By contrast, in the insider's presence, information about profitability is incorporated into prices very quickly. For small h, the expected estimation error reaches zero within any positive time t, and not only when t goes to infinity. Thus, the rate of convergence 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 estimation 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.[7]

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, (ii) the public revelation of her information through the dividend, and (iii) the obsolescence of her information through the mean-reversion in the firm's profitability.

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.[8]

That impatience induces 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, (ii) no information is revealed publicly through the dividend, i.e., the noise σ_D^2 in the dividend process is infinite, and (iii) profitability follows a random walk, i.e., the mean-reversion κ is zero. In Appendix D we extend our model to the case of no time-discounting, defining the insider's objective as the long-run average of per-period payoffs.[9] We then show that if $\sigma_D^2 = \infty$ and $\kappa = 0$, the insider prefers β to be as close to zero as possible. Thus, a patient insider prefers to trade slowly. Any of the three sources of impatience, however, suffices for the insider to trade quickly. For example, Proposition 3 shows that when $r > 0$, β is of order $\sqrt{h}$ regardless of the values of σ_D^2 and κ. Therefore, a positive interest rate can alone induce the insider to trade quickly. In Appendix D (where $r = 0$) we show that when $\kappa > 0$, β is of order $\sqrt{h}$ regardless of the value of σ_D^2, and when $\kappa = 0$ and $\sigma_D^2 < \infty$, β is of order $h^{\frac{2}{3}}$. In both cases β is larger than order h, implying that mean-reverting profitability or public revelation of information can alone induce the insider to trade quickly.

Although impatience induces the insider to trade quickly, it cannot be the full explanation. 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 possible, however, to introduce impatience in his model and show that the insider still trades slowly.[10]

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 mean-reverting profitability, or a combination of the three).

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

$$\left[B(g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C\right] A_1 \equiv B'(g_{\ell-1} - \widehat{g}_{\ell-1})^2 + C', \tag{29}$$

i.e., the value function times the scaling factor A_1 that was dropped for simplicity when writing the insider's objective as (17) instead of (9). 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 $\sigma_g \sigma_u / r$. The parameter C' also converges to a positive limit since A_1 converges to $\frac{\nu}{(r+\nu)(r+\kappa)}$. Thus, the insider's profits remain positive despite the market's 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 $x_\ell(g_{\ell-1} - \widehat{g}_{\ell-1})A_1$. 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 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.[11]

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 and Empirical Implications

The results of the previous section are asymptotic, i.e., they 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, the standard deviation σ_D of dividends, σ_g of profitability, and σ_u of noise trading, and the mean-reversion ν of dividends and κ of profitability. Before selecting values for these parameters, we define and compute our measure of market efficiency. It turns out that this measure is independent of several parameters, simplifying the calibration.

We measure efficiency by the speed at which private information is incorporated into prices. As in Proposition 4, we assume that at time zero the insider learns that profitability g_0 differs from the market maker's expectation $\widehat{g}_0$. We then examine how quickly the deviation between the insider's valuation and the price converges to zero, in expectation conditional on the insider's time-zero information. From the proof of Proposition 4, the convergence dynamics for given h are

$$E\left(v_{\frac{t}{h}} - p_{\frac{t}{h}} \middle| \mathcal{F}_0^i\right) = \left[1 - (\kappa + \nu\lambda_D)h - \lambda_x\beta\right]^{\frac{t}{h}} (v_0 - p_0). \tag{30}$$

Our measure of market efficiency is the time t_χ by which the deviation drops to a percentage χ of its original value. This time is defined by

$$E\left(v_{\frac{t}{h}} - p_{\frac{t}{h}} \middle| \mathcal{F}_0^i\right) = (1 - \chi)(v_0 - p_0). \tag{31}$$

Comparing Equations (30) and (31), we find

$$t_\chi = \frac{h \log(1 - \chi)}{\log\left[1 - (\kappa + \nu\lambda_D)h - \lambda_x\beta\right]}. \tag{32}$$

We next observe that t_χ is independent of σ_u. Indeed, consider the solution $(\lambda_D, \lambda_x, \Sigma_g, \beta, B, C)$ to Equations (13), (14), (16) and (20)-(22). 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 (32) 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, t_χ depends on $(\sigma_D, \sigma_g, \nu)$ only through the ratio $\nu\sigma_g/\sigma_D$. Indeed, suppose that $(\sigma_D, \sigma_g, \nu)$ is replaced by $(yz\sigma_D, y\sigma_g, z\nu)$ for $y, z > 0$. Then, the new solution to Equations (13), (14), (16) and (20)-(22) is $(\lambda_D/z, y\lambda_x, y\Sigma_g, \beta/y, B/y, yC)$. Equation (32) then implies that t_χ stays constant, meaning that t_χ depends on $(\sigma_D, \sigma_g, \nu)$ only through $\nu\sigma_g/\sigma_D$. Thus, we can normalize $(\sigma_u, \sigma_g, \nu)$ to one without loss of generality, and are left with the task of selecting values for (r, h, σ_D, κ).[12]

We set the interest rate r to 2%. Larger values of r would reinforce our results because the insider would trade faster. To calibrate the time between consecutive periods h, we use the average time between transactions in the stock market. Since this time varies greatly across stocks, we pick two specific stocks, one large-cap and one small-cap. In the context of these stocks, we also select values for the remaining parameters (σ_D, κ).

Large-Cap Stock

Our large-cap stock is Coca-Cola, ticker symbol KO, one of the top twenty US companies in the NYSE composite index (as of October 2005). To determine the time between transactions, we use the TAQ database and consider the four weeks 6-31/10/2003. The average number of transactions per day is 3963. With 252 trading days per year, this translates to one transaction every $h = 1/(252 \times 3963) = 0.000001$ year.

One approach to calibrating (σ_D, κ) is through the stochastic process followed by actual dividends or earnings. We focus on earnings, given that they might be better described by the stochastic process (3)-(4).[13] We consider annual changes in earnings per share (EPS) and use the method of moments, matching moments derived from the stochastic process (3)-(4) to their actual counterparts.[14] The moments we select are the standard deviation of annual earnings changes and the correlation of annual changes with each of the first five lags.[15] To increase estimation precision, we compute moments for each of the first twenty US companies in the NYSE composite index, and use cross-sectional averages. Moreover, for each company we consider the entire EPS history available from COMPUSTAT. The estimation results are $\sigma_D = 1.06$, $\sigma_g = 0.62$, $\nu = 1.47$, $\kappa = 0$. Thus, when (σ_g, ν) are normalized to one, σ_D is set equal to $1.06/(0.62 \times 1.47) = 1.17$.

To interpret the magnitudes of σ_D and κ, we map them into the speed at which the market maker learns about shocks to profitability in the insider's absence. We examine, in particular, how quickly a time-zero estimation error converges to zero, in expectation over the subsequent shocks. Convergence is faster the smaller σ_D is because the dividend process contains less noise and is thus more informative about profitability. Convergence is also faster the larger κ is because profitability decays faster to its long-run mean. We characterize convergence by the time $\bar{t}_\chi$ by which the market maker's estimation error drops to a percentage χ of its original value. From the proof of Proposition 4, this is

$$\bar{t}_\chi = \frac{h \log(1-\chi)}{\log(1-\nu\bar{\lambda}_D h)}. \tag{33}$$

We focus on the time $\bar{t}_{0.5}$ that corresponds to the half-life of the convergence dynamics. For $(\sigma_D, \sigma_g, \nu, \kappa) = (1.17, 1, 1, 0)$, learning in the insider's absence has a half-life of $\bar{t}_{0.5} = 0.81$ year.[16] While such a half-life seems reasonable, we show in our tables results for half-lives ranging from fifteen days to two years. This is both because of the noise in estimating moments, and because the actual half-life could be shorter than implied by our model. Indeed, investors can receive more information about a company than just earnings (e.g., from analysts' forecasts). In generating the values of $\bar{t}_{0.5}$, we vary σ_D but keep κ to zero for simplicity.

Table 1 compares the half-life $\bar{t}_{0.5}$ in the insider's absence to the half-life $t_{0.5}$ in her presence. The table shows that the insider speeds information revelation dramatically. Suppose, for example, that the half-life in the insider's absence is one year. The insider reduces this to 1.57 days, a factor of 365/1.57=232. We should emphasize that the insider can choose to reveal her information slowly, stretching the half-life closer to one year. Our main result, however, is that she prefers to reveal it quickly. Indeed, a half-life of 1.57 days ensures a minimal price impact, while allowing the insider to reap the benefits associated with impatience.

INSERT TABLE 1 SOMEWHERE HERE

The results of Table 1 stand in sharp contrast to Kyle (1985). Indeed, in Kyle the insider reveals her information at a constant rate over time. Thus, the half-life of information revelation in the insider's presence is one-half of the time it would take for the information to be announced publicly. For example, if the information is to be announced in one year, the insider will take six months to incorporate half of it into the price.

Small-Cap Stock

Our small-cap stock is Bairnco, ticker symbol BZ, which is traded in the NYSE but is not part of the NYSE composite index. The average number of transactions per day over the four weeks 6-31/10/2003 is four. With 252 trading days per year, this translates to one transaction every $h = 1/(252 \times 4) = 0.000992$ year. We set $\kappa = 0$, and choose σ_D to obtain the same values of the half-life $\bar{t}_{0.5}$ as for the large-cap stock. The results are in Table 2.

INSERT TABLE 2 SOMEWHERE HERE

Table 2 shows that the insider speeds information revelation substantially, even for very infrequently traded stocks. For example, when the half-life in the insider's absence is one year, the insider reduces this to 23.98 days, a factor of 365/23.98=15. Of course, the lower trading frequency has an important effect since for the large-cap stock the factor is 232. Nevertheless, the factor is still much larger than Kyle's factor of two.

Empirical Implications

To translate Tables 1 and 2 into empirically measurable quantities, we consider the correlation between the change in insider holdings over two consecutive intervals $[t-s, t]$ and $[t, t+s]$. Since the change in holdings over an interval equals the net amount of shares bought during that interval, the correlation is

$$\text{Corr}\left(\sum_{i=1}^{\frac{s}{h}} x_{\frac{t-s}{h}+i}, \sum_{i=1}^{\frac{s}{h}} x_{\frac{t}{h}+i}\right).$$

This quantity is measurable: for example if the insider is a mutual fund, then s is a quarter, the frequency at which mutual funds disclose their holdings. Using our model, we can compute the quarterly correlation, and we report it in Table 3 for the selected large-cap and small-cap stocks:[17]

INSERT TABLE 3 SOMEWHERE HERE

Table 3 shows that the correlation is closely related to the half-life of information revelation. Consider, for example, the large-cap stock, where the insider reveals her information with a half-life not exceeding two days. Since each piece of information is exploited over a very short interval, the correlation in insider trades over consecutive quarters is very small. For the small-cap stock, the correlation can be more substantial, reaching 0.361 when the insider reveals her information with a half-life of 34.71 days.

From Table 3, an implication of our model is that changes in mutual-fund holdings over consecutive quarters should be uncorrelated for large stocks, but might exhibit positive correlation for small stocks. Our model also implies that changes in fund holdings of large stocks over one quarter should be uncorrelated with stock returns over the next quarters. This is because the information on which the insider trades is incorporated very quickly into the price. For small stocks, the correlation might be positive and larger.[18]

6 Conclusion

In this paper we consider a discrete-time, infinite-horizon model with a monopolistic insider. The insider observes privately the expected growth rate of asset dividends in each period, and can trade with competitive market makers in the presence of noise traders. We find 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 an infinite horizon, we adopt the information structure in Wang (1993).

The insider in our model can be best interpreted as a proprietary-trading desk, hedge fund, or mutual fund, generating a continuous flow of private information through their superior research. Given that there are many 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 monopolists do. 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, and one could possibly use it to study the continuous-time limit.

A Price

Defining the parameter n by $e^{-nh} \equiv 1 - \nu h$, we can write Equation (1) as

$$D_\ell = e^{-nh} D_{\ell-1} + (1 - e^{-nh}) g_{\ell-1} + \varepsilon_{D,\ell}.$$

Therefore,

$$D_{\ell'} = e^{-n(\ell'-\ell)h} D_\ell + (1 - e^{-nh}) \sum_{z=\ell}^{\ell'-1} e^{-n(\ell'-1-z)h} g_z + \sum_{z=\ell+1}^{\ell'} e^{-n(\ell'-z)h} \varepsilon_{D,z}. \tag{A.1}$$

Likewise, defining the parameter k by $e^{-kh} \equiv 1 - \kappa h$, we can write Equation (2) as

$$g_\ell = e^{-kh} g_{\ell-1} + (1 - e^{-kh}) \overline{g} + \varepsilon_{g,\ell}.$$

Therefore,

$$g_{\ell'} = e^{-k(\ell'-\ell)h} g_\ell + (1 - e^{-k(\ell'-\ell)h}) \overline{g} + \sum_{z=\ell+1}^{\ell'} e^{-k(\ell'-z)h} \varepsilon_{g,z}. \tag{A.2}$$

Taking expectations in Equation (A.1), we find

$$E(D_{\ell'} | \mathcal{F}_\ell^m) = e^{-n(\ell'-\ell)h} D_\ell + (1 - e^{-nh}) \sum_{z=\ell}^{\ell'-1} e^{-n(\ell'-1-z)h} E(g_z | \mathcal{F}_\ell^m).$$

Combining with Equation (A.2), we find

$$\begin{aligned} E(D_{\ell'} | \mathcal{F}_\ell^m) &= e^{-n(\ell'-\ell)h} D_\ell + (1 - e^{-nh}) \sum_{z=\ell}^{\ell'-1} e^{-n(\ell'-1-z)h} E(e^{-k(z-\ell)h}(g_\ell - \overline{g}) + \overline{g} | \mathcal{F}_\ell^m) \\ &= e^{-n(\ell'-\ell)h} D_\ell + (1 - e^{-nh}) \frac{e^{-n(\ell'-\ell)h} - e^{-k(\ell'-\ell)h}}{e^{-nh} - e^{-kh}} \left[E(g_\ell | \mathcal{F}_\ell^m) - \overline{g} \right] + \left[1 - e^{-n(\ell'-\ell)h} \right] \overline{g}. \end{aligned}$$

Substituting into Equation (6), the price is equal to

$$\begin{aligned} p_\ell &= \sum_{\ell'=\ell}^{\infty} E(D_{\ell'} | \mathcal{F}_\ell^m) h e^{-r(\ell'-\ell)h} \\ &= \frac{h}{1 - e^{-(r+n)h}} \left[D_\ell + \frac{(1 - e^{-nh}) e^{-rh}}{1 - e^{-(r+k)h}} \left[E(g_\ell | \mathcal{F}_\ell^m) - \overline{g} \right] + \frac{(1 - e^{-nh}) e^{-rh}}{1 - e^{-rh}} \overline{g} \right]. \end{aligned}$$

This is as in Equation (7) with

$$A_0 \equiv \frac{h}{1-e^{-(r+n)h}},$$

$$A_1 \equiv A_0 \frac{(1-e^{-nh})e^{-rh}}{1-e^{-(r+k)h}},$$

$$A_2 \equiv A_0 \frac{(1-e^{-nh})(1-e^{-kh})e^{-2rh}}{(1-e^{-rh})(1-e^{-(r+k)h})}.$$

B 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 drop the conditioning set $\mathcal{F}^m_{\ell-1}$ for now on, except for moments that are conditional on other information. The signals observed by the market maker in period ℓ are

$$D_\ell = (1-\nu h)D_{\ell-1} + \nu h g_{\ell-1} + \varepsilon_{D,\ell}$$

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 belief about $g_{\ell-1}$ is of the form

$$\begin{aligned} g_{\ell-1} &= E(g_{\ell-1}) + \frac{\lambda_D}{1-\kappa h}\left(D_\ell - E\left(D_\ell\right)\right) + \frac{\lambda_x}{1-\kappa h}\left(x_\ell + u_\ell - E\left(x_\ell + u_\ell\right)\right) + \eta_\ell \\ &= \widehat{g}_{\ell-1} + \frac{\lambda_D}{1-\kappa h}\left(D_\ell - (1-\nu h)D_{\ell-1} - \nu h \widehat{g}_{\ell-1}\right) + \frac{\lambda_x}{1-\kappa h}\left(x_\ell + u_\ell\right) + \eta_\ell, \end{aligned} \quad \text{(B.1)}$$

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 belief about

$$g_\ell = (1-\kappa h)g_{\ell-1} + \kappa h \overline{g} + \varepsilon_{g,\ell}$$

thus is

$$g_\ell = (1-\kappa h)\widehat{g}_{\ell-1} + \kappa h \overline{g} + \lambda_D\left(D_\ell - (1-\nu h)D_{\ell-1} - \nu h \widehat{g}_{\ell-1}\right) + \lambda_x\left(x_\ell + u_\ell\right) + (1-\kappa h)\eta_\ell + \varepsilon_{g,\ell}.$$

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

$$\mathrm{Cov}\left(D_\ell, g_{\ell-1}\right) = \frac{1}{1-\kappa h}\mathrm{Cov}\left(D_\ell, \lambda_D D_\ell + \lambda_x\left(x_\ell + u_\ell\right)\right), \tag{B.2}$$

$$\mathrm{Cov}\left(x_\ell + u_\ell, g_{\ell-1}\right) = \frac{1}{1-\kappa h}\mathrm{Cov}\left(x_\ell + u_\ell, \lambda_D D_\ell + \lambda_x\left(x_\ell + u_\ell\right)\right). \tag{B.3}$$

Since

$$\mathrm{Cov}\left(D_\ell, g_{\ell-1}\right) = \mathrm{Cov}\left((1-\nu h)D_{\ell-1} + \nu h g_{\ell-1} + \varepsilon_{D,\ell}, g_{\ell-1}\right) = \mathrm{Var}\left(g_{\ell-1}\right)\nu h = \Sigma_g^2 \nu h, \tag{B.4}$$

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

$$\begin{aligned}
\mathrm{Var}\left(D_\ell\right) &= \mathrm{Var}\left((1-\nu h)D_{\ell-1} + \nu h g_{\ell-1} + \varepsilon_{D,\ell}\right) \\
&= \mathrm{Var}\left(g_{\ell-1}\right)\nu^2 h^2 + \mathrm{Var}\left(\varepsilon_{D,\ell}\right) = \Sigma_g^2\nu^2 h^2 + \sigma_D^2 h, \\
\mathrm{Cov}\left(D_\ell, x_\ell + u_\ell\right) &= \mathrm{Cov}\left((1-\nu h)D_{\ell-1} + \nu h g_{\ell-1} + \varepsilon_{D,\ell}, \beta(g_{\ell-1} - \widehat{g}_{\ell-1}) + u_\ell\right) \\
&= \beta\mathrm{Var}\left(g_{\ell-1}\right)\nu h = \beta\Sigma_g^2\nu h, \\
\mathrm{Var}\left(x_\ell + u_\ell\right) &= \mathrm{Var}\left(\beta(g_{\ell-1} - \widehat{g}_{\ell-1}) + u_\ell\right) = \beta^2\mathrm{Var}\left(g_{\ell-1}\right) + \mathrm{Var}\left(u_\ell\right) = \beta^2\Sigma_g^2 + \sigma_u^2 h,
\end{aligned}$$

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

$$\begin{aligned}
\Sigma_g^2\nu h &= \frac{1}{1-\kappa h}\left[\lambda_D(\Sigma_g^2\nu^2 h^2 + \sigma_D^2 h) + \lambda_x\beta\Sigma_g^2\nu h\right], \\
\beta\Sigma_g^2 &= \frac{1}{1-\kappa h}\left[\lambda_D\beta\Sigma_g^2\nu h + \lambda_x(\beta^2\Sigma_g^2 + \sigma_u^2 h)\right].
\end{aligned}$$

The solution to this linear system is given by Equations (13) and (14). Therefore, the posterior expectation of g_ℓ is as in Equation (11).

The posterior variance of g_ℓ is

$$\mathrm{Var}(g_\ell | F_\ell^m) = \mathrm{Var}((1-\kappa h)\eta_\ell + \varepsilon_{g,\ell}) = (1-\kappa h)^2\mathrm{Var}(\eta_\ell) + \sigma_g^2 h. \tag{B.6}$$

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

$$\begin{aligned}
\mathrm{Var}(\eta_\ell) &= \mathrm{Var}(g_{\ell-1}) - \frac{1}{(1-\kappa h)^2}\mathrm{Var}\left(\lambda_D D_\ell + \lambda_x (x_\ell + u_\ell)\right) \\
&= \mathrm{Var}(g_{\ell-1}) - \frac{\lambda_D}{(1-\kappa h)^2}\mathrm{Cov}\left(D_\ell, \lambda_D D_\ell + \lambda_x (x_\ell + u_\ell)\right) \\
&\quad - \frac{\lambda_x}{(1-\kappa h)^2}\mathrm{Cov}\left(x_\ell + u_\ell, \lambda_D D_\ell + \lambda_x (x_\ell + u_\ell)\right) \\
&= \mathrm{Var}(g_{\ell-1}) - \frac{\lambda_D}{1-\kappa h}\mathrm{Cov}\left(D_\ell, g_{\ell-1}\right) - \frac{\lambda_x}{1-\kappa h}\mathrm{Cov}\left(x_\ell + u_\ell, g_{\ell-1}\right) \\
&= \Sigma_g^2 - \frac{\lambda_D}{1-\kappa h}\Sigma_g^2 \nu h - \frac{\lambda_x}{1-\kappa h}\beta \Sigma_g^2,
\end{aligned}$$

where the third step follows from (B.2) and (B.3), and the fourth from (B.4) and (B.5). Using (13) and (14) to substitute for λ_D and λ_x, and plugging $\mathrm{Var}(\eta_\ell)$ back into (B.6), we find (15).

C Proof of Propositions 1-4

Proof of Proposition 1: We will show that the system of Equations (13), (14), (16) and (20)-(22) 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 (21) can be written as

$$B^2 - \frac{e^{rh}}{\lambda_x}B + \frac{e^{rh}\left[1-(\kappa+\nu\lambda_D)h\right]^2}{4\lambda_x^2} = 0.$$

This quadratic equation in B has the two solutions

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

The solution with the plus sign cannot be part of a solution to the overall system. Indeed, Equation (14) implies that $\lambda_x\beta > 0$, which from Equations (13) and (20) 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

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

Plugging into Equation (20), we find

$$\frac{\lambda_x \beta}{1-(\kappa+\nu\lambda_D)h} = \frac{\sqrt{1-e^{-rh}\left[1-(\kappa+\nu\lambda_D)h\right]^2}}{1+\sqrt{1-e^{-rh}\left[1-(\kappa+\nu\lambda_D)h\right]^2}}. \tag{C.2}$$

Substituting for λ_D and λ_x using Equations (13) and (14), we can write Equation (C.2) as

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

We can reduce Equation (C.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 (16), 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 (13), (14), (16) and (20)-(22) has a solution, we note that the left-hand side (LHS) of Equation (C.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 (C.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 (16), (13), (14), (C.1) and (22), respectively. The insider's second-order condition is the requirement that the problem (19) be concave, i.e., $1-e^{-rh}B\lambda_x > 0$. This inequality is satisfied because of Equation (C.1).

To show that the solution is unique, we will show that the LHS of Equation (C.3) is increasing in β, while the RHS is decreasing. The LHS is increasing in β if $\Sigma_g^2\beta^2$ is increasing. Equation (16) implies that

$$\Sigma_g^2\beta^2\sigma_D^2 = \frac{(1-\kappa h)^2\Sigma_g^2\sigma_D^2\sigma_u^2 h}{\Sigma_g^2-\sigma_g^2 h} - \Sigma_g^2\nu^2\sigma_u^2h^2 - \sigma_D^2\sigma_u^2 h. \tag{C.4}$$

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

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

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

$$Z = 1 - \frac{\nu^2 h(\Sigma_g^2 - \sigma_g^2 h)}{(1-\kappa h)^2 \sigma_D^2}.$$

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

Proof of Proposition 2: When $\beta = 0$, Equations (13) and (16) become

$$\overline{\lambda}_D = \frac{(1-\kappa h)\overline{\Sigma}_g^2 \nu}{\overline{\Sigma}_g^2 \nu^2 h + \sigma_D^2} \tag{C.5}$$

and

$$\left(\overline{\Sigma}_g^2 - \sigma_g^2 h\right)\left(\overline{\Sigma}_g^2 \nu^2 h + \sigma_D^2\right) - (1-\kappa h)^2 \overline{\Sigma}_g^2 \sigma_D^2 = 0, \tag{C.6}$$

respectively. To study the system of (C.5) and (C.6) for small h, we divide Equation (C.6) by h:

$$(\overline{\Sigma}_g^2 - \sigma_g^2 h)\overline{\Sigma}_g^2 \nu^2 + \frac{1-(1-\kappa h)^2}{h}\overline{\Sigma}_g^2 \sigma_D^2 - \sigma_g^2 \sigma_D^2 = 0. \tag{C.7}$$

For $h = 0$, Equation (C.7) becomes

$$\overline{\Sigma}_g^4 \nu^2 + 2\kappa \overline{\Sigma}_g^2 \sigma_D^2 - \sigma_g^2 \sigma_D^2 = 0,$$

and has the unique positive solution

$$\overline{\Sigma}_g^2 = \frac{\sigma_D^2(\rho - \kappa)}{\nu^2},$$

where $\rho \equiv \sqrt{\kappa^2 + \frac{\nu^2 \sigma_g^2}{\sigma_D^2}}$. Substituting into Equation (C.5), we find

$$\overline{\lambda}_D = \frac{\overline{\Sigma}_g^2 \nu}{\sigma_D^2} = \frac{\rho - \kappa}{\nu}.$$

By continuity, this is also the limit of the solution when h goes to zero.[19] ∎

Proof of Proposition 3: We first solve the system of Equations (16) and (C.3) in the unknowns Σ_g^2 and β. To study this system for small h, we set $\Sigma_g^2 \equiv S_g^2 \sqrt{h}$ and $\beta \equiv b\sqrt{h}$, and divide Equation (16) by h^2. This results in the system

$$\left(S_g^2 - \sigma_g^2 \sqrt{h}\right) S_g^2 (b^2 \sigma_D^2 + \nu^2 \sigma_u^2 h) + \frac{1-(1-\kappa h)^2}{\sqrt{h}} S_g^2 \sigma_D^2 \sigma_u^2 - \sigma_g^2 \sigma_D^2 \sigma_u^2 = 0$$

and

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

For $h = 0$, the system becomes

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

and

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

and has the solution $S_g^2 = \sigma_g^2 / \sqrt{r + 2\kappa}$ and $b = \sigma_u \sqrt{r + 2\kappa} / \sigma_g$. By continuity, this is also the limit of the solution when h goes to zero. This establishes the limits (25) and (26) since

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

and

$$\lim_{h \to 0} \frac{\beta}{\sqrt{h}} = \lim_{h \to 0} b = \frac{\sigma_u \sqrt{r + 2\kappa}}{\sigma_g}.$$

To prove the limits (23), (24), (27), and (28), we write Equations (13), (14), (22), and (C.1) in terms of S_g^2 and b, and use the limits of S_g^2 and b. ∎

Proof of Proposition 4: Plugging x_ℓ from Equation (12) into (18), we find

$$g_\ell - \widehat{g}_\ell = [1 - (\kappa + \nu\lambda_D)h - \lambda_x \beta]\,(g_{\ell-1} - \widehat{g}_{\ell-1}) - \lambda_D \varepsilon_{D,\ell} - \lambda_x u_\ell + \varepsilon_{g,\ell}.$$

We next take expectations conditional on the insider's time-zero information:

$$E\left(g_\ell - \widehat{g}_\ell \middle| \mathcal{F}_0^i\right) = [1 - (\kappa + \nu\lambda_D)h - \lambda_x \beta]\, E\left(g_{\ell-1} - \widehat{g}_{\ell-1} \middle| \mathcal{F}_0^i\right),$$

and iterate from period 0 to ℓ:

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

To prove the proposition, we need to determine the limit of $[1 - (\kappa + \nu\lambda_D)h - \lambda_x \beta]^{\frac{t}{h}}$ when h goes to zero. When the insider is not trading, the limit is $e^{-\rho t}$ because λ_D converges to $(\rho - \kappa)/\nu$ 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. ∎

D No Time-Discounting

For $r = 0$, we define the insider's objective as the long-run average of the normalized per-period payoffs $x_\ell(g_\ell - \widehat{g}_\ell)$. 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 (12) for some constant β. Then, λ_D, λ_x and Σ_g are given as a function of β by Equations (13), (14) and (16), 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 (18), and $x_\ell = \beta(g_{\ell-1} - \widehat{g}_{\ell-1})$. Substituting for $g_\ell - \widehat{g}_\ell$ and x_ℓ, we find

$$B_L = \frac{\beta\left[1 - (\kappa + \nu\lambda_D)h - \lambda_x\beta\right] + (L-1)B_{L-1}\left[1 - (\kappa + \nu\lambda_D)h - \lambda_x\beta\right]^2}{L}$$

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

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

$$B_L = \frac{\beta\sum_{k=0}^{L-1}\left[1 - (\kappa + \nu\lambda_D)h - \lambda_x\beta\right]^{2k+1}}{L} \tag{D.1}$$

$$C_L = \frac{\beta\left(\lambda_D^2\sigma_D^2 + \lambda_x^2\sigma_u^2 + \sigma_g^2\right)h\sum_{k=0}^{L-2}(L-1-k)\left[1 - (\kappa + \nu\lambda_D)h - \lambda_x\beta\right]^{2k+1}}{L}. \tag{D.2}$$

Suppose that $\sigma_D = \infty$ and $\kappa = 0$. Equations (13) and (14) imply that $\lambda_D = 0$ and

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

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

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

When $\beta = 0$, we have $\Pi = 0$ since the insider is not trading. 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.

Suppose next that $\kappa > 0$ or $\sigma_D^2 < \infty$. Equations (13) and (14) imply that

$$1 - (\kappa + \nu \lambda_D) h - \lambda_x \beta = \frac{(1 - \kappa h) \sigma_D^2 \sigma_u^2 h}{\Sigma_g^2 \left(\beta^2 \sigma_D^2 + \nu^2 \sigma_u^2 h^2\right) + \sigma_D^2 \sigma_u^2 h}.$$

Since $\kappa > 0$ or $\sigma_D^2 < \infty$, we have $0 < 1 - (\kappa + \nu \lambda_D) h - \lambda_x \beta < 1$. Equations (D.1) and (D.2) then imply that $\lim_{L\to\infty} B_L = 0$ and

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

Therefore, maximizing Π is equivalent to maximizing

$$\frac{\beta \left[1 - (\kappa + \nu \lambda_D) h - \lambda_x \beta\right]}{1 - \left[1 - (\kappa + \nu \lambda_D) h - \lambda_x \beta\right]^2}.$$

The first-order condition is

$$\begin{aligned} & 1 - (\kappa + \nu \lambda_D) h - 2\lambda_x \beta - \left[1 - (\kappa + \nu \lambda_D) h\right] \left[1 - (\kappa + \nu \lambda_D) h - \lambda_x \beta\right]^2 = 0 \\ \Leftrightarrow\ & (1 - \kappa h)(\lambda_x \beta)^2 = \nu \lambda_D h \left[\left[1 - (\kappa + \nu \lambda_D) h - \lambda_x \beta\right]^2 + (1 - \kappa h)\left[2(1 - \kappa h) - \nu \lambda_D h - 2\lambda_x \beta\right] - 1\right] \\ & \qquad + \left[1 - (1 - \kappa h)^2\right] (1 - \kappa h - 2\lambda_x \beta). \end{aligned} \tag{D.3}$$

In equilibrium, λ_D, λ_x, Σ_g and β are the solution to the system of Equations (13), (14), (16) and (D.3). To determine the solution for small h, we distinguish between the cases $\kappa > 0$ and $\kappa = 0$. When $\kappa > 0$, we set $\lambda_D = l_d\sqrt{h}$, $\Sigma_g^2 \equiv S_g^2\sqrt{h}$ and $\beta \equiv b\sqrt{h}$. We can then check that the resulting system in l_d, λ_x, S_g and b has the solution

$$b = \frac{\sigma_u\sqrt{2\kappa}}{\sigma_g}$$

$$S_g^2 = \frac{\sigma_g^2}{\sqrt{2\kappa}}$$

$$\lambda_x = \frac{\sigma_g}{\sigma_u}$$

$$l_d = \frac{\sigma_g^2\nu}{\sigma_D^2\sqrt{2\kappa}}$$

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

$$\lim_{h\to 0}\frac{\beta}{\sqrt{h}} = \lim_{h\to 0} b = \frac{\sigma_u\sqrt{2\kappa}}{\sigma_g}.$$

Therefore, β is of order $\sqrt{h}$. When $\kappa = 0$ and $\sigma_D^2 < \infty$, we set $\lambda_D = l_d h^{\frac{1}{3}}$, $\Sigma_g^2 \equiv S_g^2 h^{\frac{1}{3}}$ and $\beta \equiv b h^{\frac{2}{3}}$. We can then check that the resulting system in l_d, λ_x, S_g and b has the solution

$$b = \frac{2^{\frac{1}{3}}\nu^{\frac{2}{3}}\sigma_u}{\sigma_D^{\frac{2}{3}}\sigma_g^{\frac{1}{3}}}$$

$$S_g^2 = \frac{\sigma_D^{\frac{2}{3}}\sigma_g^{\frac{4}{3}}}{2^{\frac{1}{3}}\nu^{\frac{2}{3}}}$$

$$\lambda_x = \frac{\sigma_g}{\sigma_u}$$

$$l_d = \frac{\nu^{\frac{1}{3}}\sigma_g^{\frac{4}{3}}}{2^{\frac{1}{3}}\sigma_D^{\frac{4}{3}}}$$

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

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

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

References

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

Back, K. and S. Baruch, 2004, "Information in Securities Markets: Kyle meets Glosten and Milgrom," *Econometrica,* 72, 433–465.

Back, K., H. Cao and G. Willard, 2000, "Imperfect Competition among Informed Traders," *Journal of Finance,* 55, 2117–2155.

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

Brennan, M. and Y. Xia, 2001, "Stock Price Volatility and Equity Premium," *Journal of Monetary Economics,* 47, 249–283.

Chau, M., 2003, "Market-Making when Traders Follow Dynamic Order-Placement Strategies," working paper, ESSEC.

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

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

Fudenberg, D. and J. Tirole, 1991, *Game Theory,* MIT Press, Cambridge MA.

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

Kothari, S.P., 2001, "Capital Markets Research in Accounting," *Journal of Accounting and Economics,* 31, 105–231.

Kyle, A., 1985, "Continuous Auctions and Insider Trading," *Econometrica,* 53, 1315–1335.

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

Spiegel, M. and A. Subrahmanyam, 1995, "On Intraday Risk Premia," *Journal of Finance,* 50, 319–339.

Taub, B., D. Bernardt and P. Seiler, 2005, "Speculative Dynamics," working paper, University of Illinois.

Vayanos, D., 1999, "Strategic Trading and Welfare in a Dynamic Market," *Review of Economic Studies,* 66, 219–254.

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

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

Wang, J., 1994, "A Model of Competitive Stock Trading Volume," *Journal of Political Economy,* 102, 127–168.

Notes

[1]Kyle assumes no impatience, but it is possible to introduce impatience in his model and show that the insider still trades slowly. For example, Back and Baruch (2004) assume that the insider's information is announced publicly at a Poisson rather than a deterministic time. The insider then becomes impatient because she can lose her informational advantage at any time, but is shown to still trade slowly.

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

[3]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.

[4]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.

[5]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 profitability. 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.

[6]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.

[7]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).

[8]See, however, Vayanos (1999,2001) 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.

[9]This type of objective is standard in the literature on repeated games with no discounting. See, for example, Fudenberg and Tirole (1991). In addition to assuming the long-run average, we re-normalize per-period payoffs to $x_\ell(g_\ell - \widehat{g}_\ell)$ rather than $x_\ell(v_\ell - p_\ell)$.

[10]See, for example, Back and Baruch (2004) who assume that the insider's information is announced publicly at a Poisson rather than a deterministic time.

[11]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.

Note that while the number of shares traded by the insider goes to infinity, the insider generates a negligible fraction of total trading volume. Indeed, the volume u_ℓ generated by noise traders in period ℓ is of order $\sqrt{h}$. Since x_ℓ is of order $h^{\frac{3}{4}}$, the ratio x_ℓ/u_ℓ converges to zero when h goes to zero.

[12]The parameters (σ_u, σ_g) can be calibrated through the aggregate trading volume and the bid-ask spread. See Chau (2003) for an example of such a calibration.

[13]In particular, there is substantial evidence that earnings exhibit mean-reversion. See, for example, the survey by Kothari (2001).

[14]Brennan and Xia (2001) follow a similar approach when calibrating a dividend process with unobservable time-varying drift. In calculating the model-implied moments, we assume for consistency that annual earnings are not the realization of D_ℓ at year-end, but a capitalized sum of all values of D_ℓ over the year. The calculations are available upon request.

[15]We consider more moments than parameters to increase estimation accuracy. We select the

parameters that minimize the sum of squared deviations between the model-implied and the actual moments.

[16]As with t_χ, the time $\bar{t}_\chi$ depends on $(\sigma_D, \sigma_g, \nu)$ only through the ratio $\nu\sigma_g/\sigma_D$. It is thus the same whether $(\sigma_D, \sigma_g, \nu, \kappa)$ equals (1.17,1,1,0) or (1.06,0.62,1.47,0).

[17]The correlation is given by

$$\frac{a\left(1-a^{\frac{s}{h}}\right)^2}{\frac{s}{h}(1-a)^2+2a\left(\frac{s}{h}-1+a^{\frac{s}{h}}-\frac{s}{h}a\right)},$$

where $a \equiv 1-(\kappa+\nu\lambda_D)h-\lambda_x\beta$ and s/h is an integer. The calculations are available upon request.

[18]An alternative hypothesis is that funds trade for non-informational reasons (e.g., liquidity, hedging), and do so more slowly for small stocks because of the larger price impact. For example, Vayanos (2001) shows that a strategic hedger trades more slowly if he faces more risk-averse market makers. Under this hypothesis, however, the correlation between changes in holdings and subsequent returns should eventually turn negative, as fund trading would impact the stock price but not the fundamental value.

[19]Continuity follows from the implicit function theorem. We divide Equation (C.6) by h so that the resulting system is well-behaved for $h=0$.

$\bar{t}_{0.5}$	$t_{0.5}$	$t_{0.5}/\bar{t}_{0.5}$
15 days	0.30 days	2.02%
60 days	0.72 days	1.21%
0.5 year	1.27 days	0.70%
1 year	1.57 days	0.43%
2 years	1.72 days	0.24%

Table 1: Speed of information revelation for a large-cap stock.

$\bar{t}_{0.5}$	$t_{0.5}$	$t_{0.5}/\bar{t}_{0.5}$
15 days	3.08 days	20.51%
60 days	7.67 days	12.78%
0.5 year	15.70 days	8.63%
1 year	23.98 days	6.57%
2 years	34.71 days	4.76%

Table 2: Speed of information revelation for a small-cap stock.

$\bar{t}_{0.5}$	Large-cap		Small-cap	
	$t_{0.5}$	Correlation	$t_{0.5}$	Correlation
15 days	0.30 days	0.002	3.08 days	0.026
60 days	0.72 days	0.006	7.67 days	0.070
0.5 year	1.27 days	0.010	15.70 days	0.161
1 year	1.57 days	0.013	23.98 days	0.256
2 years	1.72 days	0.014	34.71 days	0.361

Table 3: Correlation between changes in insider holdings over consecutive quarters.

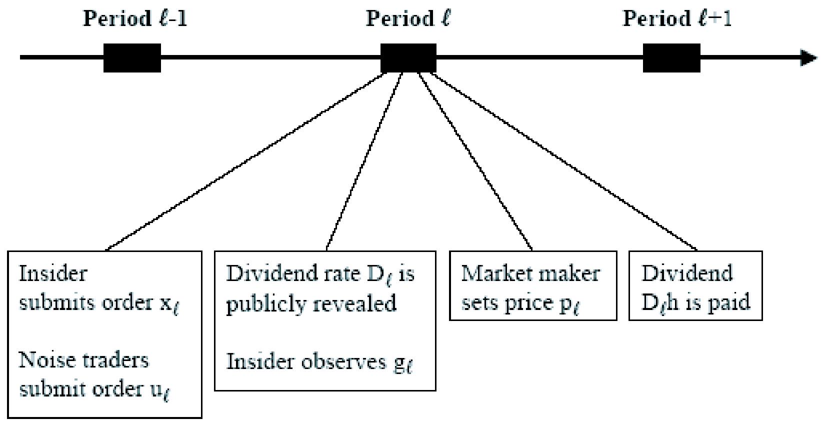

Figure 1: Timing of events in period ℓ.