



ELSEVIER

Journal of Financial Markets 5 (2002) 419–450

Journal of
FINANCIAL
MARKETS

www.elsevier.com/locate/econbase

Strategic trading and learning about liquidity[☆]

Harrison Hong^{a,*}, Sven Rady^b

^a *Graduate School of Business, Stanford University, 518 Memorial Way, Stanford, CA 94305-5015, USA*

^b *Department of Economics, University of Munich, D-80539 Munich, Germany*

Abstract

Many practitioners point out that the speculative profits of institutional traders are eroded by the difficulty in gauging the price impact of their trades. In this paper, we develop a model of strategic trading where speculators face such a dilemma because of incomplete information about time-varying market liquidity. Unlike the competitive market makers that they trade against, informed traders do not know the distribution of liquidity (“noise”) trades. Instead, they have to learn about liquidity from past prices and trading volume. This learning implies that strategic trades and market statistics such as informational efficiency are path-dependent on past market outcomes. Our paper also has normative implications for practitioners. © 2002 Elsevier Science B.V. All rights reserved.

JEL classification: G10; G14

Keywords: Strategic trading; Learning; Liquidity

1. Introduction

An issue fundamental to the analysis of asset markets is the determinants of speculative trades by large traders (e.g., active money management). To begin with, institutional ownership of common stocks is quite substantial and institutions

[☆]We thank an anonymous referee for many insightful comments. We would also like to thank Lanier Benkard, Larry Blume, Gilles Chemla, Thierry Foucault, Terence Lim, Jiang Wang and seminar participants at the Econometric Society Eighth World Congress, European Summer Symposium at Gerzensee, London School of Economics, Oxford University, the Stanford Finance Lunch, the Western Finance Association 2000 Meetings, and the University of Zürich for helpful comments and discussions.

*Corresponding author. Tel.: +1-650-725-6939; fax: +1-650-725-9932.

E-mail addresses: hghong@leland.stanford.edu (H. Hong), sven.rady@lrz.uni-muenchen.de (S. Rady).

account for a large fraction of trading volume on a number of exchanges.¹ Importantly, their trades can significantly affect price dynamics. Holthausen et al. (1990) find a price impact of about one percent for the largest buy and sell trades for randomly selected NYSE firms in 1983, while Keim and Madhavan (1996) find a price impact of around eight percent for block trades on small NYSE, AMEX and NASDAQ firms from 1985 to 1992.² The conventional wisdom is that these traders have an informational advantage over all other market participants (e.g., market makers) and hence can optimally speculate on their private information by taking into account the price impact of their trades.

However, a number of empirical findings call into question the completeness of this conventional wisdom. If large traders have superior information, then they ought to outperform the market or various passive benchmarks. This, however, does not appear to be the case.³ Many point out that this failure is due in part to the costliness of executing large trades, or an “implementation shortfall” (e.g., Perold, 1988; Chan and Lakonishok, 1993). This implementation shortfall is not due to obvious trading costs (bid–ask spreads, trading commissions) per se but rather traders’ difficulty in gauging the potential price impact of their trades.

Indeed, traders generally have incomplete information regarding the trading environment (i.e., non-fundamentals) that can significantly affect the profitability of their trades. For instance, they do not have information on market makers’ inventory fluctuations or other characteristics of order flow which would affect market liquidity and hence price impact (e.g., Schwartz and Whitcomb, 1988). Such liquidity shocks are naturally more directly observable to market makers than large traders who participate in the market on a less frequent basis. Importantly, market liquidity for a stock may vary substantially across trading days: using various measures of liquidity such as spreads and depth, Chordia et al. (2000) document substantial variability over time for all their liquidity measures. Time-varying liquidity is especially likely to be relevant for more thinly traded, small stocks where price impact is of central importance for institutional traders.

Many institutions therefore expend considerable resources on trading facilities and personnel engaged in acquiring information not only about asset payoffs (i.e., fundamentals) but also about the trading environment, that is, the various costs of executing large trades. Of course, the degree to which incomplete information regarding liquidity matters varies across institutions. For institutions that are resource rich (so have access to order flows) or that do not demand immediacy (so they can use limit orders), uncertainty about liquidity is less of an issue. But as

¹For instance, large institutions held discretionary control over more than half of the U.S. equity market at the end of 1996 (e.g., Gompers and Metrick, 2001) and accounted for over seventy percent of the trading volume among the New York Stock Exchange, the London Stock Exchange and the Tokyo Stock Exchange in 1990 (e.g., Schwartz and Shapiro, 1992).

²See also Kraus and Stoll (1972), Scholes (1972), Holthausen et al. (1987), Hausman et al. (1992), and Chan and Lakonishok (1993).

³Numerous studies have documented portfolio managers’ inability to outperform various passive benchmarks, despite considerable effort to analyze and select stocks (e.g., Fama, 1991; Chevalier and Ellison, 1999).

practitioners point out and empirical evidence confirms, most institutional speculators want to maintain anonymity and demand immediacy (and so use primarily market orders)—for these traders the issue of price impact uncertainty is a very relevant one.⁴

In response, such traders purchase information technology systems which forecast the price impact of trades.⁵ Such automated trading programs provide money managers with estimates of price impact based on past market outcomes such as a stock's past price fluctuations and trading volume. Fund managers then take into account these estimates and the error bands around these estimates in formulating their best execution strategies.

Motivated by these stylized facts, we develop an equilibrium model of speculation by large (informed) traders who face incomplete information about liquidity. The set-up is similar to a version of the Kyle (1985) model, in which risk-neutral informed traders strategically trade against risk-neutral and competitive market makers to exploit their private information. Noise trades are generated from a distribution whose variance evolves over time according to a Markov chain with two states, high or low. The true variance (or the true supply of liquidity) is known to market makers but imperfectly observed by informed traders.⁶ In addition, there are inter-dealer trades as market makers trade amongst themselves for exogenous reasons. In equilibrium, the sensitivity of price to order flow depends on the true state of liquidity (it is larger in the low than in the high liquidity state). Since informed traders have incomplete information about the true state of liquidity, they face uncertainty regarding the price impact of their trades.

However, these traders can extract valuable information about the true state of liquidity by learning from past market prices and trading volume.⁷ As a result of this learning, their beliefs about the true state of liquidity change over time. We solve in closed form for how informed traders update their beliefs conditionally on their private information about asset payoffs and the history of prices and trading volume (see Section 4).

Based on this solution, we establish a few simple results regarding the revisions of beliefs. After observing a price that deviates significantly from the forecast of the fundamental asset value based on public news, informed traders are more likely to revise their beliefs in favor of the low liquidity state than in favor of the high liquidity state. In addition to past prices, past trading volumes also provide valuable

⁴In this regard, Keim and Madhavan (1995) find that the majority of orders (approximately 87% of the total number and 90% of total value) in a dataset of 21 institutions are executed using market (or market-not-held) orders. We will revisit this issue in more detail in Section 2.5.

⁵A commonly used system among funds is developed by Investment Technology Group. Those with more resources may hire consulting firms specializing in executing large trades. A notable example is Plexus Group based in Los Angeles which provides clients with systems to estimate price impact and advice on how best to execute large trades.

⁶This assumption is a simple way of capturing large traders' incomplete information about aspects of the order flow better known to market makers (see Section 2.5).

⁷The presence of inter-dealer trades keeps price and trading volume from fully revealing the true liquidity state to informed traders each period.

information on the state of liquidity. Given a non-zero price deviation from the forecasted fundamental value, trading volume has a higher mean in the high liquidity state. Hence, observations of low trading volume tend to indicate an illiquid market and lead to corresponding revisions in beliefs. Jointly, large price deviations along with low trading volume suggest even more dramatic revisions of beliefs in favor of low liquidity.

The dependence of revisions in beliefs on past prices and trading volume implies that strategic trading will also be path dependent. When informed traders are uncertain about the liquidity in the market, they trade as if the price impact were an average of those in the high and low variance states, weighted by their beliefs about the likelihood of each state occurring. The more confident they are that the true state is low, the lower the liquidity they expect and the less aggressively they trade on their private information. If recent price deviations from fundamentals have been large, or volume has been low, then informed traders tend to revise their beliefs in favor of a less liquid market and end up trading less aggressively on their private information. Through their dependence on strategic trades, market statistics such as informational efficiency and trading volume then become dependent on the path of prices and trading volume as well.

The contributions of our paper are two-fold. First, our model yields testable implications on how past prices and trading volume help to predict time variation in strategic trading and other relevant market statistics. In fact, the behavior of the informed traders in our model (formulating trading decisions based on forecasts of liquidity) seems to fit very nicely with descriptions of fund managers formulating their trading decisions based on information technology systems that forecast price impact from past market outcomes. How one could test these predictions using data sets on the trading strategies of institutional investors (e.g., Chan and Lakonishok, 1993, 1995; Keim and Madhavan, 1995) will be discussed in Section 6.

Second and most importantly, our model addresses a problem that many institutional speculators care about—how to forecast the price impact of their trades. Most existing methods employed by practitioners tend to be statistical in nature (e.g., a black box approach of fitting price impact to a variety of market statistics and using these regressions to generate forecasts). In contrast, our model provides practitioners a consistent and parsimonious equilibrium framework with which to think about this problem. This model can perhaps be the basis for a more economically informed structural model to estimate and forecast the price impact of trades. In other words, our model has a number of normative implications for practitioners concerned with estimating the price impact of their trades.

In the following section, we develop a simple dynamic model to capture the ideas outlined above. Equilibrium trading and pricing strategies are derived in Section 3. The emphasis of Section 4 is on understanding how large traders learn from prices and trading volume about market liquidity. In Section 5, we discuss the time variation in strategic trading induced by learning about liquidity. In Section 6, we draw out the empirical implications of our model. We contrast our work with related papers in Section 7. Section 8 concludes. All proofs are in the appendix.

2. The model

In this section, we analyze a simple model of strategic trading. We consider a setup similar to that of Kyle (1985). One risky security is traded by three types of traders: informed, risk-neutral traders who possess identical private information about the liquidation value of the risky security, liquidity (“noise”) traders who trade for idiosyncratic or liquidity reasons, and competitive, risk-neutral market makers.

2.1. Market structure

The single asset is traded over a span of time equal to T trading periods. It is assumed that the ex post liquidation value of the asset at the end of period T is exogenously given by

$$v = \sum_{t=1}^T v_t, \tag{1}$$

where v_t for $t = 1, \dots, T$ are independently distributed normal random variables each having a mean zero and variance σ_v^2 . Any trader holding a share of the asset at the end of period T receives a liquidating dividend of v dollars. The value of the asset is gradually resolved over time. v_t becomes public information to all market participants in period t . By the end of period T , the value of the asset will be known by all.

The quantity traded by noise traders in period t , denoted by u_t , is drawn from a normal distribution with mean zero and a variance σ_u^2 that follows an exogenous two-state Markov process with possible values $\sigma_1^2 > \sigma_0^2 > 0$. Let ρ_k ($k = 0, 1$) denote the probability of staying in state σ_k^2 for one more period: $\rho_k = \Pr(\sigma_{t+1}^2 = \sigma_k^2 | \sigma_t^2 = \sigma_k^2)$. We assume that $\rho_k < 1$ for $k = 0, 1$ and that $\rho_0 + \rho_1 > 1$. Thus, there is no absorbing state and the Markov chain exhibits persistence in the sense that a given state is more likely to be reached from that same state than from the other. The variables v_1, \dots, v_T are independent of the noise trades and, conditional on the realizations of the variance process, the variables u_1, \dots, u_T are independent as well.

The total net quantity of inter-dealer trades, denoted by e_t , is drawn from a normal distribution with mean zero and variance σ_e^2 and is i.i.d. over time and independent of the v 's, u 's and the Markov chain for the variance of noise trades. The distribution of the inter-dealer trades is common knowledge.

2.2. Timing of trades

In period t , trading takes place in three steps as follows. In step one, the N informed traders submit market orders $x_{n,t}$ ($n = 1, \dots, N$). At the beginning of period t , v_t is observed by the informed traders. One can think of v_t as public news, which the informed traders get to peek at the period before it is released. The informed

traders know v_t but do not know u_t when placing the market order, nor do they know from which of the two possible distributions u_t was drawn.

In step two, the market makers determine the price p_t at which they trade the quantity necessary to clear the market. When doing so, they observe the total order flow $y_t = \sum_{n=1}^N x_{n,t} + u_t$, but not the $x_{n,t}$ or u_t separately. However, unlike the informed traders, they do know from which distribution u_t was drawn. That is, they know the realization of σ_t^2 .

In step three, market makers may trade amongst themselves for exogenous reasons unrelated to the pricing of the risky asset.⁸ Once all trades are completed at the end of period t , the total trading volume $q_t = |y_t + e_t|$ becomes public information along with the innovation v_t to the fundamental asset value.

2.3. Pricing

The competitive, risk-neutral market makers determine the price in period t based on the history of public information, past and current order flows, and past inter-dealer trades. The zero expected profit condition implies that p_t , the price set in period t by the market makers, is

$$p_t = E[v|\sigma_t^2, V_{t-1}, Y_t, E_{t-1}], \quad (2)$$

where $V_{t-1} = (v_1, v_2, \dots, v_{t-1})$, $Y_t = (y_1, y_2, \dots, y_t)$ and $E_{t-1} = (e_1, e_2, \dots, e_{t-1})$ are the histories of public news and private observations available to the market makers.

2.4. Informed trading and learning

In each period, a different and new cohort of N informed traders is active in the market. So the N informed traders in the market in period t trade on their private information v_t , then a different cohort of N informed traders participates in the market in period $t + 1$ to trade on their private information v_{t+1} , and so on. More formally, after observing the market history and the current realization of v_t , each of the N traders in period t chooses a trade $x_{n,t}$ so as to maximize his expected profits, i.e.,

$$E[x_{n,t}(v - p_t)|V_t, P_{t-1}, Q_{t-1}], \quad (3)$$

where $P_{t-1} = (p_1, \dots, p_{t-1})$ and $Q_{t-1} = (q_1, q_2, \dots, q_{t-1})$ are the histories of prices and volumes, respectively. Here, issues related to the timing of informed trading, which are important in Kyle (1985), do not arise since informed traders participate infrequently in the market. However, different trading periods are linked through the informed traders' updating of beliefs about the current variance of noise trades. In equilibrium, the revelation of v_t at the end of round t makes the trades $x_{n,t}$ publicly known. The new cohort of informed traders then uses this information and the observation of price and trading volume to update their belief about the variance of

⁸This assumption is simply a modeling device analogous to assuming that market makers can identify certain parts of the order flow as uninformative (see Section 2.5).

noise trades according to Bayes' law. In other words, informed traders in period t can learn from past prices and trading volume in all the periods preceding t .

Suppose that informed traders begin trading round t with the common belief assigning probability π_t to the event that $\sigma_t^2 = \sigma_1^2$. At the end of round t , this belief is updated to a belief π_t'' on the basis of the observed price p_t and trading volume q_t given the informed orders $x_{n,t}$.⁹ The belief π_{t+1} that is then taken into trading round $t + 1$ is obtained by adjusting π_t'' for the possibility of a state change between rounds t and $t + 1$. This process repeats itself every round. At the end of each round, the market makers know everything that the informed traders know, so they can infer the informed traders' updated belief.

2.5. Comments on the model

Next, we discuss the various assumptions of our model. We will explain how one can think about the real world problem of uncertain liquidity in terms of our model. Then we highlight the types of institutional traders and the market settings that our model really speaks to.

Several key assumptions drive our results. The first is that market makers have more information regarding the variance of liquidity trades than informed traders. This assumption is a simple way to model the asymmetry in information between market makers and traders regarding various factors that affect price impact.¹⁰ According to various descriptions of transactions costs and institutional investor trading strategies (e.g., Schwartz and Whitcomb, 1988), this assumption is eminently realistic since market makers tend to have more information about limit orders and other characteristics of order flow related to liquidity and price impact.¹¹

The second key assumption is that the liquidity state is more likely to persist than to change. While this assumption is irrelevant to the updating of beliefs at the end of a given period, it matters for the belief that is then held at the beginning of the next period. In fact, we shall see that π_{t+1} is strictly increasing in π_t'' if there is persistence ($\rho_0 + \rho_1 > 1$) and decreasing otherwise. The persistence assumption is a plausible one given the short to intermediate horizons that we are considering (e.g., Breen et al., 2000).

We also assume that there are inter-dealer trades in the background which are irrelevant for the purposes of pricing. This assumption is a simple way to keep the combination of price and trading volume from fully revealing the true state of liquidity in each period. There are other ways to prevent information about liquidity from being fully revealed each period. For instance, we could have also assumed that

⁹The updated belief π_t'' reflects both price and volume in round t ; in Section 4 we will also consider an updated belief π_t' that reflects just the price.

¹⁰Like many other rational expectations models, we assume that competitive market makers do not sell or compete on their private information. One can think of several different rationales behind this; see the many related models of trading on private information following Grossman and Stiglitz (1980).

¹¹There are some exceptions. Perhaps proprietary traders at Goldman Sachs may be able to peek at the order flow handled by the brokerage wing of Goldman. But such peeking is generally frowned upon and is certainly the exception.

market makers are able to identify certain parts of order flow as being uninformative or that informed traders observe trading volume with some noise. Inter-dealer trades just provide a more specific rationale for why price and trading volume need not fully reveal all information about factors affecting price impact.

Perhaps most importantly, we utilize the Kyle trading mechanism in modeling price impact and uncertain liquidity. We think this mechanism is a good way to capture the real world problem of uncertain liquidity. In practice, informed traders have more information about the execution price of their trades than is suggested by the Kyle model. For instance, some traders may be able to get quotes from brokers for execution prices on trades of various sizes. Most will have incomplete information about this price schedule. But even if they were certain of the price schedule this moment, they would still be uncertain about what liquidity will look like down the road. Imagine that the informed traders in our model lived for multiple periods and had some longer-lived information. Then their optimal trades today would depend on their expectations about future liquidity. So even if they were certain of liquidity this period, their uncertainty about liquidity in future periods would affect their trading decisions today. Kyle's batching process and the uncertainty over market depth in our model is a good way to capture this effect.

Indeed, empirical evidence speaks very clearly to the fact that most institutions face such incomplete information regarding the price impact of their trades and that this affects their decision making. As we mentioned in the introduction, the remedy used by many institutions consists of information technology systems that forecast the price impact of their trades based on past market outcomes for a stock. They base their trading decisions on these forecasts. This type of behavior very closely mirrors that of the informed traders in our model. Moreover, consulting practices have risen up to address precisely such issues. Interestingly, exchanges, noting the profitability of such services, have gotten into the act by promising to set up trading systems which allow a typical institutional trader more access to information regarding liquidity.¹²

The importance of price impact uncertainty may vary across different types of institutions. For the rare institutions with very sophisticated trading desks or access to order flow, this problem is mitigated somewhat precisely because they have the resources to acquire information about liquidity. But based on the descriptions of the business of active money management, such institutions are the exception rather than the norm.

Price impact uncertainty is also likely to be more important for institutions that speculate on short-lived information and hence require immediacy. Such speculation forces them into the more costly strategy of market orders and away from limit orders. Interestingly, a surprisingly high number of the orders placed by institutional traders (e.g., technical traders whose information decays quickly) are in fact market orders (e.g., Keim and Madhavan, 1995).

¹²For instance, the NYSE is promising an electronic trading tool for institutional investors called Institutional Express (e.g., Ruyter, 1999). The product will use software and connections to the NYSE floor to give investors more information about the order book.

Moreover, institutions with access to upstairs markets may want to negotiate their block trades, once again mitigating the effect of price impact uncertainty for such institutions. However, institutions who want to speculate on private information prefer to keep their anonymity and hence shy away from upstairs markets in which they have to reveal their identity before negotiating trades. Upstairs markets are also very expensive because there is limited capital available. So except for the most difficult of trades, most traders stay away from the upstairs market.

Finally, we assume that there is a new cohort of informed traders in each period t who only have access to private information for that period. The idea here is that we think of the T trading periods as being a long period of time and of each period as being a day or a week. So, over time, different informed traders have different pieces of private information and participate with their trades. Traders that come later inherit an endowment of information and beliefs from previous generations of traders.

This assumption provides tractability and allows us to focus on the information content of past prices and trading volume about liquidity. However, the cost is that the model cannot speak to another potentially interesting effect that arises out of uncertain liquidity. Suppose instead that each cohort of informed traders is allowed to trade twice before being completely replaced by a new cohort. (Different cohorts do not overlap.) Otherwise, the rest of the assumptions of the model stay the same.

In this set-up, the first-period action impacts on second-period profit because it affects the informed investors' second-period belief about liquidity. If there were no liquidity uncertainty, then the equilibrium would simply be the same as the two-period Kyle equilibrium. Otherwise, informed traders may have an incentive to experiment and sacrifice first-period trading profit for more precise information about liquidity, which may improve the profitability of their second-period trades. Unfortunately, even this simple set-up cannot be solved in closed form.

While this experimentation effect is an interesting one, the incentives to experiment are dampened by a number of factors in such a model. An important one is that the market makers will react to the experimentation of the informed traders. That is, if the informed traders experiment to get more precise information about liquidity, market makers will recognize that they have better information about liquidity and undertake actions to counteract their advantage and hence dampen their incentives to experiment in the first place.¹³ In other words, a number of our results are likely to be robust even in the face of possible experimentation, though more extensive research is needed.

¹³We can make this argument rigorous by considering the following comparative static. We exogenously increase the precision of the informed traders' signals (observed price and quantity) in the first period and ask how this affects their second-period profit. If the expected second-period profit (or value function) is convex in second-period beliefs, then there is value to experimentation. If it is concave in these beliefs, then there is no value. When investigating this set-up, we found that the relationship can indeed be concave over some range of beliefs.

3. Equilibrium

3.1. Prices and trading strategies

We characterize the unique linear equilibrium of the trading game. In a linear equilibrium, the informed traders place orders $x_{n,t} = \beta_{n,t}v_t$ and the market makers employ a pricing rule of the form

$$p_t = \begin{cases} \sum_{j=1}^{t-1} v_j + \lambda_{0,t}y_t & \text{if } \sigma_t^2 = \sigma_0^2, \\ \sum_{j=1}^{t-1} v_j + \lambda_{1,t}y_t & \text{if } \sigma_t^2 = \sigma_1^2, \end{cases} \quad (4)$$

where y_t is the combined order flow from informed traders and noise traders. Our notation conforms with that in Kyle (1985). Here, $\lambda_{k,t}$ ($k = 0, 1$) measures the responsiveness of price to order flow when the noise trades are drawn from the distribution with variance σ_k^2 . As news v_t about the terminal value of the asset gets revealed over time, the expected liquidation value of the asset changes. Market makers adjust the price in response to the part of current order flow that is not inter-dealer trades, hence conveys some of the informed traders' short-lived private information.

Given the pricing rule (4), the informed traders' uncertainty about the variance of noise trades translates directly into price impact uncertainty. The average $\lambda_t = \pi_t\lambda_{1,t} + (1 - \pi_t)\lambda_{0,t}$ measures the informed traders' expected price impact. As in Kyle (1985), the inverse $1/\lambda_t$ captures the market depth dimension of market liquidity: this is the order size that, from the perspective of an informed trader, will on average move the price by one unit. Of course, actual market depth depends on the true variance of noise trades and is measured by $1/\lambda_{k,t}$ with $k = 0$ or 1 .

As inter-dealer trades do not convey information about asset values, the market makers disregard these orders when setting the price. Therefore, inter-dealer trades do not contribute to the informed traders' price impact uncertainty and have no effect on their trading strategies.

The following result describes equilibrium strategies as a function of beliefs. To state it, we define

$$\bar{\sigma}_t^2 = \pi_t\sigma_1^2 + (1 - \pi_t)\sigma_0^2. \quad (5)$$

This is the variance of uninformed order flow from noise traders that informed traders expect in trading round t given that their belief at the beginning of the trading round is π_t .

Lemma 3.1. *The dynamic trading game has a unique linear equilibrium. In this equilibrium, all informed traders submit identical market orders, $x_t = \beta_t v_t$, expecting to achieve trading profits of $\beta_t v_t^2 / (N + 1)$ per trader. Market makers use a pricing strategy as in (4) with the informed traders' expected price impact given by*

$$\lambda_t = \pi_t\lambda_{1,t} + (1 - \pi_t)\lambda_{0,t} = \frac{1}{(N + 1)\beta_t}. \quad (6)$$

Given the belief π_t with which the informed traders enter trading round t , the equilibrium β_t , $\lambda_{0,t}$ and $\lambda_{1,t}$ are

$$\beta_t = \beta(\pi_t) = \frac{\sqrt{\gamma_t}}{\sqrt{N}\sigma_v} \tag{7}$$

and

$$\lambda_{k,t} = \lambda_k(\pi_t) = \frac{\sqrt{N}\gamma_t\sigma_v}{N\gamma_t + \sigma_k^2}, \tag{8}$$

where $\gamma_t = \gamma(\pi_t)$ is the unique positive root of the quadratic equation

$$N\gamma^2 + [\sigma_0^2 + \sigma_1^2 - (N + 1)\bar{\sigma}_t^2]\gamma - \sigma_0^2\sigma_1^2 = 0. \tag{9}$$

While this lemma pins down equilibrium behavior for any given belief of the informed traders, it does not say anything about the evolution of these beliefs. We will derive their equilibrium law of motion in Section 4. Before that, we wish to take note of some basic properties of the equilibrium and interpret the variable γ_t that is crucial to the calculation of equilibrium strategies.

3.2. Simple properties of the equilibrium

First, we point out some properties of the equilibrium price function. Since the market makers always know the true distribution from which the liquidity trades are drawn, the price set by the market makers is more sensitive to order flow (and the market is less deep) when the variance of noise trades is low: $\lambda_{0,t} > \lambda_{1,t}$. Furthermore, competition between the risk-neutral market makers ensures that, conditional on the true variance of noise trades, price is an unbiased estimator. So by taking the expectation over that variance, we see that price is also an unbiased forecast unconditionally. These results are stated formally in the following proposition.

Proposition 3.1. *In the equilibrium with price impact uncertainty, the price is more sensitive to order flow when the variance of noise trades is low. However, price is always an unbiased forecast of the liquidation value:*

$$p_t = E[v|V_{t-1}, p_t]. \tag{10}$$

We next point out a few basic properties of the informed traders' equilibrium strategies. To develop some intuition, consider first the special case where the informed traders have full information on the state of the variance of noise trades, i.e., $\sigma_t^2 = \sigma_k^2$ and $\pi_t = k$ where $k \in \{0, 1\}$ is the true state. Then $\gamma_t = \sigma_k^2$ and $\beta_t = \sigma_k/(\sqrt{N}\sigma_v)$, which is exactly the equilibrium strategy for the N -player one-period Kyle model with commonly known variance σ_k^2 . For non-degenerate beliefs, Eq. (7) shows that the informed traders behave as if they were in an N -player one-period Kyle model with known variance γ_t . In this sense, γ_t is the ‘‘certainty equivalent variance’’ for the informed traders. A second useful interpretation of the variable γ_t emerges when we calculate the variance of the informed order flow:

$\text{Var}[N\beta_t v_t | V_{t-1}, P_{t-1}, Q_{t-1}] = N^2 \beta_t^2 \sigma_v^2 = N\gamma_t$. Thus, the variance of informed order flow is proportional to γ_t , the factor of proportionality being the number of informed traders. Finally, we note from Eq. (6) that the trading aggressiveness β_t of an informed trader is inversely related to the *expected* price impact of his trades, λ_t . This generalizes a well-known property of the standard Kyle model to the case of price impact uncertainty.

To gain a better understanding of the effects of liquidity uncertainty on strategic trading, we fix an expected level of liquidity $\bar{\sigma}_t^2$ and vary the spread between the two possible variances of noise trades. The higher this spread, the more uncertain informed traders are about liquidity. We have the following result.

Proposition 3.2. *As the informed traders' uncertainty about liquidity increases (holding expected liquidity constant), they expect a higher price impact of their trades, trade less aggressively, and expect lower profits.*

In fact, the proof shows that, given an expected variance of noise trades $\bar{\sigma}_t^2$, the certainty equivalent variance γ_t is strictly decreasing in the difference $\sigma_1^2 - \sigma_0^2$. In view of Eq. (7), therefore, the aggressiveness of informed trading, measured by β_t , also decreases in $\sigma_1^2 - \sigma_0^2$. In particular, informed traders trade less aggressively than if the variance of noise trades were $\bar{\sigma}_t^2$ for sure.

This *uncertainty effect*, which is already discussed in Lindsey (1992) and Forster and George (1992), is easy to understand. Suppose that the informed traders are uncertain about liquidity, but use the trading strategy $x_t = \hat{\beta}_t v_t$ with $\hat{\beta}_t = \sqrt{\bar{\sigma}_t^2} / (\sqrt{N}\sigma_v)$ that would be the equilibrium strategy if the variance of noise trades were commonly known to be $\bar{\sigma}_t^2$. The market makers, who know the true variance σ_t^2 , determine the slope of their pricing strategy by estimating the asset value on the basis of the order flow. By the projection formula for normal variables, this slope coefficient turns out to be strictly convex in σ_t^2 . By Jensen's inequality, therefore, the expected slope $\lambda_t = \pi_t \lambda_{1,t} + (1 - \pi_t) \lambda_{0,t}$ exceeds $\tilde{\lambda}_t = \sqrt{N}\sigma_v / [(N + 1)\sqrt{\bar{\sigma}_t^2}]$, which would be the equilibrium slope if the variance of noise trades were $\bar{\sigma}_t^2$ for sure. Since it is the expected slope that determines the price response anticipated by the informed traders, they face higher expected costs to trading, and are thus better off trading less aggressively. At the same time, higher expected trading costs mean lower expected profits.

The other effect of liquidity uncertainty on strategic trading, which will be the focus of our paper, is how trading aggressiveness depends on beliefs about liquidity. The following proposition proves an important result that will be used later in the paper.

Proposition 3.3. *The informed traders trade more aggressively as they become more confident that the variance of noise trades is high.*

In fact, γ_t and in turn β_t are strictly increasing functions of the beliefs about liquidity, π_t . One can characterize this as the *level effect* of liquidity uncertainty on strategic trading.

Informed traders change their trading strategies over time as they learn from past prices and trading volume and update their beliefs about the distribution that generates the noise trades. The market makers condition their pricing rule on the variance of noise trades as well as the beliefs of the informed traders. Consequently, the equilibrium price functions also change over time. How beliefs are updated is the focus of the next section; how the time variation in beliefs affects strategic trading and other market statistics will then be the focus of the remainder of the paper.

4. Learning about liquidity

In this section, we study how informed investors can learn about the variance of liquidity trades and hence form forecasts of price impact to optimally determine their strategic trades. This analysis is interesting in light of the automated trading systems (described in the introduction) that many institutional traders use to help them formulate their trading decisions.

We begin by defining the *price innovation*

$$z_t = p_t - \sum_{j=1}^{t-1} v_j. \tag{11}$$

We call z_t a price innovation because $\sum_{j=1}^{t-1} v_j$ represents the best forecast of the fundamental value or fair price of the asset entering period t , and z_t represents the deviation of the new equilibrium price from this forecast. We note that in the absence of inter-dealer trades ($\sigma_e^2 = 0$), observing the equilibrium price innovation and the transaction volume q_t is sufficient to determine the current level of liquidity. This is because without inter-dealer trades, pricing rule (4) implies $|z_t| = \lambda_{k,t} q_t$ when the current variance of noise trades is σ_k^2 , so knowledge of z_t and q_t implies knowledge of $\lambda_{k,t}$. And since $\lambda_{0,t} \neq \lambda_{1,t}$, this fully reveals the current level of liquidity.

When there are inter-dealer trades ($\sigma_e^2 > 0$), informed traders cannot perfectly infer current liquidity from the observed price and volume since it is impossible to infer $\lambda_{k,t}$ from z_t and q_t without knowing e_t . However, the joint distribution of equilibrium price and volume when liquidity is high will in general be different from the distribution when liquidity is low, so informed traders can still extract valuable information about current liquidity from their observation of the market outcome.

4.1. The basics of updating

For expositional purposes, we will decompose the updating of beliefs in three steps. In the first step, informed traders update their belief from π_t to π'_t in response to the observed price innovation z_t . In the second step, they update from π'_t to π''_t in response to the observed transaction volume q_t . In the third step, they form the belief π_{t+1} by adjusting π''_t for the possibility of a change in the variance of noise trades between periods t and $t + 1$. We analyze each of these steps in turn.

Given the belief π_t held at the beginning of trading round t , and conditional on the realization of v_t and the variance $\sigma_t^2 = \sigma_k^2$, the price innovation $z_t = \lambda_{k,t}v_t = \lambda_{k,t}(N\beta_t v_t + u_t)$ is normally distributed with mean

$$E[z_t|\pi_t, v_t, \sigma_k^2] = \lambda_{k,t}N\beta_t v_t = \frac{N\gamma_t}{N\gamma_t + \sigma_k^2}v_t =: \mu_{k,t}v_t \tag{12}$$

and variance

$$\text{Var}[z_t|\pi_t, v_t, \sigma_k^2] = \lambda_{k,t}^2\sigma_k^2 = \frac{N\gamma_t\sigma_v^2\sigma_k^2}{(N\gamma_t + \sigma_k^2)^2} =: \Sigma_{k,t}^2 \tag{13}$$

We denote the corresponding density function by $f_k(z_t|\pi_t, v_t)$. By Bayes' rule, the belief held after observing v_t and z_t (but not q_t) is

$$\pi'_t = \frac{\pi_t f_1(z_t|\pi_t, v_t)}{\pi_t f_1(z_t|\pi_t, v_t) + (1 - \pi_t)f_0(z_t|\pi_t, v_t)} = \left[1 + \frac{1 - \pi_t}{\pi_t} \frac{1}{\ell_z(z_t|\pi_t, v_t)} \right]^{-1}, \tag{14}$$

where $\ell_z(z_t|\pi_t, v_t) = f_1(z_t|\pi_t, v_t)/f_0(z_t|\pi_t, v_t)$ is the likelihood ratio for z_t given π_t and v_t . Note that π'_t increases strictly in $\ell_z(z_t|\pi_t, v_t)$. In particular, we have $\pi'_t > \pi_t$ if and only if $\ell_z(z_t|\pi_t, v_t) > 1$, i.e., if and only if the observed price is more likely to have been generated from the distribution associated with the high variance of noise trades.

Next, conditional on z_t and $\sigma_t^2 = \sigma_k^2$, transaction volume

$$q_t = |N\beta_t v_t + u_t + e_t| = \left| \frac{z_t}{\lambda_{k,t}} + e_t \right| \tag{15}$$

is distributed like the absolute value of a normal random variable with mean $z_t/\lambda_{k,t}$ and variance σ_e^2 . Let $g_k(q_t|\pi_t, z_t)$ be the corresponding density function. By Bayes' rule again, the belief that the informed traders hold at the end of trading round t (i.e., after observing v_t, z_t and q_t) is

$$\pi''_t = \left[1 + \frac{1 - \pi'_t}{\pi'_t} \frac{1}{\ell_q(q_t|\pi_t, z_t)} \right]^{-1}, \tag{16}$$

where $\ell_q(q_t|\pi_t, z_t) = g_1(q_t|\pi_t, z_t)/g_0(q_t|\pi_t, z_t)$ is the likelihood ratio for q_t given π_t and z_t . Note that $\pi''_t > \pi'_t$ if and only if $\ell_q(q_t|\pi_t, z_t) > 1$, i.e., if and only if the observed order flow is more likely to have been generated from the distribution associated with the high variance of noise trades. It is straightforward to see that $\pi''_t > \pi_t$ if and only if the product $\ell_z(z_t|\pi_t, v_t)\ell_q(q_t|\pi_t, z_t)$ exceeds 1. Of course, this product is the likelihood ratio for (z_t, q_t) given π_t and v_t —the likelihood ratio one would use to update from π_t to π''_t in a single step.

Given the belief π''_t held at the end of period t , the informed traders assign probability $\pi''_t\rho_1$ to the event that $\sigma_{t+1}^2 = \sigma_t^2 = \sigma_1^2$, and probability $(1 - \pi''_t)(1 - \rho_0)$ to the event that $\sigma_t^2 = \sigma_0^2$ and $\sigma_{t+1}^2 = \sigma_1^2$. So the belief that these traders take into the subsequent trading round is

$$\pi_{t+1} = \pi''_t\rho_1 + (1 - \pi''_t)(1 - \rho_0). \tag{17}$$

It is instructive to re-write this as

$$\pi_{t+1} = \bar{\pi} + (\rho_0 + \rho_1 - 1)(\pi_t'' - \bar{\pi}), \tag{18}$$

where $\bar{\pi} = (1 - \rho_0)/(2 - \rho_0 - \rho_1)$ is the expected long-run fraction of time that the variance process will spend in state σ_1^2 . By assumption, $0 < \rho_0 + \rho_1 - 1 < 1$, so π_{t+1} is increasing in π_t'' , yet closer to $\bar{\pi}$ than π_t'' was. Thus, the possibility of a change in liquidity introduces mean reversion into the informed traders' beliefs: from the end of period t to the beginning of period $t + 1$, beliefs are adjusted towards the long-run average $\bar{\pi}$. This mean reversion is the stronger, the smaller $\rho_0 + \rho_1 - 1$, i.e., the lower the persistence of the variance process.

Since π_{t+1} is monotonic with respect to π_t'' , we can focus on the updating from π_t to π_t'' when studying the dynamics of beliefs. Our next aim is to provide some qualitative insights into these dynamics. We consider learning from prices first.

4.2. Learning from prices

Given the belief π_t and the private information v_t , the price signal z_t is drawn from one of two possible normal distributions. As $\mu_{1,t} < \mu_{0,t}$, the two possible distributions of z_t have different means (unless $v_t = 0$, which is a null event). The informed traders thus expect to see price innovations closer to zero when the market is deeper. If the two possible price distributions had the same variance, informed traders would therefore put more weight on the high liquidity state whenever they saw a price innovation close to zero. Examination of (13) shows, however, that the two distributions will in general have different variances.

Small absolute price innovations will still be evidence in favor of the deeper market if the deeper market has a *smaller* price variance. Yet, as we shall see shortly, this need not be the case. Nor is the deeper market necessarily associated with a *higher* price variance. While this would hold true if the market makers used the same pricing strategy in both states, they actually choose a flatter pricing strategy when the variance of noise trades is high ($\lambda_{1,t} < \lambda_{0,t}$). As the variance of price innovations is the product of the squared slope of the pricing strategy and the variance of noise trades, it is not clear, without further analysis, in what state the variance of price innovations will be higher.

To determine which of the two possible price distributions has the higher variance we have to compare the difference in the squared slopes of pricing strategies, $\lambda_{0,t}^2 - \lambda_{1,t}^2$, with the given difference in the variances of noise trades, $\sigma_1^2 - \sigma_0^2$. It is straightforward to show that $\lambda_{0,t}^2 - \lambda_{1,t}^2$ decreases monotonically in $N\gamma_t$, the variance of informed order flow. This is very intuitive: a growing variance of informed order flow means that orders from the informed traders become a larger component of total order flow, while the orders from noise traders become less important; consequently, the market makers' incentive to vary their pricing strategy according to the true variance of noise trades diminishes. The higher the variance of informed order flow, therefore, the less pronounced is the flattening of the market makers' pricing strategy when the true variance of noise trades is high, and the higher is the price variance in the high liquidity state relative to the price variance in the low

liquidity state. More precisely, the ranking of these two price variances depends on whether $N\gamma_t$ exceeds $\sigma_0\sigma_1$ or not. We thus need to distinguish between these two cases.

To figure out how traders revise their beliefs after observing price innovations, we will calculate the likelihood ratio. Consider the first case in which the variance of informed order flow is small, $N\gamma_t < \sigma_0\sigma_1$. In this instance, $\Sigma_{1,t}^2 < \Sigma_{0,t}^2$, and so the deeper market is associated with a *smaller* variance of price innovations. To see the implication of this for the updating of beliefs from π_t to π'_t , we compute the corresponding likelihood ratio:

$$\ell_z(z_t|\pi_t, v_t) = \frac{\Sigma_{0,t}}{\Sigma_{1,t}} \exp\left(-\frac{1}{2}\left[\left(\frac{z_t - \mu_{1,t}v_t}{\Sigma_{1,t}}\right)^2 - \left(\frac{z_t - \mu_{0,t}v_t}{\Sigma_{0,t}}\right)^2\right]\right). \quad (19)$$

For $\Sigma_{1,t}^2 < \Sigma_{0,t}^2$, the quadratic in the exponent of (19) is strictly concave in z_t , with a global maximum at $z_t = \mu_t v_t$ with

$$\mu_t = \frac{\Sigma_{0,t}^2 \mu_{1,t} - \Sigma_{1,t}^2 \mu_{0,t}}{\Sigma_{0,t}^2 - \Sigma_{1,t}^2} = \frac{N^2 \gamma_t^2}{N^2 \gamma_t^2 - \sigma_0^2 \sigma_1^2}. \quad (20)$$

Thus, the informed traders update the stronger in favor of the state where the market is deeper, the closer the realized innovation is to $\mu_t v_t$. Price innovations far off this mark are more likely to have been generated from the price distribution with higher variance, $\Sigma_{0,t}^2$, hence lead the traders to update in favor of σ_0^2 .

The second case is when the variance of informed order flow is large, $N\gamma_t > \sigma_0\sigma_1$. In this instance, $\Sigma_{1,t}^2 > \Sigma_{0,t}^2$, and so the deeper market is associated with a *higher* variance of price innovations. The quadratic in the exponent of (19) is now strictly convex in z_t , with a global minimum at $\mu_t v_t$. Price innovations far off this mark are more likely to stem from the price distribution with variance $\Sigma_{1,t}^2$, hence lead the traders to update in favor of σ_1^2 .

Note that as $N\gamma_t$ decreases, the factor μ_t defined above tends to 0; as $N\gamma_t$ increases, μ_t tends to 1. For sufficiently small and sufficiently large variances of informed order flow, therefore, a price innovation that is far off the mark $\mu_t v_t$ is also large in absolute value. This immediately yields the following proposition.

Proposition 4.1. *When the variance of informed order flow is small, then large price deviations from the forecast of the asset value based on public information lead to revisions of beliefs in favor of the low liquidity state. When the variance of informed order flow is large, then large price deviations from the forecast lead to revisions of beliefs in favor of the high liquidity state.*

The intuition behind this is simple. For relatively large and relatively small variances of informed order flow, the difference between the means of the two possible price distributions is small, so learning about the variance of prices is the dominant force in the updating. Large price innovations then simply speak in favor of whatever distribution has the higher variance.

Over an intermediate range of the variance of informed flow, namely for $N\gamma_t$ close to $\sigma_0\sigma_1$, the difference between the variances of the two possible price distributions is small, and learning about the means of prices is the dominant force in the updating. In that case, a price innovation close to zero is indicative of a deeper market. This is most easily seen by examining the borderline between the two cases discussed above, $N\gamma_t = \sigma_0\sigma_1$. By Eq. (9), this is equivalent to $\sigma_0^2 + \sigma_1^2 - (N + 1)\bar{\sigma}_t^2 - (N - 1)\sigma_0\sigma_1 = 0$, which holds if and only if the current belief is

$$\pi_t = \frac{\sigma_1 - N\sigma_0}{(N + 1)(\sigma_1 - \sigma_0)}. \tag{21}$$

In this instance, $\Sigma_{1,t}^2 = \Sigma_{0,t}^2 = \sigma_0\sigma_1\sigma_v^2/(\sigma_0 + \sigma_1)^2$, and the likelihood ratio simplifies to

$$\ell_z(z_t|\pi_t, v_t) = \exp\left(-\frac{\sigma_1^2 - \sigma_0^2}{\sigma_0\sigma_1\sigma_v^2}v_t\left[z_t - \frac{1}{2}v_t\right]\right). \tag{22}$$

This is strictly decreasing in z_t if $v_t > 0$, and strictly increasing if $v_t < 0$, in accordance with our earlier statement that for identical price variances, a price innovation close to zero is indicative of a deeper market.¹⁴

Note that the right-hand side of (21) equals $\frac{1}{2}$ for $N = 1$, decreases in N and σ_0 , and increases in σ_1 . If $\sigma_1 \leq N\sigma_0$ (which requires at least two traders, and will always hold for sufficiently large N) the right-hand side of (21) is non-positive, so $N\gamma_t > \sigma_0\sigma_1$ at all beliefs π_t . As the number of traders or the spread between the two possible variances of uninformed order flow grows, therefore, there is a larger range of beliefs where the state with a high variance of noise trades (and thus a deeper market) is characterized by more variable equilibrium prices.

In summary, belief revisions due to price observations depend in a non-trivial way on the variance of informed order flow, $N\gamma_t$, which in turn is determined by the number of traders and the belief with which they enter the trading period. In comparison, learning from volume, which we address next, turns out to be simpler.

4.3. Learning from trading volume

Given the belief π_t and the observed price innovation z_t , transaction volume q_t is the absolute value of a draw from one of two possible normal distributions with common variance σ_e^2 and means $z_t/\lambda_{k,t}$ ($k = 0, 1$).¹⁵

If $z_t = 0$, these means are identical, so the observation of q_t does not contain any new information, and $\pi_t'' = \pi_t'$. This is easy to explain. Given $z_t = 0$, the informed traders know already that the order flow $y_t = N\beta_t v_t + u_t$ must have been zero. Since the informed traders also know v_t , they can perfectly infer the noisy order flow u_t . This is as much information as they could possibly hope for—observing a second signal cannot improve matters.

¹⁴Of course, if $v_t = 0$, then the two possible distributions for z_t are identical, and there is no updating at all.

¹⁵Alternatively, we could consider the informationally equivalent volume signal q_t^2 . The two possible distributions from which this signal is drawn are non-central χ^2 distributions of degree 1 with non-centrality parameter $b_k = z_t^2/(\lambda_{k,t}^2\sigma_e^2)$ ($k = 0, 1$), hence mean $1 + b_k$ and variance $2 + 4b_k$.

Almost surely, however, $z_t \neq 0$. Then the distribution of q_t in the low liquidity state has a lower mean and a lower variance than in the high liquidity state. This suggests that informed traders will interpret high volume as evidence of high liquidity. To confirm this, we compute the relevant likelihood ratio:

$$\ell_q(q_t|\pi_t, z_t) = \exp\left(-\frac{1}{2\sigma_e^2}\left[\frac{1}{\lambda_{1,t}^2} - \frac{1}{\lambda_{0,t}^2}\right]z_t^2\right) \times \frac{\exp\left(\frac{1}{\sigma_e^2}\frac{z_t}{\lambda_{1,t}}q_t\right) + \exp\left(-\frac{1}{\sigma_e^2}\frac{z_t}{\lambda_{1,t}}q_t\right)}{\exp\left(\frac{1}{\sigma_e^2}\frac{z_t}{\lambda_{0,t}}q_t\right) + \exp\left(-\frac{1}{\sigma_e^2}\frac{z_t}{\lambda_{0,t}}q_t\right)}. \tag{23}$$

We first note that $\ell_q(q_t|\pi_t, z_t)$ is invariant to a sign change of z_t ; in other words, the likelihood ratio depends on the price innovation only through its absolute value, $|z_t|$.¹⁶ Second, the inequality $\lambda_{1,t} < \lambda_{0,t}$ implies that the likelihood ratio is increasing in q_t when $z_t \neq 0$. Thus, we have the very intuitive result that higher observed volume makes informed traders more optimistic about market liquidity. Third, if $z_t \neq 0$, then the likelihood ratio is smaller than 1 at $q_t = 0$; by continuity, this also holds for sufficiently small observed transaction volume. This means that very small volume is unambiguous evidence in favor of the low liquidity state and leads to a downward belief revision from π_t' to π_t'' .

We summarize these results in the following proposition:¹⁷

Proposition 4.2. *After a small volume of transactions has been observed, beliefs are revised in favor of the low liquidity state. Optimism about liquidity increases with observed transaction volume.*

The implications of this proposition and the results of Section 4.2 for strategic trading, informational efficiency and trading volume will be spelled out in Sections 5.2 and 5.3.

4.4. Speed of learning

The speed of learning, i.e., how much adjustment in beliefs one should expect to see in any given trading round, depends on how much information about the variance of noise trades the informed traders can extract from the equilibrium price and volume.

A measure of the information content of the price signal z_t is the relative entropy

$$e_t^z = \int_{-\infty}^{\infty} f_1(z_t|\pi_t, v_t) \ln \frac{f_1(z_t|\pi_t, v_t)}{f_0(z_t|\pi_t, v_t)} dz_t. \tag{24}$$

¹⁶This is obvious when one uses the alternative signal q_t^z since the non-centrality parameters of the relevant distributions depend on the price innovation only through z_t^2 .

¹⁷In this and the following propositions, we neglect the null event of a zero price innovation. Thus all statements are meant for $z_t \neq 0$.

The entropy is always non-negative, and equals zero if and only if the two density functions coincide. Moreover, a higher entropy means a higher information content of observed prices, and hence a stronger revision of beliefs from π_t to π'_t .¹⁸ It is easy to evaluate (24) for normal distributions; for the means and variances (12) and (13), the result is

$$e^{\bar{z}_t} = \frac{1}{2} \left[\frac{\Sigma_{1,t}^2 + (\mu_{1,t} - \mu_{0,t})^2 v_t^2}{\Sigma_{0,t}^2} - \ln \frac{\Sigma_{1,t}^2}{\Sigma_{0,t}^2} - 1 \right]. \tag{25}$$

This explicit representation allows us to perform comparative statics with respect to various factors that affect the speed of learning.

Proposition 4.3. *The speed of learning from prices increases with the absolute value of the realization of the informed traders' private information and with the difference between the two possible variances of noise trades. It decreases in the number of informed traders.*

The statement for $|v_t|$ is obvious from (25), and that for $\sigma_1^2 - \sigma_0^2$ and N follows from some straightforward but tedious computations. We therefore omit the proof.

The intuition behind Proposition 4.3 is the following. First, a larger (in absolute value) realization of the private information drives the means of the two possible price distributions further apart while keeping their variances unchanged, and so makes the observed price more informative. Second, the larger the spread between the two possible variances of noise trades, the larger is the spread between the corresponding price volatilities and means, and the easier it becomes to distinguish the two states. Third, as the number of informed traders grows, the distribution of noise trades becomes less important to the overall order flow and the formation of prices. As a consequence, the difference between the two possible price distributions shrinks, and price observations reveal less about the variance of noise trades.

The same analysis can be carried out for the information content of trading volume given the price observation.

Proposition 4.4. *The speed of learning from transaction volume (after observing the price) increases with the absolute value of the price innovation and with the difference between the two possible variances of noise trades. It decreases in the variance of inter-dealer trades and the number of informed traders.*

The proof is again omitted.

The intuition behind this result is clear. A larger (in absolute value) realization of the price innovation or a larger spread between the two possible variances of noise

¹⁸ In a simple example with binary uncertainty about a time-invariant fundamental asset value, binary signals (sale or purchase) and i.i.d. trades, O'Hara (1995, pp. 82–86) shows that beliefs converge exponentially at a rate equal to the entropy. Our set-up here is more complicated insofar as the entropy itself changes as beliefs change, but the basic relationship between speed of learning and entropy carries over.

trades drives the means and variances of the two possible distributions of transaction volume further apart. This makes volume more informative. As the variance of inter-dealer trades or the number of informed traders grows, on the other hand, noise trades become a less important component of total volume. This makes volume less informative.

5. Strategic trades and market statistics

Having established an understanding of the dynamics of beliefs, we turn to the implications of learning by informed traders for their trading strategies and in turn for market statistics such as informational efficiency and trading volume.

5.1. Time variation in beliefs about liquidity

So far, we have studied updating by an insider who has the informational advantage of knowing v_t . In this sub-section, we shall consider updating by an outside observer who enters trading round t with the belief π_t and observes the market outcome, but not the insiders' private information v_t .¹⁹ This is an important and empirically relevant case because the perspective of the outside observer is the appropriate one when we consider how past market outcomes affect future trading behavior and informational efficiency.

Conditional on the true variance of noise trades being $\sigma_t^2 = \sigma_k^2$, such an observer anticipates the equilibrium price innovation, $z_t = \lambda_{k,t}(N\beta_t v_t + u_t)$, to be normally distributed with mean zero and variance $\hat{\Sigma}_{k,t}^2 = \lambda_{k,t}^2(N^2\beta_t^2\sigma_v^2 + \sigma_k^2) = N\gamma_t\sigma_v^2/(N\gamma_t + \sigma_k^2)$. Note that this variance is unambiguously smaller in the state where the market is deeper: $\hat{\Sigma}_{1,t}^2 < \hat{\Sigma}_{0,t}^2$. This is because from the perspective of the outsider, the variance of price innovations has two components: one stemming from informed order flow ($\lambda_{k,t}^2 N\gamma_t$), the other from uninformed order flow ($\lambda_{k,t}^2 \sigma_k^2$).²⁰ When the variance of informed order flow is small ($N\gamma_t < \sigma_0\sigma_1$), both components are smaller in the state where the market is deeper: $\lambda_{1,t}^2 N\gamma_t < \lambda_{0,t}^2 N\gamma_t$ and $\lambda_{1,t}^2 \sigma_1^2 < \lambda_{0,t}^2 \sigma_0^2$. When the variance of informed order flow is large ($N\gamma_t \geq \sigma_0\sigma_1$), the “informed component” is still smaller in the state where the market is deeper, but the “uninformed component” is larger in this state: $\lambda_{1,t}^2 N\gamma_t < \lambda_{0,t}^2 N\gamma_t$ and $\lambda_{1,t}^2 \sigma_1^2 \geq \lambda_{0,t}^2 \sigma_0^2$. Yet precisely because the variance of informed order flow is large, the “informed component” dominates, and we obtain the same ranking of the two possible price variances as before. For someone who does not condition on the private information v_t , therefore, a deeper market unambiguously means a lower variance of equilibrium prices.

Consequently, an observer who sees a price innovation close to zero will put more weight on the state where the variance of uninformed order flow is high. Conversely, a very high or very low price innovation is ascribed to a lack of market depth, and

¹⁹The observer can calculate the belief π_t from the informed traders' prior belief at $t = 1$ and the public history of fundamentals, prices and volumes ($V_{t-1}, P_{t-1}, Q_{t-1}$).

²⁰For an insider who knows v_t , there is only the second component; cf. Section 4.2.

more weight is put on the state where the variance of uninformed order flow is low. In fact, the same arguments as above imply that in response to seeing an innovation z_t , the observer updates his belief to

$$\hat{\pi}_t = \left[1 + \frac{1 - \pi_t}{\pi_t} \sqrt{\frac{N\gamma_t + \sigma_0^2}{N\gamma_t + \sigma_1^2}} \exp\left(\frac{(\sigma_1^2 - \sigma_0^2)z_t^2}{2N\gamma_t\sigma_v^2}\right) \right]^{-1}, \tag{26}$$

which is strictly decreasing in $|z_t|$.

After observing the market price, the outside observer views the belief π_t' held by the informed traders as a random variable whose realization depends on the realization of the private information v_t . On average, the outsider expects this belief to equal the one he holds himself: $E[\pi_t' | \pi_t, z_t] = \hat{\pi}_t$. So the fact that $\hat{\pi}_t$ decreases in $|z_t|$ suggests that after a small price innovation, an informed trader will tend to be more confident that liquidity is high. The next proposition confirms this.

Proposition 5.1. *After a small absolute price innovation, an informed trader is more likely to revise his beliefs in favor of the high liquidity state than in favor of the low liquidity state.*

More precisely, the conditional probability of the event $\pi_{t+1} > \pi_t$ given π_t and z_t exceeds one half if $|z_t|$ lies below some threshold.

Since a small price innovation means a low information content of transaction volume, the previous result is independent of observed volume and holds even for the lowest possible volume, $q_t = 0$. Still, a higher volume observation increases the probability that the outside observer assigns to an upward revision of beliefs by the informed traders. This follows immediately from the monotonicity of the likelihood ratio $\ell_q(q_t | \pi_t, z_t)$ with respect to q_t .

Proposition 5.2. *For a given price innovation, the probability that the informed traders revise their belief in favor of the high liquidity state increases with the observed transaction volume.*

Formally, the conditional probability of the event $\pi_{t+1} > \pi_t$ given π_t, z_t and q_t (but not v_t) increases with q_t .

5.2. Time variation in strategic trades

The first implication of our model is that past market outcomes help predict how aggressively informed traders trade on their private information. Recall from Proposition 3.3 that as the informed investors' confidence in the high variance of liquidity trades increases, they trade more aggressively. Putting this together with the results of the previous sub-section, we immediately obtain the following proposition.

Proposition 5.3. *After a small absolute price innovation in a given period, the probability that the informed traders trade more aggressively on their private*

information in the following period exceeds one half. Moreover, this probability increases with the observed transaction volume.

The intuition for this result is clear. Small absolute price innovations tend to indicate a deeper market, and informed traders take advantage of a deeper market by trading more aggressively. Moreover, the informed traders' optimism about liquidity, and hence their trading aggressiveness in the following period, increases with the observed transaction volume.

5.3. Time variation in trading volume

We next consider the effects of liquidity uncertainty and learning on trading volume. Unlike the case of complete information in which past outcomes are uninformative as to the future level of volume, past outcomes do help predict future levels of trading volume when there is liquidity uncertainty. For example, we saw that a small price innovation tends to make informed traders more optimistic about the level of liquidity, hence more aggressive in their trading on private information. This causes an increase in the trading volume generated by informed traders. By definition, increased confidence in the high liquidity state also implies higher expected trading volume from uninformed traders. Finally, the market makers on average have to accommodate more trades as the order flows from informed and uninformed traders increase.

This yields the following result.

Proposition 5.4. *After a small absolute price innovation in a given period, the probability that trading volume is higher in the following period exceeds one half. Moreover, this probability increases with the observed transaction volume.*

5.4. Time variation in informational efficiency

Finally, we consider the effects of liquidity uncertainty and learning on informational efficiency. When the informed traders face uncertainty about the variance of liquidity trades, they are no longer able to trade exactly the "right" amount, and informational efficiency depends on the extent to which informed traders over- or underestimate market depth. If they overestimate it, they will trade too aggressively and reveal too much information. If they underestimate market depth, they will trade too gingerly and reveal too little information. After seeing the market outcome, an outside observer can draw inferences as to which of the two scenarios is more likely, and how much information has been revealed through price and volume.

In this regard, a measure of the informational efficiency of the equilibrium price in trading round t is

$$r_t^p = \text{Var}[v_t | \pi_t, V_{t-1}, p_t] = \text{Var}[v_t | \pi_t, z_t]. \quad (27)$$

Higher informational efficiency means a lower “residual variance” r_t^p . This is the appropriate measure for an observer whose information set is the same as that of an informed trader, except that he has no privileged information about v_t at the beginning of each trading period.

It is well known that in the Kyle model, the informed traders trade in such a way as to reveal exactly the fraction $N/(N+1)$ of their private information, i.e., the above measure of informational efficiency equals $\sigma_v^2/(N+1)$. Interestingly, under certainty about liquidity, informational efficiency is constant, and it does not depend at all on the (known) variance of liquidity trades. The main intuition behind this result is that the informed traders trade just aggressively enough to take advantage of any additional variance of noise trades.

With uncertainty about liquidity, however, informed traders are unable to fine-tune their trades in this way, and informational efficiency becomes a stochastic variable depending on beliefs and price innovations. In fact, when the absolute price innovation $|z_t|$ is large, it is quite likely, given the analysis in Section 5.1, that the market is less liquid. So the informed traders probably overestimated the variance of noise trades (and in turn market depth) and traded too aggressively. Hence, when $|z_t|$ is large, it is likely that the informed traders revealed too much information and so prices are more informationally efficient than in the certainty benchmark. Conversely, when $|z_t|$ is small, it is quite likely that the informed traders underestimated the variance of noise trades (and in turn market liquidity) and traded too gingerly. Hence, when $|z_t|$ is small, the informed traders probably did not reveal enough information and so prices are less informationally efficient than if there were complete information about liquidity trades. The following proposition confirms this.

Proposition 5.5. *With price impact uncertainty, informational efficiency of prices, as measured in (27), depends on the price and is strictly increasing in the absolute magnitude of the price innovation.*

More precisely, the proof shows that there is a cutoff level $|\bar{z}_t|$ such that for absolute price innovations above $|\bar{z}_t|$, informational efficiency of the price is higher than under complete information, i.e., $r_t^p < \sigma_v^2/(N+1)$; whereas for absolute innovations below $|\bar{z}_t|$, informational efficiency is lower than in the benchmark, i.e., $r_t^p > \sigma_v^2/(N+1)$. The proof of the proposition also shows that *average* informational efficiency is exactly as under certainty about liquidity: $E[r_t^p|\pi_t] = \sigma_v^2/(N+1)$.

6. Empirical implications

In this section, we draw out the empirical implications of our model. Without belaboring the point, our model matches a number of basic stylized facts about strategic trading that are absent from many of the existing models. Our informed traders face an implementation shortfall associated with incomplete information

about market liquidity and hence price impact. As a result, their speculative profits are eroded (see results in Section 3) and they use forecasts of the price impact of their trades (based on past market outcomes) in formulating their strategic trades (see results in Section 4). These findings fit very well with (1) what practitioners tell us is an important source of the under-performance (relative to passive benchmarks) on the part of active money management and (2) what institutional traders use (black-box systems to forecast price impact based on market outcomes) to address this uncertainty regarding price impact.

Beyond these general findings, our model also generates a number of testable implications. The most basic is how past prices affect future strategic trading. Extreme price innovations, reflecting a market with low liquidity, will lead to less aggressive trading by informed traders in the next period. One could test this prediction using data sets on institutional trades similar to those in Keim and Madhavan (1995) and Chan and Lakonishok (1993, 1995).

For instance, Keim and Madhavan use data on buyer and seller initiated trades for the period of January 1991 to March 1993. This data includes the date when the trading decision was made, the desired number of shares in the order at the time of the trading decision with a buy–sell indicator, the number of broker releases per order, the duration before orders were filled and the choice of order type (active market orders or more passive limit orders). Keim and Madhavan study various aspects of trading behavior motivated by models such as Kyle (1985) and more generic trading behavior such as feedback trading.

Here, we propose that one could use this data to get a reasonable measure of aggressiveness of trading on private information (i.e., estimate β_t using the desired market order size) and test the prediction of the model by regressing this measure on a measure of illiquidity derived from past prices (higher moments or some measure of deviation from forecasted fundamentals) and trading volume. Of course, one would need to control for a variety of factors, but seeing how future strategic trading depends on past prices and trading volume would be interesting. In a similar vein, it would not be hard to implement the predictions regarding the path dependence of trading volume on past prices. The result on informational efficiency is of course more difficult to test, though not entirely impossible with more data.

7. Comparison with related models

The goal of our paper is broadly related to the literature on strategic trading.²¹ For the most part, this literature—following Kyle (1985)—has maintained the assumption that traders are only uncertain about at most the first moment of any distribution. Some examples include the following: Back (1992) extends Kyle's

²¹ Admati (1991) and O'Hara (1995) provide surveys of this growing literature on strategic trading. Relatedly, a number of other papers in the market-microstructure literature have shown that strategic trading is an important part of explanations for a number of empirical findings regarding intra-day return and volume patterns (e.g., Admati and Pfleiderer, 1988; Foster and Viswanathan, 1990).

results to more general distributions of asset payoffs; Holden and Subrahmanyam (1992) consider the case of many informed traders who have the same information; Foster and Viswanathan (1996) and Back et al. (2000) assume many informed traders who have different pieces of information.

There are two emerging literatures which consider the impact of uncertainty over the second moments of distributions on equilibrium. Our paper is part of a literature which analyzes the effect of uncertainty over the second moment of noise trades.²² The paper by Lindsey (1992) comes closest to our model. Like us, he considers a market where informed traders do not know the variance of noise trades whereas market makers do. Unlike us, he develops a dynamic model with long-lived private information in which the low variance state can only take on the extreme value of zero variance (i.e., no noise trades whatsoever).²³ In each period, there is some probability that the low variance state is drawn—everyone finds out immediately and the informed traders make no trading profits after that point. So, the informed trader always knows that he is in the high variance state as long as the trading game continues. There is little learning from past prices and trading volume. The model is like Kyle (1985) except that there is a certain probability that the game will end each period and informed traders lose out on further speculative opportunities. Hence, informed traders tend to trade more aggressively on their private information than in the Kyle benchmark.

In our model, the low variance state is not necessarily zero variance, so there is scope for genuine learning from past prices and volume. Hence, our results on learning about the variance of noise trades are absent from Lindsey, as are all of our implications regarding the path dependence of strategic trading, informational efficiency and trading volume on past market outcomes. The existence of a link from past prices and trading volume to future actions is perhaps the most distinguishing feature of our model. This link generates a number of new empirical implications.

Another set of papers considers the impact of uncertainty over the precision of informed traders' information on equilibrium. Gervais (1997) considers a set-up in which market makers have incomplete information about the precision of the informed traders' information. The greater this precision, the more aggressively the traders speculate on their information. As such, the bid–ask spread and insider profits both increase with a first-order stochastic shift of the information precision, but decrease with a second order stochastic shift. The effect of these shifts on trading volume is in general ambiguous. Spiegel and Subrahmanyam (2000) also develop a

²²Forster and George (1992) consider a static model in which market makers have better information regarding liquidity trades. They examine the consequences of anonymity of liquidity trading for various welfare measures. Kumar and Seppi (1994) examine a model of arbitrage in index futures where the precision of heterogeneous signals received by market makers at different geographic locations is private information. Madrigal (1996) considers a Kyle-like model in which there are a set of traders that have private information about non-fundamentals (e.g., they know the realizations of noise trades unlike other informed traders who have information about fundamentals). He derives the effect of the interaction between these two sets of traders on equilibrium prices and market liquidity.

²³Lindsey (1992) also develops a static model in which the low variance state need not be zero. This static version pre-dates the static version of our model.

model in which market makers are uncertain about the precision of the informed traders' information. They show that the linear equilibria that often characterize papers in the literature (including our model) no longer exist. For very small trades, the price does not move very much since market makers take these trades to imply that the variance of the traders' information is small. However, for larger trades, the market makers markedly revise their beliefs about the variance upward, and so the supply curve steepens. Their model generates relationships between price impact and trade size that are more consistent with empirical findings than existing models. Moreover, uncertainty over the informational variance leads to price variances being serially correlated in time, similar to what would be characterized by ARCH models.

Whereas our model emphasizes the informed traders' uncertainty about liquidity (which is known to market makers), the models of Gervais and Spiegel and Subrahmanyam emphasize the uncertainty on the part of market makers regarding informational variance (known to informed traders). The most interesting results of our model are how traders learn about liquidity over time and how this learning affects their trading strategies. Our model does not generate many novel implications regarding price behavior because the market makers who set the price are perfectly informed about all second moments. In contrast, the models of Gervais and Spiegel and Subrahmanyam develop a number of interesting implications regarding price behavior because changes in the market makers' beliefs about informational variance have significant effects on the price. Unlike us, these models do not have many implications for the behavior of informed traders.

Our model is also related to a small but important literature on learning from trading volume in addition to prices. In the classic rational expectations models along the lines of Grossman and Stiglitz (1980), investors extract information from prices but trading volume contains no additional information. Yet, there appears to be ample anecdotal evidence such as technical analysis which indicates that investors do learn from trading volume. The first paper to provide a rationale for learning from volume is Blume et al. (1994), who consider a model of trading on private information in which traders have heterogeneous (rather than homogeneous) information quality. They show that trading volume provides information on information quality not deducible from the price statistic. Bernardo and Judd (1997) generalize Grossman and Stiglitz (1980) by considering non-normal distributions for asset payoffs, which results in non-linear price equilibria where trading volume contains additional information. In these two models, volume is informative about asset payoffs, whereas our paper emphasizes the information content of volume for non-payoff variables such as liquidity.

8. Conclusion

In this paper, we develop a model to study the effects of price impact uncertainty on the optimal trading strategies of large traders as well as the equilibrium feedback to prices. In our model, risk-neutral, informed traders strategically trade against risk-neutral competitive market makers to exploit their private information. Unlike

market makers, they have incomplete information about the distribution from which liquidity (“noise”) trades are drawn. As a result, they face uncertainty about the price impact of their trades (i.e., market liquidity). They optimally take into account this uncertainty in their trades and learn about market liquidity from past prices and trading volume.

To summarize briefly, we find the following. Extreme past price realizations and low trading volume tend to lead to revisions of informed traders’ beliefs in favor of a low liquidity state. In turn, strategic trades and other market statistics are dependent on the path of past prices and trading volume.

Even though the model is highly stylized, it does exhibit an attractive property: past prices and trading volume affect future strategic trades through a learning effect, which in turn feeds back to market statistics. This feature is missing in most of the existing models of strategic trading. Hence, a careful study of the dynamics may yield more insights into return and trading patterns.

Finally, our model is limited in its focus on the speculative motive for trade by large traders. In reality, large traders probably trade to both speculate and hedge. This dual trading motive is captured in a number of papers such as Admati and Pfleiderer (1988), Foster and Viswanathan (1990), Seppi (1990), and Vayanos (2001). Additionally, we ignore the dynamics in the unwinding of large positions, which is an important problem encountered by large institutions. This problem is addressed by Bertsimas and Lo (1998) who study the dynamic strategies of large traders who have a fixed time horizon to complete a trade. Relatedly, it would be interesting to consider in future research strategic experimentation on the part of the large traders. Large traders with longer trading horizons may have an important incentive to trade in such a way as to sacrifice short-term trading profits for more precise information about liquidity which will benefit them in the longer run. To fully understand the importance of uncertainty about market liquidity, a model should incorporate these other motives for trade as well.

Appendix. Proofs

Proof of Lemma 3.1. Given the informed traders’ strategies $x_{n,t} = \beta_{n,t}v_t$ ($n = 1, \dots, N$) in trading round t , the market makers can calculate the conditional expectation of v_t as

$$E \left[v_t \middle| \sigma_t^2, \sum_{n=1}^N \beta_{n,t}v_t + u_t = y_t \right] = \frac{\text{Cov}[v_t, \sum_{n=1}^N \beta_{n,t}v_t + u_t | \sigma_t^2]}{\text{Var}[\sum_{n=1}^N \beta_{n,t}v_t + u_t | \sigma_t^2]} y_t. \tag{A.1}$$

Thus, if $\sigma_t^2 = \sigma_k^2$ with $k \in \{0, 1\}$, the market makers use the pricing rule $p_t = \sum_{j=1}^{t-1} v_j + \lambda_{k,t}y_t$ where

$$\lambda_{k,t} = \frac{\sum_{n=1}^N \beta_{n,t} \sigma_v^2}{(\sum_{n=1}^N \beta_{n,t})^2 \sigma_v^2 + \sigma_k^2}. \tag{A.2}$$

Taking this pricing rule and the linear strategies of other informed traders as given, informed trader n 's objective in round t is to maximize

$$\begin{aligned}
 & E[x_{n,t}(v - p_t) | V_t, P_{t-1}, Q_{t-1}] \\
 &= E \left[x_{n,t} \left(v_t + \sum_{j=t+1}^T v_j - \lambda_{k,t} \left[x_{n,t} + \sum_{m \neq n} \beta_{m,t} v_t + u_t \right] \right) \middle| V_t, P_{t-1}, Q_{t-1} \right] \\
 &= x_{n,t} \left(v_t - \lambda_t \left[x_{n,t} + \sum_{m \neq n} \beta_{m,t} v_t \right] \right) \tag{A.3}
 \end{aligned}$$

with

$$\lambda_t = E[\lambda_{k,t} | V_t, P_{t-1}, Q_{t-1}] = E[\lambda_{k,t} | V_{t-1}, P_{t-1}, Q_{t-1}] = \pi_t \lambda_{1,t} + (1 - \pi_t) \lambda_{0,t}. \tag{A.4}$$

(The last equality in (A.3) uses the fact that $E[\lambda_{k,t} u_t | V_t, P_{t-1}, Q_{t-1}] = 0$, which is obvious since $E[\lambda_{k,t} u_t | V_t, P_{t-1}, Q_{t-1}, \sigma_k^2] = 0$ for $k = 0, 1$.) This expression for expected profits has the following interpretation. Informed traders face incomplete information about the type of market maker they are trading with, where the ‘‘type’’ is given by $\lambda_{k,t}$. In a linear equilibrium, they effectively trade with the average type, which is λ_t as given in (A.4).

The first-order condition yields $x_{n,t} = [1 - \lambda_t \sum_{m \neq n} \beta_{m,t}] v_t / [2\lambda_t]$, hence $\beta_{n,t} = [1 - \lambda_t \sum_{m \neq n} \beta_{m,t}] / [2\lambda_t]$ or $1 = \lambda_t [2\beta_{n,t} + \sum_{m \neq n} \beta_{m,t}] = \lambda_t [\beta_{n,t} + \sum_{m=1}^N \beta_{m,t}]$. This implies that $\beta_{n,t}$ is the same for all informed traders and equal to

$$\beta_t = \frac{1}{(N + 1)\lambda_t}. \tag{A.5}$$

By (A.3), therefore, each informed trader's expected profit in trading round t (conditional on v_t) equals $[\beta_t v_t](v_t - \lambda_t [N\beta_t v_t]) = \beta_t v_t^2 / (N + 1)$.

Using (A.4) and (A.2), we can rewrite (A.5) as

$$\beta_t = \frac{1}{N + 1} \left(\pi_t \frac{N\beta_t \sigma_v^2}{N^2 \beta_t^2 \sigma_v^2 + \sigma_1^2} + (1 - \pi_t) \frac{N\beta_t \sigma_v^2}{N^2 \beta_t^2 \sigma_v^2 + \sigma_0^2} \right)^{-1}. \tag{A.6}$$

Simplifying this equation, we see that it is a quartic in β_t , with only the even powers of β_t appearing. Writing $\gamma = N\beta_t^2 \sigma_v^2$, we find that (A.6) transforms into quadratic equation (9). As the left-hand side of (9) is negative at $\gamma = 0$, this equation has a unique positive root. Since a negative β_t would imply a negative λ_t in contradiction to the informed traders' second-order condition, we then obtain (7) and (8). This is the unique linear equilibrium. \square

Proof of Proposition 3.1. The claim that $\lambda_{0,t} > \lambda_{1,t}$ follows directly from Eq. (8).

In equilibrium, $p_t = E[v | \sigma_t^2, V_{t-1}, y_t]$. As $p_t = \sum_{j=1}^{t-1} v_j + \lambda_{k,t} y_t$ when $\sigma_t^2 = \sigma_k^2$ ($k = 0$ or 1), conditioning on $(\sigma_t^2, V_{t-1}, y_t)$ is equivalent to conditioning on $(\sigma_t^2, V_{t-1}, p_t)$, so $p_t = E[v | \sigma_t^2, V_{t-1}, p_t]$. This in turn implies $p_t = E[v | V_{t-1}, p_t]$. \square

Proof of Proposition 3.2. Define

$$\eta_{k,t} = \frac{d\gamma_t}{d\sigma_k^2} \Big|_{\sigma_i^2 = \text{constant}}$$

for $k = 0, 1$. Differentiating Eq. (9) with respect to σ_k^2 while holding σ_i^2 fixed, we obtain

$$(2N\gamma_t + \sigma_k^2)\eta_{k,t} = \sigma_{k'}^2 - \gamma_t,$$

where k' is the element of $\{0, 1\}$ different from k . At non-degenerate beliefs, $\sigma_0^2 < \gamma_t < \sigma_1^2$, hence $\eta_{1,t} < 0 < \eta_{0,t}$. The proposition now follows from Lemma 3.1. \square

Proof of Proposition 3.3. Eq. (A.6) characterises β_t as the unique fixed point of the function

$$F_t(\beta) = \frac{1}{N + 1} \left(\pi_t \frac{N\beta\sigma_v^2}{N^2\beta^2\sigma_v^2 + \sigma_1^2} + (1 - \pi_t) \frac{N\beta\sigma_v^2}{N^2\beta^2\sigma_v^2 + \sigma_0^2} \right)^{-1}.$$

As $F_t(\beta)$ increases in π_t , so does the fixed point β_t , and with it γ_t . \square

Proof of Proposition 5.1. We want to show that the median of π_{t+1} conditional on π_t and z_t exceeds π_t when $|z_t|$ is sufficiently small. By continuity, it suffices to show this for $z_t = 0$. In this case, the likelihood ratio simplifies to

$$\ell_z(0|\pi_t, v_t) = \frac{(N\gamma_t + \sigma_1^2)\sigma_0}{(N\gamma_t + \sigma_0^2)\sigma_1} \exp\left(\frac{N\gamma_t}{2} \left[\frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}\right] \frac{v_t^2}{\sigma_v^2}\right).$$

The random variable v_t^2/σ_v^2 has a $\chi^2(1)$ distribution with median m lying strictly between 0 and 1.²⁴ The median of $\ell_z(0|\pi_t, v_t)$ is

$$\frac{(N\gamma_t + \sigma_1^2)\sigma_0}{(N\gamma_t + \sigma_0^2)\sigma_1} \exp\left(\frac{N\gamma_t}{2} \left[\frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}\right] m\right),$$

and we are done if we can show that it always exceeds 1.

To this end, we consider the expression

$$L = \frac{(A + \sigma_1^2)\sigma_0}{(A + \sigma_0^2)\sigma_1} \exp\left(\frac{A}{2} \left[\frac{1}{\sigma_0^2} - \frac{1}{\sigma_1^2}\right] m\right)$$

for arbitrary $A > 0$. Clearly, $L \rightarrow 1$ as $\sigma_1 \rightarrow \sigma_0$. It is therefore enough to prove that L is strictly increasing in σ_1 . Now, straightforward computation shows that the partial derivative of L with respect to σ_1 has the same sign as the quadratic $Q = A^2m + A(1 - m)\sigma_1^2 + \sigma_1^4$, which is obviously positive. \square

Proof of Proposition 5.4. Given π_{t+1} (but not v_{t+1} or z_{t+1}), volume q_{t+1} is the absolute value of a draw from a normal distribution with mean zero and variance

$$\text{Var}[N\beta_{t+1}v_{t+1} + u_{t+1} + e_{t+1}|\pi_{t+1}] = N\gamma_{t+1} + \bar{\sigma}_{t+1}^2 + \sigma_\varepsilon^2.$$

²⁴A more precise numerical estimate is $0.45493 < m < 0.45494$.

By Proposition 4.2 and the definition of $\bar{\sigma}_{t+1}^2$ in Eq. (5), this variance is strictly increasing in π_{t+1} . A higher π_{t+1} therefore implies higher transaction volume q_{t+1} in the sense of first-order stochastic dominance. The result thus follows from Propositions 5.1 and 5.2. \square

Proof of Proposition 5.5. Given the belief π_t held by informed traders at the beginning of round t , and conditional on $\sigma_t^2 = \sigma_k^2$, the variables v_t and $z_t = \lambda_{k,t}(N\beta_t v_t + u_t)$ are jointly normal with $E[v_t|\pi_t, \sigma_t^2 = \sigma_k^2] = E[z_t|\pi_t, \sigma_t^2 = \sigma_k^2] = 0$, $\text{Var}[v_t|\pi_t, \sigma_t^2 = \sigma_k^2] = \sigma_v^2$, $\text{Var}[z_t|\pi_t, \sigma_t^2 = \sigma_k^2] = \lambda_{k,t}^2(N^2\beta_t^2\sigma_v^2 + \sigma_k^2)$ and $\text{Cov}[v_t, z_t|\pi_t, \sigma_t^2 = \sigma_k^2] = \lambda_{k,t}N\beta_t\sigma_v^2$. By the Projection Theorem,

$$\text{Var}[v_t|\pi_t, \sigma_t^2, z_t] = \text{Var}[v_t|\pi_t, \sigma_t^2] - \frac{\text{Cov}[v_t, z_t|\pi_t, \sigma_t^2]^2}{\text{Var}[z_t|\pi_t, \sigma_t^2]}$$

hence

$$\text{Var}[v_t|\pi_t, \sigma_t^2, z_t] = \sigma_v^2 - \frac{N^2\beta_t^2\sigma_v^4}{N^2\beta_t^2\sigma_v^2 + \sigma_t^2} = \sigma_v^2 \frac{\sigma_t^2}{N\gamma_t + \sigma_t^2}, \tag{A.7}$$

with the last equality following from Eq. (7).

Since $E[v_t|\pi_t, \sigma_t^2, z_t] = z_t$ by (10), we have $\text{Var}[E[v_t|\pi_t, \sigma_t^2, z_t]|\pi_t, z_t] = 0$, so the Law of Iterated Variances reduces to $\text{Var}[v_t|\pi_t, \pi_t, z_t] = E[\text{Var}[v_t|\pi_t, \sigma_t^2, z_t]|\pi_t, z_t]$. In Section 5.1 we calculated the conditional probability $\hat{\pi}_t$ that $\sigma_t^2 = \sigma_1^2$ given z_t . By (A.7), we then have

$$\begin{aligned} \text{Var}[v_t|\pi_t, z_t] &= \sigma_v^2 E\left[\frac{\sigma_t^2}{N\gamma_t + \sigma_t^2} \middle| \pi_t, z_t\right] \\ &= \sigma_v^2 \left[\hat{\pi}_t \frac{\sigma_1^2}{N\gamma_t + \sigma_1^2} + (1 - \hat{\pi}_t) \frac{\sigma_0^2}{N\gamma_t + \sigma_0^2} \right], \end{aligned} \tag{A.8}$$

which is strictly increasing in $\hat{\pi}_t$. Straightforward algebra using (9) shows that $\text{Var}[v_t|\pi_t, z_t] = \sigma_v^2/(N + 1)$ if and only if $\hat{\pi}_t = \pi_t$. By Eq. (26), $\hat{\pi}_t$ is strictly decreasing in $|z_t|$ and greater than π_t for $z_t = 0$, so there is a unique cutoff level $|\bar{z}_t|$ for the absolute price innovation such that $\hat{\pi}_t = \pi_t$ if and only if $|z_t| = |\bar{z}_t|$, and $\text{Var}[v_t|\pi_t, z_t\sigma_v^2/(N + 1)]$ if and only if $|z_t| < |\bar{z}_t|$. Finally, note that $E[\hat{\pi}_t|\pi_t] = \pi_t$, hence $E[\text{Var}[v_t|\pi_t, z_t]|\pi_t] = \sigma_v^2/(N + 1)$ by Eq. (A.8) and the remark immediately after it. \square

References

Admati, A.R., 1991. The informational role of prices: a review essay. *Journal of Monetary Economics* 28, 347–361.
 Admati, A.R., Pfleiderer, P., 1988. A theory of intraday patterns: volume and price variability. *Review of Financial Studies* 1, 3–40.
 Back, K., 1992. Insider trading in continuous time. *Review of Financial Studies* 5, 387–410.

- Back, K., Cao, H., Willard, G., 2000. Imperfect competition among informed traders. *Journal of Finance* 55, 2117–2155.
- Bernardo, A., Judd, K., 1997. Volume and price formation in an asset trading model with asymmetric information. Mimeo, UCLA.
- Bertsimas, D., Lo, A., 1998. Optimal control of execution costs. *Journal of Financial Markets* 1, 1–50.
- Blume, L., Easley, D., O'Hara, M., 1994. Market statistics and technical analysis: the role of volume. *Journal of Finance* 49, 153–181.
- Breen, W., Hodrik, L.S., Korajczyk, R.A., 2000. Predicting equity liquidity. Mimeo, Northwestern University.
- Chan, L., Lakonishok, J., 1993. Institutional trades and intra-day stock price behavior. *Journal of Financial Economics* 33, 173–199.
- Chan, L., Lakonishok, J., 1995. The behavior of stock prices around institutional trades. *Journal of Finance* 50, 1147–1174.
- Chevalier, J., Ellison, G., 1999. Are some mutual fund managers better than others? Cross-sectional patterns in behavior and performance. *Journal of Finance* 54, 875–899.
- Chordia, T., Roll, R., Subrahmanyam, A., 2000. Commonality in liquidity. *Journal of Financial Economics* 56, 3–28.
- Fama, E., 1991. Efficient capital markets II. *Journal of Finance* 46, 1575–1617.
- Forster, M.M., George, T.J., 1992. Anonymity in securities markets. *Journal of Financial Intermediation* 2, 168–206.
- Foster, F.D., Viswanathan, S., 1990. A theory of interday variations in volume, variance, and trading costs in securities markets. *Review of Financial Studies* 3, 593–624.
- Foster, F.D., Viswanathan, S., 1996. Strategic trading when agents forecast the forecasts of others. *Journal of Finance* 51, 1437–1478.
- Gervais, S., 1997. Market microstructure with uncertain information precisions: a new framework. Mimeo, University of Pennsylvania.
- Gompers, P.A., Metrick, A., 2001. Institutional investors and equity prices. *Quarterly Journal of Economics* 116, 229–259.
- Grossman, S., Stiglitz, J., 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70, 393–408.
- Hausman, J., Lo, A., MacKinlay, C., 1992. An ordered probit analysis of transaction stock prices. *Journal of Financial Economics* 31, 319–379.
- Holden, C., Subrahmanyam, A., 1992. Long-lived private information and imperfect competition. *Journal of Finance* 47, 247–270.
- Holthausen, R., Leftwich, R., Mayers, D., 1987. The effect of large block transactions on security prices: a cross-sectional analysis. *Journal of Financial Economics* 19, 237–267.
- Holthausen, R., Leftwich, R., Mayers, D., 1990. Large block transactions, the speed of response and temporary and permanent stock-price effects. *Journal of Financial Economics* 26, 71–95.
- Keim, D., Madhavan, A., 1995. Anatomy of the trading process: empirical evidence on the behavior of institutional traders. *Journal of Financial Economics* 37, 371–398.
- Keim, D., Madhavan, A., 1996. The upstairs market for large-block transactions: analysis and measurement of price effects. *Review of Financial Studies* 9, 1–36.
- Kraus, A., Stoll, H., 1972. Price impacts of block trading on the New York Stock Exchange. *Journal of Finance* 27, 569–588.
- Kumar, P., Seppi, D., 1994. Information and index arbitrage. *Journal of Business* 67, 481–509.
- Kyle, A., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315–1336.
- Lindsey, R., 1992. Market Makers, Asymmetric Information and Price Formation. Unpublished Ph.D. thesis, University of California, Berkeley.
- Madrigal, V., 1996. Non-fundamental speculation. *Journal of Finance* 51, 553–578.
- O'Hara, M., 1995. *Market Microstructure Theory*. Blackwell Publishers, Cambridge, MA.
- Perold, A., 1988. The implementation shortfall: paper versus reality. *Journal of Portfolio Management* 14, 4–9.

- Ruyter, T., 1999. The dinosaur, the rock and the hard place. Plan Sponsor (online magazine found at <http://www.assetpub.com/psmarch99/psmarch99.html>).
- Scholes, M., 1972. The market for securities: substitution versus price pressure and the effects of information on share prices. *Journal of Business* 45, 179–211.
- Schwartz, R., Whitcomb, D., 1988. Transaction costs and institutional investor trading strategies. Mimeo, New York University.
- Schwartz, R., Shapiro, J., 1992. The challenge of institutionalism for the equity markets. Mimeo, New York University.
- Seppi, D., 1990. Equilibrium block trading and asymmetric information. *Journal of Finance* 45, 73–94.
- Spiegel, M., Subrahmanyam, A., 2000. Asymmetric information and disclosure rules. *Journal of Financial Intermediation* 9, 363–403.
- Vayanos, D., 2001. Strategic trading in a dynamic noisy market. *Journal of Finance* 56, 131–171.