Econometrica, Vol. 84, No. 4 (July, 2016), 1441–1475

INSIDER TRADING, STOCHASTIC LIQUIDITY, AND EQUILIBRIUM PRICES

BY PIERRE COLLIN-DUFRESNE AND VYACHESLAV FOS[1]

We extend Kyle's (1985) model of insider trading to the case where noise trading volatility follows a general stochastic process. We determine conditions under which, in equilibrium, price impact and price volatility are both stochastic, driven by shocks to uninformed volume even though the fundamental value is constant. The volatility of price volatility appears 'excessive' because insiders choose to trade more aggressively (and thus more information is revealed) when uninformed volume is higher and price impact is lower. This generates a positive relation between price volatility and trading volume, giving rise to an endogenous subordinate stochastic process for prices.

KEYWORDS: Kyle model, insider trading, asymmetric information, liquidity, price impact, market depth, stochastic volatility, volume, subordinate process, execution costs, continuous time.

1. INTRODUCTION

IN HIS SEMINAL CONTRIBUTION, Kyle (1985) derived the equilibrium price dynamics in a model where a large trader possesses long-lived private information about the value of a stock that will be revealed at some known date, and optimally trades continuously into the stock to maximize his expected profits. Risk-neutral market makers try to infer from aggregate order flow the information possessed by the insider. Because order flow is also driven by uninformed 'noise traders,' who trade solely for liquidity purposes, prices are not fully revealing. Instead, prices respond linearly to order flow. Kyle's lambda, which measures the equilibrium price impact of order flow, is constant in the model. Price volatility is also constant and independent of noise trading volatility. It depends only on the prior uncertainty about fundamentals.

In this paper, we generalize Kyle's (1985) model to allow the volatility of noise trading to change stochastically over time. Empirically, it is clear that trading volume and its volatility fluctuate widely over time. Given the amount of resources spent by market participants (witness, in particular, the growth of the high-frequency industry) to separate the component of order flow that is

[1]We thank Bruno Biais, Kerry Back, Darrell Duffie, Thierry Foucault, Gur Huberman, Larry Glosten, Wei Jiang, Semyon Malamud, Ioanid Rosu (WFA discussant), Norman Schuerhoff, James Thompson (NFA discussant), and seminar participants at Bocconi University, Copenhagen Business School, Cass Business School, Duke University, ETH Zurich, University of Helsinki, Kellogg School of Management, University of Illinois at Urbana-Champaign, ESADE, LSE, LBS, Unversity of Lugano, NYU, Stanford Business School, Haas School of Business, The Swedish House of Finance, University of Toronto, University of Toulouse, 9th Annual Cowles General Equilibrium Conference, AFFI 2013 (Best paper award), China International Conference (TCW 2013 Best paper award), the WFA 2014 Annual Meetings, the SOFIE 2015 meetings in Aarhus, the SFI-Finrisk meetings and, especially, Julien Hugonnier, for many helpful comments.

 DOI: 10.3982/ECTA10789

informed from that which is uninformed, it seems useful to understand how stochastic noise trading volatility impacts equilibrium price dynamics and liquidity. We ask the following questions. How does the informed investor adapt his optimal trading strategy to account for stochastic volatility of uninformed trading volume? How are the equilibrium price and volatility dynamics affected by shocks to the volatility of uninformed trading volume? How are adverse selection costs paid by uninformed traders affected?

Below, we characterize the equilibrium for general noise trading volatility dynamics.

First, we find that the equilibrium price follows a new class of bridge process that converges almost surely at maturity to the value known ex ante only to the insider. This guarantees that all private information will have been incorporated into the equilibrium price at maturity and generalizes the result proved in Back (1992) that the equilibrium price in the continuous-time Kyle model follows a standard (i.e., constant volatility) Brownian Bridge. Under plausible conditions which we derive in the paper, the equilibrium price volatility is stochastic. In that case, our model generates 'excess stochastic volatility,' since non-payoff-relevant shocks that affect noise trading volatility may affect the stock price volatility because the market maker rationally anticipates that in periods where noise traders are more active, the informed trader is more aggressive. This generates a positive volume-volatility relation, consistent with the empirical evidence reported in Gallant, Rossi, and Tauchen (1992) and different from the standard Kyle model, where price volatility is constant and independent of the level of noise trading volatility. The equilibrium price is a subordinate process with an endogenous time-change process driven by uninformed volume volatility as proposed by Clark (1973) and consistent with the results in Lamoureux and Lastrapes (1990) that ARCH effects tend to disappear when volume is included in the variance dynamics.

Second, we find that the optimal trading strategy for the insider is to trade proportionally to the undervaluation of the asset at a rate that is inversely related to price impact but increasing in a measure of the current 'state of liquidity' that we identify below. We thus show that it is optimal for the insider to 'time liquidity.' This is consistent with the empirical evidence in Collin-Dufresne and Fos (2015), who documented, using data from SEC filings, that informed activist-shareholder investors trade more aggressively when abnormal volume is higher and measured price impact is lower.

Third, we find that price impact is stochastic and negatively correlated with uninformed volume volatility. In fact, we find that price impact is a submartingale, expected to increase over time. This contrasts our model from much of the previous literature. In the original Kyle model, price impact is constant. In extensions of that model (Back (1992), Back and Pedersen (1998), Baruch (2002), Back and Baruch (2004)), price impact is either a martingale or a supermartingale. The intuition why price impact is a submartingale in our model is that, with stochastic noise trading volatility, the insider has a 'liquidity timing' option to wait for higher noise trading volatility to trade. In equilibrium, price impact must increase on average to entice the insider to trade early and give up his option to wait for better liquidity states.

The prediction that price impact is expected to rise on average seems consistent with the empirical evidence in Madhavan, Richardson, and Roomans (1997) who found that estimated execution costs rise significantly on average over the day.

Last, in our model the rate of price discovery increases when noise trading volatility increases and price impact decreases, which is different from the standard Kyle model, where private information is revealed at a constant rate that is independent of the level of noise trading volatility. This also implies that aggregate execution (or slippage) costs incurred by uninformed liquidity traders are stochastic and path-dependent in our model, unlike in Kyle (1985) or Back and Pedersen (1998), where unconditional expected execution costs always equal realized costs path-wise. As we show below, when price impact is stochastic, then aggregate execution costs are poorly proxied by the time-average of price impact, which is often used in practice to measure such costs.

Related Literature

This paper is related to a long list of papers that study the impact of asymmetric information on asset prices, volatility, volume, and market liquidity. Admati and Pfleiderer (1988), Foster and Viswanathan (1990, 1993) first proposed extensions of the Kyle model to dynamic economies with myopic agents (essentially a sequence of one-period Kyle models where agents only trade on short-lived information), where they generated time variation in price volatility and trading volume. In their models, price volatility changes because new (short-lived) private information is produced every period in response to deterministic changes in noise trading volatility.

Foster and Viswanathan (1995) developed and empirically tested a model with stochastic variation in noise trading volume to study the joint behavior of price volatility, volume, and price impact.[2] However, as they pointed out (p. 380), "an important limitation of [their] analysis is that they assume that information is short-lived," as in Admati and Pfleiderer (1988) and Foster and Viswanathan (1990). Our model provides the more general characterization with long-lived information.

Back and Pedersen (1998) (BP) extended Kyle's original model to allow for deterministically changing noise trading volatility to capture intra-day patterns (clustering) of liquidity trading, effectively analyzing a non-myopic version of Admati and Pfleiderer (1988). They obtained that price volatility follows the same deterministic pattern as noise trading. However, because they assumed

[2]Watanabe (2008) also captured GARCH features in equilibrium prices, by directly incorporating stochastic volatility in the (short-lived) private information process.

noise trading is deterministic, their model cannot generate any systematic patterns in price impact (e.g., it is a constant in the Gaussian model). This implies that in their model, 'expected execution costs of liquidity traders do not depend on the timing of their trades' (BP, p. 387). As we show below, only if noise trading volatility is stochastic can the insider's 'liquidity timing' option generate stochastic price volatility and a nonzero correlation between price volatility, market depth, and (uninformed) trading volume.

Our model is also related to a large literature in financial econometrics initiated by Clark (1973), which models the price change process as subordinated to the normal distribution, with a directing process related to 'informed volume' usually modeled via some latent 'information variable' extracted from volume and price volatility (see Andersen (1996), Tauchen and Pitts (1983), Richardson and Smith (1993) for discussions of the so-called MDM 'mixture of distributions model').

In our model, price follows such a subordinate process (in fact, a time-changed Brownian motion). Interestingly, the model provides the endogenous directing (or time-change) process as a function of the uninformed trading volume process. In the classic MDM model, it is typically assumed that the subordinating process is driven by an information flow variable that represents new information arrival and that is orthogonal to uninformed volume (e.g., Andersen (1996)). As such, this information flow variable is latent and has to be estimated from volume and volatility. Our model shows that even if there is no exogenous information flow, but informed investors can strategically time their trades, then stochastic uninformed volume will drive the directing process and thus price volatility, albeit in a nonlinear fashion which reflects the informed trader's option to time his trades as a function of the "liquidity state."

Section 2 introduces the general model and solves for an equilibrium. Section 3 investigates two specific cases to emphasize key features of the equilibrium. Section 4 discusses extensions such as more general correlation structures and the link to the MDM literature. Section 5 concludes. All proofs are in the Appendix.

2. INFORMED TRADING WITH STOCHASTIC LIQUIDITY SHOCKS

We extend Kyle's (1985) model (in the continuous-time formulation given by Back (1992)) to allow for time-varying volatility of noise trading. As in Kyle, we assume there is an insider trading in the stock with perfect knowledge of the terminal value v. The insider is risk-neutral and maximizes the expectation of his terminal profit:

$$\max_{\theta_t \in \mathcal{A}} \mathrm{E}\left[\int_0^T (v - P_t)\theta_t \, dt \middle| \mathcal{F}_t^Y, v\right], \tag{1}$$

where we denote by $\mathcal{F}_t^Y$ the information filtration generated by observing the entire past history of aggregate order flow Y (which we denote by $Y^t = \{Y_s\}_{s \leq t}$).

In addition, the insider knows the actual value v of the stock, and, of course, his own trading. Following Back (1992), we assume that the insider chooses an absolutely continuous trading rule θ that belongs to an admissible set $\mathcal{A} = \{\theta \text{ s.t. } \mathrm{E}[\int_0^T \theta_s^2\, ds] < \infty\}$.[3]

The market maker is also risk-neutral, but does not observe the terminal value. Instead, he has a prior that the value v is normally distributed $N(P_0, \Sigma_0)$. The market maker observes the aggregate order flow arrival, which is the sum of the insider's demand and the 'noise-trader' demand:

$$(2) \qquad dY_t = \theta_t\, dt + \sigma_t\, dZ_t,$$

where Z_t is a standard Brownian motion independent of v. The market maker absorbs the total order flow by trading against it at a price set so as to break even on average. Since we assume the market maker is risk-neutral, equilibrium break-even requires that the market clearing price is

$$(3) \qquad P_t = \mathrm{E}\big[v|\mathcal{F}_t^Y\big].$$

If noise trading volatility were constant, then this setup would be the Kyle–Back model. Instead, we assume that the noise trading volatility, σ_t, follows a general stochastic process. Specifically, we assume there is a (possibly discontinuous) martingale W such that

$$(4) \qquad \frac{d\sigma_t}{\sigma_t} = m\big(t, \sigma^t\big)\, dt + \nu\big(t, \sigma^t\big)\, dW_t,$$

where the growth rate (m) and volatility (ν) of σ can depend on its past history, which we denote by σ^t, but not on the history of Y (or Z). Further, we assume they satisfy standard integrability requirements for the SDE to admit a unique strong solution (e.g., Liptser and Shiryaev (2001, Theorem 4.6)). For simplicity, we assume for now that W_t and Z_t are independent, but we relax this assumption in Section 4. The main (economic) restriction we require throughout is that the noise trading volatility process is independent of the insider's private information and that it may not be Granger-caused by order flow (i.e., m and ν cannot depend on v or on Y^t).

We assume that both the market maker and the insider observe the history of σ perfectly. Since, in continuous time, observing aggregate order flow allows to

[3]As in Kyle's model, we assume the insider submits market orders dX_t that will be filled by the market maker at price $P_{t+dt} = P_t + dP_t$. Thus, his profits are, assuming a zero risk-free rate, $\int_0^T (v - P_{t+dt})\, dX_t = \int_0^T (v - P_t)\, dX_t - \int_0^T dP_t\, dX_t$. Back (1992) showed that it is optimal to choose an absolutely continuous $dX_t = \theta_t\, dt$ so that the second term is zero (otherwise, that term is always negative due to price impact). Note that this requires that volatility be common knowledge. Else there could be multiple equilibria where the strategy of the insider is not absolutely continuous.

measure its quadratic variation perfectly, the filtration $\mathcal{F}_t^Y$ contains both histories of order flow (Y^t) and of volatility (σ^t). The assumption that market participants observe order flow may be partially justified by the fact that volume and order-book information are available in many markets. In the standard Kyle model, assuming that the informed investor observes total order flow is innocuous, since if the insider only observes equilibrium prices, he can typically recover the total order flow (and, given his own trading, the noise trading order flow). However, when uninformed order flow has stochastic volatility, observing only prices may not be sufficient to recover noise trading volatility, as we give some examples below, where the equilibrium price is independent of noise trading volatility, even though the insider's trades depend on it.

An equilibrium is a price process and an admissible trading strategy, (P_t, θ_t), that satisfy the market maker's rationality condition (3) while solving the insider's optimality condition (1).

To solve for an equilibrium, we proceed in a few steps. First, we derive the dynamics of the stock price consistent with the market maker's risk-neutral filtering, conditional on a conjectured strategy rule followed by the insider. Then, we solve the insider's optimal portfolio choice problem, given the assumed dynamics of the equilibrium price. Finally, we show that the conjectured rule by the market maker is indeed consistent with the insider's optimal choice. Before stating the full theorem (proved in the Appendix), we provide a sketch of proof to give some intuition and highlight the main differences with respect to previous literature.

In a linear equilibrium where the insider chooses a trading strategy of the form $\theta_t = \beta_t(v - P_t)$ for some $\mathcal{F}_t^Y$ adapted process β_t, market makers will move prices linearly with order flow:

$$dP_t = \lambda_t \, dY_t, \tag{5}$$

where $\lambda_t = \frac{\beta_t \Sigma_t}{\sigma^2}$. This follows from the usual Kalman-filter equations, which also imply that the posterior variance for the market maker $\Sigma_t = \mathrm{E}[(v - P_t)^2 | \mathcal{F}_t^Y]$ satisfies

$$d\Sigma_t = -\lambda_t^2 \sigma_t^2 \, dt. \tag{6}$$

We conjecture that the value function of the insider is quadratic of the form $J_t = \frac{(v-P_t)^2 + \Sigma_t}{2\lambda_t}$. Then, applying Itô's rule and using (5) and (6) above, we find

$$\begin{aligned} &dJ_t + (v - P_t)\theta_t \, dt \\ &\quad = \frac{(v-P_t)^2 + \Sigma_t}{2} \, d\frac{1}{\lambda_t} - (v - P_t) \, dP_t \, d\frac{1}{\lambda_t} - (v - P_t)\sigma_t \, dZ_t. \end{aligned}$$

It follows that if we can find a market depth process ($\frac{1}{\lambda_t}$) that is (i) a martingale ($\mathrm{E}[\frac{1}{\lambda_t}] = 0$), and (ii) independent of price changes ($dP_t\, d\frac{1}{\lambda_t} = 0$), then our conjectured value function satisfies

$$J_0 = \mathrm{E}\left[\int_0^T (v - P_s)\theta_s\, ds + J_T\right].$$

Since $J_T \geq 0$, this implies that J_0 will be the optimal value function if there exists a trading strategy θ_t^* such that $P_T = v$ a.s. (which implies $J_T = 0$).[4] Theorem 1 below establishes that such a strategy exists. Note that it implies a third restriction on the market depth process, namely that (iii) the equilibrium is fully revealing at T ($\Sigma_T = 0$). The new feature of our equilibrium relative to previous literature is the characterization of the market depth process consistent with conditions (i), (ii), (iii) above and given the stochastic noise trading volatility process (4). Specifically, this requires solving a 'forward-backward' system:

$$\mathrm{E}\left[d\frac{1}{\lambda_t}\right] = 0, \tag{7}$$

$$d\Sigma_t = -\lambda_t^2 \sigma_t^2\, dt, \tag{8}$$

given initial condition Σ_0 and terminal condition $\Sigma_T = 0$. The nice feature of this system is that if we define G_t by setting $\lambda_t = \sqrt{\frac{\Sigma_t}{G_t}}$, then the system 'decouples':

$$\mathrm{E}[d\sqrt{G_t}] = -\frac{\sigma_t^2}{2\sqrt{G_t}}\, dt, \tag{9}$$

$$\frac{d\Sigma_t}{\Sigma_t} = -\frac{\sigma_t^2}{G_t}\, dt, \tag{10}$$

subject to terminal condition $G_T = 0$ and initial condition Σ_0. Thus, the central new feature of our equilibrium construction is this quantity G_t which solves the 'backward' recursive equation

$$\sqrt{G_t} = \mathrm{E}\left[\int_t^T \frac{\sigma_s^2}{2\sqrt{G_s}}\, ds \,\middle|\, \mathcal{F}_t^\sigma\right]. \tag{11}$$

Using the classic interpretation of price impact as a signal to noise ratio where signal is measured by Σ_t, the market's posterior variance of v, we see that G_t

[4]Since $J_T \geq 0$, we have $J_0 \geq \mathrm{E}[\int_0^T (v - P_s)\theta_s\, ds]$ for all θ and the inequality holds with equality for *any* θ_t^* such that $J_T = 0$, which is thus optimal. The potential non-uniqueness of the optimal strategy is pointed out in Back (1992).

measures the equilibrium noise component. It is a sufficient statistic for the expected amount of future noise trading relevant for the insider to determine how aggressively to trade based on his private information.

We first establish that, under some technical conditions, there exists a unique solution to this recursive equation.

LEMMA 1: (i) *If W_t is a Brownian motion and σ_t is uniformly bounded above by $\overline{\sigma}$ and below by $\underline{\sigma} > 0$, then there exists a bounded solution G_t to the recursive equation* (11). *Further, that solution satisfies*

$$\underline{\sigma}^2(T-t) \le G_t \le \overline{\sigma}^2(T-t). \tag{12}$$

(ii) *If a bounded solution to equation* (11) *exists, then it is unique.*

With these results established, we can now proceed to characterize the equilibrium trading strategy and price process in our economy. The equilibrium we obtain, which constitutes the main result of our paper, is summarized in the following theorem.

THEOREM 1: *If there exists a bounded solution to equation* (11) *and if σ_t is uniformly bounded above by $\overline{\sigma}$ and below by $\underline{\sigma} > 0$, then there exists an equilibrium where the price process has dynamics*

$$dP_t = \kappa_t(v - P_t)\,dt + \sqrt{\Sigma_0}e^{-\int_0^t (1/2)\kappa_s\,ds}\sqrt{\kappa_t}\,dZ_t. \tag{13}$$

The mean-reversion rate is the stochastic process

$$\kappa_t = \frac{\sigma_t^2}{G_t}, \tag{14}$$

where G_t is the unique solution to equation (11). *Note that G_t satisfies*

$$G_t = \mathrm{E}_t\left[\int_t^T \sigma_s^2\,ds\right] - \mathrm{E}_t\left[\int_t^T \Sigma_s\,d\left[\frac{1}{\lambda}\right]_s\right] \le \mathrm{E}_t\left[\int_t^T \sigma_s^2\,ds\right]. \tag{15}$$

The optimal strategy of the insider is

$$\theta_t^* = \frac{\kappa_t}{\lambda_t}(v - P_t). \tag{16}$$

The conditional posterior variance of v in the market maker's filtration, Σ_t, is given by

$$\Sigma_t = \Sigma_0 e^{-\int_0^t \kappa_u\,du}. \tag{17}$$

In equilibrium, price change is linear in order flow (*i.e.*, $dP_t = \lambda_t\, dY_t$) *with a price impact process given by*

$$\lambda_t = \sqrt{\frac{\Sigma_t}{G_t}}. \tag{18}$$

The maximized expected profit of the insider is

$$J_t = \frac{(v - P_t)^2 + \Sigma_t}{2\lambda_t}. \tag{19}$$

The unconditional expected profit of the insider (*from the point of view of the market maker*) *is* $\sqrt{\Sigma_0 G_0}$.

Further, P_t *is a martingale with respect to the market maker's filtration*. *With respect to the insider's filtration*, P_t *is a bridge process that converges in* L^2 *and almost surely to* v *at the final date* T.

Last, *market depth* $\frac{1}{\lambda_t}$ *is a martingale and thus* λ_t *is a submartingale with respect to both the market maker's and the insider's filtrations*.

We now comment on several implications of the theorem. First, the equilibrium price follows a new class of bridge process that converges (almost surely and in L^2) to the value v, known ex ante only to the insider, at maturity T. This guarantees that all private information will have been incorporated into equilibrium prices at maturity and generalizes the result proved in Back (1992) that the equilibrium price in the continuous-time Kyle model follows a standard (i.e., constant volatility) Brownian Bridge. Indeed, in our model, the equilibrium price volatility will be stochastic if the mean-reversion rate κ_t is stochastic. We show below that a necessary condition for this is that the growth rate (m_t) of noise trading volatility is stochastic. In that case, our model generates a volatility of price volatility that appears 'excessive' since shocks to noise trading volatility affect price volatility even if they are unrelated to the fundamental value, because the market maker rationally anticipates that, in periods where noise traders are more active, the insider trades more aggressively.[5]

Second, we find that the optimal trading strategy for the insider is to trade proportionally to the undervaluation of the asset ($v - P_t$) at a rate that is inversely related to her price impact (λ_t) but increasing in the current 'state of liquidity' as measured by how large the current noise trading variance is relative to the expected noise trading variance ($\kappa_t = \frac{\sigma_t^2}{G_t}$). The latter quantity reduces to the inverse of the remaining time horizon ($\frac{1}{T-t}$) in the original Kyle

[5]Of course, since market makers are risk-neutral, total integrated price volatility is always tied to fundamentals in our model as $\int_0^T (dP_t)^2 = \Sigma_0$ a.s.

model when σ_t is constant and, more generally, when σ_t is a martingale. However, the idea that the insider trades at a deterministic rate inversely related to the remaining time horizon $T - t$ does not hold outside these specific cases.

Third, our expression for the price impact generalizes both Kyle's and BP's result that price impact (the inverse of market depth) is a signal to noise ratio. The signal is measured as in the previous papers by the posterior variance of the liquidation value. But interestingly, the relevant measure of noise now solves the recursive equation (11), the solution of which is strictly smaller than the expected remaining variance of noise trading when price impact is stochastic.

Fourth, in our model, market depth is a martingale and thus price impact is a submartingale, expected to increase over time. This contrasts our framework from much of the previous literature. In the original Kyle model, price impact is constant. In extensions of that model Back (1992), Back and Pedersen (1998), Baruch (2002), Back and Baruch (2004), Caldentey and Stacchetti (2010), price impact is either a martingale or a supermartingale. The intuition why price impact is a submartingale in our model is that, with stochastic noise trading volatility, the insider has an option to wait for better liquidity (i.e., higher noise trading volatility) to trade. In equilibrium, price impact must increase on average to entice the insider to trade early and give up his option to wait for better liquidity states.[6]

Fifth, since in equilibrium $d\Sigma_t = -dP_t^2$, a high rate of information arrival implies high stock price volatility. When stock price volatility is stochastic in our model, it is positively correlated (in changes) with noise trading volatility. That is, the rate of price discovery increases when noise trading volatility increases. This is very different from the standard Kyle model, where price volatility is constant and independent of noise trading volatility, and where private information is revealed at a constant rate that is independent of the level of noise trading volatility.[7]

Sixth, in our model, unconditional profits of the informed investor depend on how much private information remains to be released to the market and on the total expected amount of noise trading as measured by the solution

[6]In Kyle (1985), Back (1992), and Back and Pedersen (1998), price impact is constant or a martingale, for else the insider would concentrate all his trading in the period where trading costs are expected to be lowest. In contrast, in Baruch (2002) and Back and Baruch (2004), Kyle's lambda is a supermartingale because the insider has an incentive to trade earlier (because the horizon is random or because he is risk-averse). Foster and Viswanathan (1996) and Back, Cao, and Willard (2000) may also generate an increase in the deterministic price impact, at least near the end of the trading horizon, because of competition among multiple informed traders (we thank a referee for pointing this out).

[7]Note that in Admati and Pfleiderer (1988) or in Back and Pedersen (1998), price volatility is also positively related to the deterministic path of noise trading volatility. We discuss the difference due to the presence of unexpected shocks in noise trading volatility in detail in the next section.

to the recursive equation for G_0, which is strictly less than the total expected noise trader variance if price impact is stochastic (i.e., whenever noise trading volatility is stochastic). It follows that stochastic noise trading volatility *reduces* the unconditional expected profits of the informed trader relative to an economy with the same total average but constant noise trading variance.

Seventh, aggregate execution or slippage costs incurred by uninformed liquidity traders, which can be defined (see Appendix A.4) by:

$$\int_0^T \sigma_t \, dZ_t \, dP_t = \int_0^T \lambda_t \sigma_t^2 \, dt \tag{20}$$

are stochastic and path-dependent in our model, unlike in Kyle and BP, for example, where unconditional expected execution costs always equal realized costs path-wise.

As we show below, when price impact is stochastic, then aggregate execution costs are poorly proxied by the time-average of price impact, which is often used in practice to measure such costs.

We now consider two specific noise trading volatility processes to further emphasize important features of the equilibrium.

3. TRADING VOLUME, PRICE VOLATILITY, AND PRICE IMPACT

As is clear from Theorem 1, much of the characterization of the equilibrium depends on G_t, which affects equilibrium price volatility and price impact. As we show below, a crucial distinction is whether the growth rate (m_t) of noise trading volatility is stochastic. Only in that case can the model generate stochastic price volatility (in addition to stochastic price impact).

3.1. *Deterministic Growth Rate of Noise Trading Volatility*

Suppose that the growth rate (m_t) of the noise trading volatility process in equation (4) is a deterministic process, but that its volatility may be stochastic with general form $\nu(t, \sigma^t)$. Then we can derive an explicit solution for the solution G_t (which does not require σ_t to be uniformly bounded or W_t to be a Brownian motion). With this solution in hand, we can derive all the equilibrium quantities as in Theorem 1 (even though we do not require a uniform upper or lower bound on σ_t). We summarize the results in the following theorem.

THEOREM 2: *If the expected growth rate of noise trading volatility* (m_t) *is deterministic such that*

$$B_t = \int_t^T e^{\int_t^u 2m_s \, ds} \, du \tag{21}$$

is bounded for all $t \in [0, T]$, *then the solution to equation* (11) *is*

$$G_t = \sigma_t^2 B_t. \tag{22}$$

Note that

$$G_t = \int_t^T \mathrm{E}\big[\sigma_s|\sigma^t\big]^2\,ds \le \mathrm{E}\bigg[\int_t^T \sigma_s^2\,ds\bigg|\sigma^t\bigg].$$

In that case, the mean-reversion rate from equation (14) *becomes deterministic* ($\kappa_t = \frac{1}{B_t}$) *and thus from equation* (17), *private information flows into prices deterministically* irrespective *of the magnitude of shocks to noise trading volatility* (ν_t). *Stock price dynamics are given by*

$$dP_t = \frac{1}{B_t}(v - P_t)\,dt + \underbrace{e^{\int_0^t m_s\,ds}\sigma_v}_{\sigma_P(t)}\,dZ_t, \tag{23}$$

where we define $\sigma_v^2 = \frac{\Sigma_0}{B_0}$. *In particular, stock price volatility is a deterministic function which depends only on the unconditional expected path of noise trading volatility*:

$$\sigma_P(t) = \lambda_0 \mathrm{E}\big[\sigma_t|\mathcal{F}_0^\sigma\big]. \tag{24}$$

In equilibrium, price impact is given by

$$\lambda_t = e^{\int_0^t m_s\,ds}\frac{\sigma_v}{\sigma_t}. \tag{25}$$

From equation (16), *the optimal trading strategy of the insider is*

$$\theta_t = \frac{1}{\lambda_t B_t}. \tag{26}$$

The unconditional expected trading rate of the insider is

$$\mathrm{E}[\theta_t|v, \mathcal{F}_0] = e^{\int_0^t 2m_s\,ds}\frac{\sigma_0}{\sigma_v}\frac{(v - P_0)}{B_0}. \tag{27}$$

The unconditional expected profit at time zero of the insider is $\sigma_v\sigma_0 B_0$.

As expected, the model reduces to the continuous-time Kyle model also derived in Back (1992) when $\sigma_t = \sigma$ is constant, in which case $m_t = \nu_t = 0$ and thus $B_t = T - t$. In that case, price volatility is constant equal to σ_v and price impact is constant equal to $\lambda = \frac{\sigma_v}{\sigma}$, where $\sigma_v^2 = \frac{\Sigma_0}{T}$ is the annualized variance of the market maker's prior. Note, in particular, that price volatility is independent of the level of noise trading volatility in the Kyle benchmark.

This theorem informs us about two important special cases: when noise trading volatility is stochastic but unpredictable (i.e., when $m_t = 0$ and $\nu_t \neq 0$) and when it is deterministic but not constant (i.e., $m_t \neq 0$ and $\nu_t = 0$) where we recover the model of Back and Pedersen (BP 1998). We discuss both in turn.

General Unpredictable Martingale Dynamics ($m_t = 0$, $\nu_t \neq 0$)

Comparing the results in Theorem 2 when $m_t = 0$ and thus $B_t = T - t$ to the Kyle benchmark, we see that the equilibrium looks formally identical to the original Kyle model where one substitutes the stochastic process σ_t for the constant noise trading volatility in the original model. This implies that, as in the original Kyle model, price volatility (and the rate at which private information is revealed) is constant and equal to σ_v.

However, both the trading strategy of the insider and the price impact are affected by stochastic noise trading volatility. Price impact is inversely related to noise trading volatility. The insider trades more aggressively when noise trading volatility increases and price impact decreases. Both effects exactly offset to leave equilibrium prices unchanged. In fact, unconditionally, the insider expects to trade continuously a constant number of shares ($\frac{(v-P_0)}{\lambda T}$; see equation (27)) per unit time as in the Kyle model even though noise trading is stochastic. In equilibrium, then, information flows into prices at the same constant rate as in the Kyle model. The expected profit level of the insider equals what it would be in the Kyle model with noise trading volatility set to the constant initial noise trading volatility σ_0.

While it may seem intuitive that, when noise trading volatility is stochastic but unpredictable, the equilibrium is similar to the original Kyle/Back model in some unconditional expected sense, it is interesting to note that it is market depth (and not price impact) that is, on average, constant in that case. That market depth is a martingale follows directly from the fact that it is proportional to noise trading volatility, which is itself a martingale in this case.

Predictable but Deterministic Dynamics ($m_t \neq 0$, $\nu_t = 0$)

When volatility is deterministic but not constant, then price impact is constant and equal to $\frac{\sigma_v}{\sigma_0}$. Price volatility is deterministic and reflects the future realized path of noise trading volatility, which is known as of time zero in this case without uncertainty about future noise trading volatility.

General Case ($m_t \neq 0$, $\nu_t \neq 0$)

In general, when noise trading volatility is stochastic ($\nu \neq 0$) and the growth rate deterministic ($m_t \neq 0$), then the solution combines features from both previous cases. In particular, price impact is stochastic and negatively correlated (in changes) with noise trading volatility. However, price volatility and the posterior variance of the fundamental value (Σ_t) are both deterministic.

This is because the trading strategy of the insider has two components: one due to market depth ($\frac{1}{\lambda_t}$), which is unpredictable, and one due to the current 'liquidity state' ($\kappa_t = \frac{1}{B_t}$), which is deterministic. Price volatility only reflects that second part.[8]

Notice that price volatility is not affected by the realized shocks in noise trading volatility when m_t is deterministic. Only the unconditional ex ante expected path of noise trading volatility matters for the rate at which information flows into prices ex post and consequently for future price volatility. Even though price impact is stochastic, price volatility is deterministic, and the model cannot generate any contemporaneous relation between changes in volume and price volatility or between price impact and price volatility. To generate such relations, we need a stochastic growth rate of noise trading volatility. The next section presents a framework which generates both 'true' stochastic volatility in prices and a meaningful correlation between price volatility and volume.

3.2. *Stochastic Expected Growth Rate of Noise Trading Volatility*

Here we consider a case where noise trading volatility follows a two-state continuous Markov chain. This case is interesting because it introduces state-dependent predictability and thus captures the case of a stochastic expected growth rate in noise trading volatility, which we have shown to be crucial to generate stochastic price volatility.[9] We assume there are two fixed values $\sigma^L < \sigma^H$ with σ_t starting at $\sigma_0 \in \{\sigma^H, \sigma^L\}$ and with dynamics

$$d\sigma_t = \left(\sigma^H - \sigma_t\right) dN_L(t) - \left(\sigma_t - \sigma^L\right) dN_H(t), \tag{28}$$

where $N_i(t)$ is a standard Poisson counting process with jump intensity η_i for $i = H, L$.

Since the volatility process is Markov, we expect that the solution to the recursive equation (11) will be of the form $G(t, \sigma_t)$. Indeed, we find the following:

THEOREM 3: *The unique bounded solution to the recursive equation* (11) *is given by* $G(t, \sigma_t) = \mathbf{1}_{\{\sigma_t=\sigma^L\}} G^L(T-t) + \mathbf{1}_{\{\sigma_t=\sigma^H\}} G^H(T-t)$, *where the deterministic functions* G^L, G^H *satisfy the system of ODE given in equations* (29)–(30)

[8]In the Supplemental Material (Collin-Dufresne and Fos (2016b, Appendix S1)), we discuss the optimal strategy of the insider in further detail.

[9]In the Supplemental Material, Appendix S2, we also analyze a diffusion case with stochastic growth rate, where specifically we assume that noise trading volatility follows a mean-reverting diffusion process. Using an expansion approximation to the solution for G, we show how uncertainty about future noise trading volatility interacts with mean-reversion to generate stochastic price volatility.

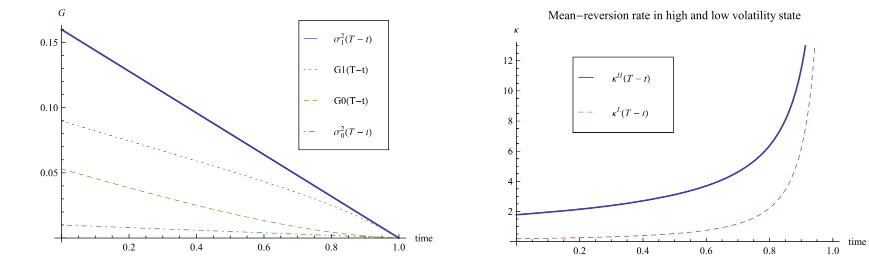

(a) $G^i(T-t)$ functions solving the system (29)–(30) along with the lines $(\sigma^i)^2(T-t)$ for $i = H, L$.

(b) Mean-reversion rates $\kappa^i(T-t)$ for $i = H, L$ from equation (32).

FIGURE 1.—Markov switching model for noise trading volatility. Parameter values are: $\Sigma_0 = 0.04$, $T = 1$, $\eta_L = \eta_H = 2$, $\sigma^L = 0.2$, and $\sigma^H = 0.4$.

below, with boundary conditions $G^L(0) = G^H(0) = 0$:

$$G^L_\tau(\tau) = \left(\sigma^L\right)^2 + 2\eta_L\left(\sqrt{G^H(\tau)G^L(\tau)} - G^L(\tau)\right), \tag{29}$$

$$G^H_\tau(\tau) = \left(\sigma^H\right)^2 + 2\eta_H\left(\sqrt{G^H(\tau)G^L(\tau)} - G^H(\tau)\right). \tag{30}$$

When there are no transitions between states $\eta_i = 0$, then the solution is $G^i(\tau) = (\sigma^i)^2(T-t)$ for $i = L, H$. When $\eta_i \neq 0$, the system of coupled differential equations for $G^i(\tau)$ for $i = L, H$ is easily solved numerically. For illustration, we choose a period length $T = 1$, $\eta_L = \eta_H = 2$, $\sigma^L = 0.2$, and $\sigma^H = 0.4$. For these parameter values, we report in Figure 1 the pair of functions (G^H, G^L). As expected, $G^H(t) \geq G^L(t)$, and as maturity approaches, the two functions converge smoothly to the lines $(\sigma^i)^2(T-t)$ for $i = L, H$ that would prevail if there were no transitions between states (i.e., the state was absorbing) and which also correspond to the original Kyle model with noise trading volatility fixed at either the high or low level.

From Theorem 1, we can then write price dynamics as

$$dP_t = \kappa(t, \sigma_t)(v - P_t)\,dt + \underbrace{\sqrt{\Sigma_0}e^{-\int_0^t (1/2)\kappa(s,\sigma_s)\,ds}\sqrt{\kappa(t, \sigma_t)}}_{\sigma_P(t)}\,dZ_t, \tag{31}$$

where, from its definition in equation (14), we obtain the mean-reversion rate process

$$\kappa(t, \sigma_t) = \mathbf{1}_{\{\sigma_t = \sigma^L\}}\kappa^L(T-t) + \mathbf{1}_{\{\sigma_t = \sigma^H\}}\kappa^H(T-t). \tag{32}$$

Price follows a mean-reverting process with stochastic volatility, where both the strength of mean-reversion and the volatility are modulated by the volatility state. Figure 1 shows the values of $\kappa^i(T-t) = \frac{(\sigma^i)^2}{G^i(T-t)}$ for $i = L, H$. We see that mean-reversion is always higher in the high noise trading volatility state, where we expect the insider to trade more aggressively and thus to contribute more to price discovery. Further, both mean-reversion coefficients increase to infinity as we approach maturity. Since maturity is fixed and finite, agents trade more and more aggressively as it approaches because they do not want to leave any money on the table. Price volatility dynamics can be calculated explictly from its definition in equation (31):

$$\begin{aligned}(33)\quad &\frac{d\sigma_P(t)}{\sigma_P(t)}\\ &\quad = \frac{1}{2}\left(\frac{\partial \log \kappa(t,\sigma_t)}{\partial t} - \kappa(t,\sigma_t)\right)dt + \left(\sqrt{\frac{\kappa^H(T-t)}{\kappa^L(T-t)}} - 1\right)dN_L(t)\\ &\qquad + \left(\sqrt{\frac{\kappa^L(T-t)}{\kappa^H(T-t)}} - 1\right)dN_H(t).\end{aligned}$$

Clearly, price volatility always jumps up (down) when noise trading volatility jumps up (down), which generates a positive relation between changes in volume and volatility. However, conditional on being in a noise trading regime, there are two counterbalancing effects driving the change in volatility. First, informed trading reveals information and decreases the amount of remaining uncertainty, which reduces price volatility. This is the negative term $(-\kappa(t, \sigma_t))$ in the drift of σ_P, which is always more negative in the high volatility state (where more information is revealed) than in the low state. Second, due to the finite maturity effect discussed above, the insider trades more aggressively irrespective of the state as maturity approaches, which tends to increase volatility. This is the time derivative $(\frac{\partial \log \kappa(t,\sigma_t)}{\partial t})$ in the drift of σ_P, which is always positive. Which effect dominates quantitatively depends on the parameters.

For illustration, we plot in Figure 2(a) four paths of the stock price volatility $\sigma_P(t)$ for the case where noise trading volatility switches to the high regime at date zero and stays there until maturity (high/high), when it starts in the low regime and stays there until maturity (low/low), and when there is a jump at $t = 0.5$ from high to low (high/low) and from low to high (low/high), respectively. We also plot the stock price volatility in the Kyle benchmarks with constant high or low noise trading volatility.

Consistent with our previous discussion, price volatility starts out higher in the high noise trading volatility state than in the low state. Subsequently, however, volatility decreases in the high regime, but *increases* in the low volatility regime, which shows that the finite maturity effect discussed above dominates in the low regime. When noise trading volatility jumps up (down), price volatility also jumps up (down), generating the positive volume volatility correlation in changes. However, volume and volatility need not be positively correlated

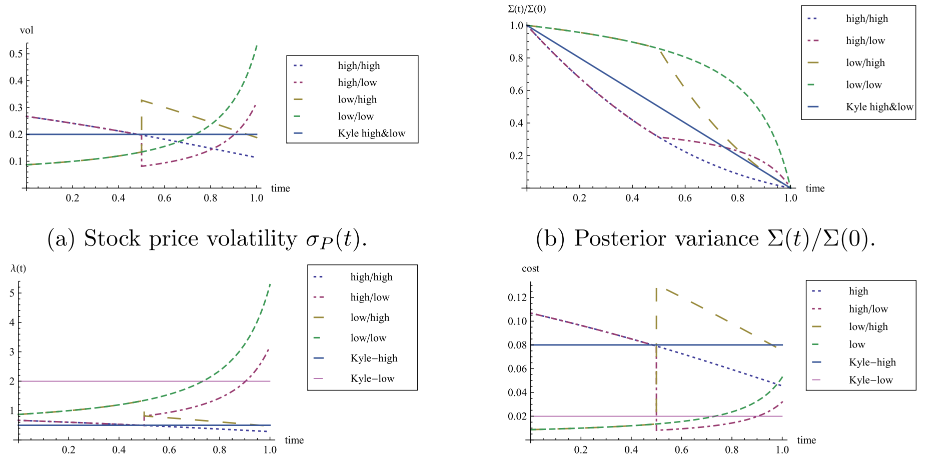

(a) Stock price volatility $\sigma_P(t)$.

(b) Posterior variance $\Sigma(t)/\Sigma(0)$.

(c) Price impact (λ_t).

(d) Realized execution costs $(\lambda_t \sigma_t^2)$.

FIGURE 2.—We graph price volatility, posterior variance, price impact, and execution costs in the Markov chain model for noise trading volatility for four separate paths of noise trading volatility corresponding to (1) start and stay in the high noise volatility regime until T, (2) start and stay in the low volatility regime until T, (3) start in the high volatility regime and switch to low volatility at $t = 0.5$, and (4) start in the low volatility regime and switch to high volatility at $t = 0.5$. We also plot these for the Kyle-benchmark models with high or low (but constant) noise trading volatility. Parameter values are: $\Sigma_0 = 0.2^2$, $T = 1$, $\eta_L = \eta_H = 2$, $\sigma^L = 0.2$, and $\sigma^H = 0.4$.

in levels due to the finite horizon effect discussed above. For example, Figure 2(a) shows that price volatility becomes highest if the economy reaches maturity having remained all the time in the low state (low/low path). Intuitively, this is a path where the insider has to trade on most of his information just prior to market close in a thin market, after having waited unsuccessfully for better liquidity states. So the market closes with low volume and high price volatility.

The figure clearly illustrates that price volatility can exceed very significantly the price volatility of the Kyle model, which is always identical to σ_v, independent of the level of noise trader volatility. Instead, in our model, price volatility is stochastic and 'excessive,' in that it is driven by stochastic switches in noise trading volatility regimes, which are by assumption unrelated to the (constant) fundamental.

In this model, price impact is also stochastic and negatively related (in changes) to noise trading volatility. Indeed, recall that $\lambda_t = \sqrt{\frac{\Sigma_t}{G_t}}$ and thus, since $\Sigma_t = \Sigma_0 e^{-\int_0^t \kappa(u,\sigma_u)\,du}$ is absolutely continuous, the immediate effect of an upward jump from G^L to G^H is to lower price impact. Subsequently, however, the path of price impact depends on the relative speed at which the posterior variance Σ_t and the G_t function decrease, which depends on the path of noise trading volatility.

Figure 2(b) plots the path of the posterior variance Σ_t for the four paths of noise trading volatility considered previously and against the Kyle benchmark. Note that, in Kyle's model, $\frac{\Sigma_t}{\Sigma_0} = \frac{T-t}{T}$, that is, information always decays linearly in time, *irrespective* of the level of noise trading volatility. Instead, with stochastic noise trading volatility, posterior variance typically drops faster when noise trading volatility is higher. In fact, the figure shows the posterior variance is a decreasing convex function of time in the high volatility regime, but becomes decreasing concave when there is a switch to the low noise trading regime. The intuition is that, in the low noise trading regime, the insider is playing a waiting game, in that he trades much less aggressively than he would in the Kyle economy with the same level of volatility. He does so hoping for the high noise trading regime to arrive, when he trades more aggressively, leading to faster information revelation. Of course, if the regime switch does not arrive, then ultimately, he will have to become more aggressive so that all his information eventually makes it into prices (see the path marked as 'low/low' on the graph). The switch from convex to concave is due to the finite horizon effect discussed above.[10]

In Figure 2(c), we plot the path of the price impact ($\lambda(t)$) process for the same four noise trading volatility scenarios. We see that if the economy starts in the high noise trading regime and stays there until maturity, then measured price impact is relatively low and decays steadily. Instead, if the economy starts in the low noise trading regime, then price impact is at first only slightly higher than in the high noise trading regime, but it increases exponentially as the economy approaches maturity without switching. Similarly, if the regime switches at some point from high to low volatility, then price impact immediately jumps up and then starts to increase along a very convex path as markets become more illiquid. Interestingly, note that if the economy is in the high noise trading regime, then measured price impact starts low and decreases steadily at the beginning, even though there is a lot of 'adverse selection' in the sense that a lot of information is getting into prices as shown in Figure 2(b).

Last, in Figure 2(d), we plot the path of realized execution costs ($\lambda_t \sigma_t^2$) for the same four scenarios. The total execution costs paid by noise traders at time T are captured by the area below each curve plotted. From the graph, it is clear that execution costs are lowest in the low volatility regime scenario, and much higher in the high noise trading volatility regime. We give the corresponding numbers in Table I. This may seem paradoxical, since, as is clear from the table, the high noise trading volatility regime is also the one where the average price impact (λ) is lower. However, the comparison is not appropriate since there is

[10] In a model where T is modeled as an unpredictable stopping time with constant arrival intensity, Σ_t always follows a decreasing convex path, with a faster decay rate in the high volatility state (see Collin-Dufresne and Fos (2016a) for further discussion and extension to random intensity).

TABLE I

PATH-DEPENDENCE IN EXECUTION COSTS: AN ILLUSTRATION[a]

	Noise Trading Volatility Paths			
	High/High (1)	Low/Low (2)	High/Low (3)	Low/High (4)
Panel A: Aggregate execution costs				
Total	0.078	0.017	0.054	0.057
Path Dependent	0.047/0.031	0.005/0.012	0.047/0.007	0.005/0.052
Panel B: 'Number' of noise traders				
Total	0.16	0.01	0.085	0.085
Path Dependent	0.08/0.08	0.005/0.005	0.08/0.005	0.005/0.08
Panel C: Normalized aggregate execution costs				
Total	0.487	1.740	0.636	0.671
Path Dependent	0.587/0.387	1/2.4	0.587/1.4	1/0.65
Panel D: Average price impact				
Total	0.487	1.740	1.023	0.853
Path Dependent	0.584/0.39	1.06/2.42	0.584/1.462	1.06/0.646
Panel E: Average stock price volatility				
Total	0.195	0.174	0.190	0.182
Path Dependent	0.234/0.156	0.106/0.242	0.234/0.146	0.106/0.258

[a]This table presents several equilibrium quantities depending on various scenarios of realized paths of noise trading volatility. Each path of realized noise trading volatility corresponds to a certain 'number' of uniformed traders arriving to the market. This 'number' is measured by the quadratic variation of the order flow. In column (1), the realized noise trading volatility is always in the high regime. In column (2), the realized noise trading volatility is always in the low regime. In column (3), the realized noise trading volatility changes from high to low regime in the middle of the trading period. In column (4), the realized noise trading volatility changes from low to high regime in the middle of the trading period. Panel A reports the realized aggregate execution costs for noise traders, $\int_0^T \lambda_t \sigma_t^2 \, dt$. Panel B reports the 'number' of noise traders, $\int_0^T \sigma_t^2 \, dt$. Panel C reports the normalized realized execution costs for noise traders, $\int_0^T \lambda_t \sigma_t^2 \, dt / \int_0^T \sigma_t^2 \, dt$. Panel D reports the average price impact, $\frac{1}{T}\int_0^T \lambda_t \, dt$. Panel E reports the average stock price volatility, $\frac{1}{T}\int_0^T \lambda_t \sigma_t \, dt$. The corresponding Kyle/Back stock volatility is constant and equal to 0.2 independent of the level of noise trading volatility.

more noise trading (as measured by the quadratic variation of the order flow) in a high volatility scenario than in the low volatility scenario and, therefore, it is natural that the aggregate execution costs paid by noise traders are higher in the high volatility scenario. However, if we compare the two other scenarios (high/low to low/high), where there are the same 'number' of noise traders along each path (in the sense that the cumulative quadratic variation of noise trader order flow is the same as is confirmed in the third row of Table I), then we see that this is not the whole story. Indeed, aggregate execution costs paid by noise traders are higher even though average price impact is lower in the low/high scenario relative to the high/low scenario. This shows that a simple average of price impact does not capture the actual level of execution costs. Instead, let us define the volume weighted average (i.e., 'normalized') execution costs to noise traders:

$$\text{(34)} \qquad \textit{Normalized Execution Costs} = \frac{\int_0^T \lambda_t \sigma_t^2 \, dt}{\int_0^T \sigma_t^2 \, dt}.$$

The normalized execution costs reduce to the simple time average of λ_t if either λ_t or σ_t is constant. This is the case when noise trading volatility is deterministic Kyle (1985), Back and Pedersen (1998), in which case normalized execution costs are constant. Instead, with stochastic noise trading volatility, they depend on the path followed by noise trading volatility. Panel C shows that normalized execution costs better capture the actual execution costs paid by the average noise trader.

4. EXTENSIONS AND FURTHER DISCUSSION

4.1. *The Dynamics of Aggregate Order Flow*

For simplicity, we have assumed that aggregate order flow and noise trading volatility are conditionally uncorrelated. This assumption can be relaxed. Consider, for example, the more general model for total order flow $\hat{Y}_t$ given by

$$\text{(35)} \qquad d\hat{Y}_t = \theta_t \, dt + \sigma_t \, dZ_t + \eta(t, \sigma^t, \hat{Y}^t) \, dW_t,$$

where we leave the dynamics of σ_t unchanged, that is, as in equation (4). In that case, to solve the equilibrium it is useful to define the process Y_t by

$$\text{(36)} \qquad dY_t = d\hat{Y}_t - \frac{\eta(t, \sigma^t, \hat{Y}^t)}{\nu(t, \sigma^t)} \big(d\sigma_t - m(t, \sigma^t) \, dt\big) = \theta_t \, dt + \sigma_t \, dZ_t.$$

Note that the dynamics of Y_t in equation (36) are identical to those in equation (2). Further, since observing $(\hat{Y}_t, \sigma_t)$ is equivalent to observing (Y_t, σ_t) for all market participants, it can be shown that all our results above (and in particular Theorem 1) are unchanged with this more general model of order flow. In particular, the price process is given by $dP_t = \lambda_t \, dY_t$ as in Theorem 1, even though total order flow is given by $\hat{Y}_t$. This shows that our previous assumption that total order flow and σ_t be conditionally uncorrelated is not crucial for our results. Instead, our equilibrium proof does rely on the assumption that order flow does not Granger-cause σ_t in the sense that the history of $\hat{Y}_t$ does not affect the future dynamics of σ_t (i.e., m and ν cannot depend on $\hat{Y}^t$ for the solution to the equation for G_t to be independent of total order flow and of the strategy of the insider).

Economically, this extension is also interesting, as it shows that price change need not be linear in total order flow, but only in the component of the order flow that is informative about the insider's actions. So, for example, in this more general model, a regression of price changes on total order flow (which is typically used to estimate Kyle's lambda in empirical papers) delivers an estimate of price impact 'regression coefficient': $\hat{\lambda}_t = \frac{dP_t\, d\hat{Y}_t}{d\hat{Y}_t^2} = \lambda_t \frac{\sigma_t^2}{\sigma_t^2 + \eta(t)^2}$, which equals λ_t only if $\eta(t) = 0$. Using this estimate will not provide an accurate estimate of actual trading costs paid by individual uninformed investors if $\eta(t) \neq 0$. This is because uninformed investors' trades come in two groups in this extended model: the group who can be used by the informed investor to hide his trades ($\sigma_t\, dZ_t$) and who therefore generate a per-trade slippage cost of λ_t, and the group ($\eta(t)\, dW_t$) who is known to be pure noise and whose demand does not generate any slippage. Indeed, if the market maker could distinguish between the two types of uninformed order flow, then they would each pay different trading costs.

4.2. *Informed Trading and Adverse Selection Measures*

Empirical measures of adverse selection typically rely on an estimate of the persistent component of the price impact of trades to measure the amount of private information in trades. In their well-known survey of the micro-structure literature, Biais, Glosten, and Spatt (2005) described the empirical relation between adverse selection and the price impact (λ) as follows: "As the informational motivation of trades becomes relatively more important, λ goes up." (page 232). Consistent with this intuition, many empirical studies rely on simple averages of price impact to sort firms into groups with different levels of adverse selection.

One implication of our model is that, while at an individual trade level price impact does represent the transaction cost paid by a noise trader, over any finite period average price impact may not be a valid measure of aggregate (or average) adverse selection costs when both σ_t and λ_t change over time. This is most evident in the Markov chain example. In columns (1) and (2) of Table I, the noise trading volatility is constant and therefore the average price impact is equal to the normalized aggregate execution cost paid by noise traders. In contrast, in columns (3) and (4), the noise trading volatility and price impact change over time, resulting in a significant difference between the average price impact and the normalized execution cost paid by the average noise traders. When σ_t and thus λ_t change over time, it is better to use 'volume' weighted price impact as a measure of aggregate adverse selection costs, as it takes into account the variation in the trading by uninformed investors and the endogenous response of the informed investor's trading.

In a recent empirical study, Collin-Dufresne and Fos (2015) investigated a large sample of trades by informed investors[11] and found that these trade much more aggressively, when measured adverse selection is low. Their study uncovers a strong negative relation between traditional measures of adverse selection (such as estimates of Kyle's lambda obtained from high-frequency data) and trading by informed investors. They also showed that these informed investors are more likely to trade when abnormal volume on the stock itself as well as other measures of market-wide liquidity, such as the abnormal volume on the S&P 500 stock index, are high, which is consistent with the economic mechanism of this paper where informed investors wait for better liquidity (as measured by $\kappa_t = \frac{\sigma_t^2}{G_t}$) to trade more. This also calls into question the widespread use of (high- or low-frequency) price-impact measures as measures of adverse selection, for example, in cross-sectional empirical asset pricing tests.

4.3. *Dynamics of Price Volatility, Bid-Ask Spreads, and Returns*

The model makes interesting predictions about the joint dynamics of price volatility, price impact, and stock returns. For example, the model predicts that typically price volatility tends to be high when volume is high and price impact is low. This goes in an opposite direction to the relation predicted, for example, by inventory models of trading costs, where higher volatility would typically be positively related to trading costs (Stoll (1978)). Further, our model predicts that the joint dynamics are path-dependent and will depend significantly on the realized path of uninformed trading. So, for example, after a long period of low noise trading, price volatility and price impact can actually both rise together. This occurs in the model if the insider approaches maturity without having been able to trade much (see Figures 2(c) and 2(a)). Eventually, he is not willing to wait anymore and trades aggressively in a low noise trading environment, leading to high price impact and high volatility. These predictions could be tested and used to improve estimates of the adverse selection component of trading costs. They could help better understand the complex joint dynamics of volatility, return, and volume often observed around daily market closes.

4.4. *Subordinated Processes and Time Changes*

There is a long tradition in financial econometrics to model the time series of price changes as subordinated to the normal distribution. Clark (1973) initiated this literature and proposed that the directing process τ_t be a form of

[11]Exploiting an SEC disclosure requirement that requires activist shareholders to file a 13D schedule in which they report past trades in target stocks when they hit the 5% ownership threshold, CF built a sample of trades by activist investors. They documented that these trades are informed, based on their abnormal realized profits, and analyzed the price impact of these trades.

'business time' measure related to the volume of trading. This idea led to several reduced-form models trying to capture the volatility-volume relationship (see Tauchen and Pitts (1983), Epps and Epps (1976), Richardson and Smith (1993), Andersen (1996), Gallant, Rossi, and Tauchen (1992)), in particular to capture the heteroscedasticity in returns. Our model provides microeconomic foundations for such a subordinate process for stock return. Indeed, the directing process is endogenous in our model and related to trading volume, as we point out in the following corollary:

COROLLARY 1: *Define the positive increasing stochastic (directing) process*:

$$\tau_t = T\left(1 - e^{-\int_0^t \sigma_u^2/G_u\,du}\right). \tag{37}$$

Then, setting $\sigma_v^2 = \frac{\Sigma_0}{T}$ *and since* $\tau_T = T$, *we have*

$$\Sigma_t = \sigma_v^2(\tau_T - \tau_t) \tag{38}$$

and

$$dP_t = \frac{(v - P_t)}{\tau_T - \tau_t}\,d\tau_t + \sigma_v\,dB_{\tau_t} \tag{39}$$

for some Brownian motion B independent of W defined by $\sigma_v\,dB_{\tau_t} = \lambda_t\sigma_t\,dZ_t$. It follows that the time-changed price process $\hat{P}(\cdot)$ defined by $\hat{P}(\tau_t) = P_t$ is a Brownian Bridge on the time-changed insider's $\mathcal{F}_{\tau_t}$-filtration and a Brownian motion on the market maker's $\mathcal{F}^Y_{\tau_t}$-filtration.

The corollary shows that the equilibrium price is a time-changed Brownian Bridge, which is reminiscent of the Kyle–Back model where the price process is a standard (constant volatility) Brownian Bridge. However, our equilibrium *cannot* simply be obtained as a time-change of that model, since price impact is constant in the Kyle–Back model and a stochastic process in ours. In the market maker's filtration, price is a time-changed Brownian motion, which belongs to the class of subordinate processes proposed in Clark (1973). Interestingly, our model gives an endogenous expression for the directing process (τ_t) that depends on the (uninformed) volume dynamics and does not require the specification of a latent 'information process' to generate stochastic volatility (see the discussion in Andersen (1996), e.g.). The model is general in that we can solve for the directing process for *any* dynamics of volume (the literature has considered various distributions such as Normal, Poisson, and Log-normal; e.g., Andersen (1996) and Richardson and Smith (1993)). Thus, our theoretical model has the potential to "jointly account for major stylized facts—serially correlated volatility, contemporaneous volume-volatility correlation and excess kurtosis of price changes" as emphasized by Gallant, Rossi, and Tauchen (1992).

5. CONCLUSION

In this paper, we have extended Kyle's (1985) model of dynamic insider trading to the case where noise trading volatility can change stochastically over time. Under certain conditions which we identify, the equilibrium price process exhibits stochastic 'excess volatility' in the sense that non-payoff-relevant shocks that change noise trading in equilibrium also drive the price volatility. This is because rational market makers anticipate that more informed trading occurs, and thus more information is revealed, when noise trading volatility is high. As a result, price impact is stochastic and negatively correlated with noise trading volatility. Further, in equilibrium, price impact is a submartingale, indicating that, on average, execution costs are expected to increase over time, reflecting the 'liquidity timing option' held by the insider.

The model makes interesting predictions about the joint dynamics of price volatility, price impact, and volume, which could be taken into account to empirically measure the adverse selection component of trading costs.

The model makes many simplifying assumptions that could be relaxed to further our understanding of how information flows into prices and how price volatility, price impact, and trading volume comove. First, we assume that the amount of private information is fixed and only noise trading volatility is time varying. Second, we assume that the horizon is fixed. Third, we assume throughout that the noise trading volatility process is observable to all. Fourth, we assume that the presence of the insider is common knowledge. And last, we assume that the insider and market makers are risk-neutral. We leave these extensions for future research.

APPENDIX

A.1. *Proof of Lemma 1: Existence*

We note that $y_t = \sqrt{G_t}$ solves the backward stochastic differential equation

$$dy_t = -f(t, y_t)\,dt - \Lambda_t\,dM_t,$$

with $f(t, y) = \frac{\sigma_t^2}{2y}$ and with terminal condition $y_T = 0$. Now $f(t, y_t) \leq \ell(y_t)$ $\forall(t, \omega)$ where we define the function $\ell(y) = \frac{\overline{\sigma}^2}{2|y|}$. We note that $\ell(y)$ is continuous and strictly positive and that $\int_0^\infty \frac{dx}{\ell(x)} = \int_{-\infty}^0 \frac{dx}{\ell(x)} = \infty$. Thus $\ell(x)$ is super-linear as shown in Lemma 1 of Lepeltier and San Martin (1997). Their Theorem 1 then applies, which gives us the existence of a maximal bounded solution for y_t (and therefore for G_t). In addition their Theorem 1 implies that there exist two solutions $L(t), U(t)$ that solve $L_t = -\int_t^T \ell(L_s)\,ds$ and $U_t = \int_t^T \ell(U_s)\,ds$ such that we have $L_t \leq G_t \leq U_t$. It is easy to calculate that $U_t^2 = -L_t^2 = \overline{\sigma}^2(T - t)$. This gives us the upper bound.

For the lower bound, we use a comparison theorem. Consider the solution to the following backward equation:

$$dx_t = -\frac{\underline{\sigma}^2}{2x_t}\,dt - \tilde{\Lambda}_t\,dW_t,$$

with terminal condition $x_T = 0$. It can be computed straightforwardly as $x_t = \underline{\sigma}\sqrt{T-t}$ (note $\tilde{\Lambda}_t = 0$). Since $\forall(t,\omega)\ f(t,y) \geq \frac{\underline{\sigma}^2}{2y}$, we can use the comparison result Corollary 2 of Lepeltier and San Martin (1997) to obtain

(40) $\qquad y_t \geq x_t \quad \forall(t,\omega),$

which gives the lower bound on the maximal solution for G_t.

A.2. *Proof of Lemma 1: Uniqueness*

Define $g_t = \sqrt{G_t}$. To prove uniqueness of the solution, assume that there are two (uniformly bounded) solutions g_t^1 and g_t^2 to the recursive equation. Then consider the difference $\Delta_t = g_t^1 - g_t^2$. It satisfies

$$\Delta_t = \mathrm{E}\left[\int_t^T -\Delta_s \frac{\sigma_s^2}{2g_s^1 g_s^2}\,ds \Big| \mathcal{F}_t^\sigma\right].$$

Thus, if we define $a_s = \frac{\sigma_s^2}{2\sqrt{g_s^1 g_s^2}} \geq 0$, we have that $e^{-\int_0^t a_s\,ds}\Delta_t$ is a bounded continuous martingale (note that Δ_t is clearly bounded since it is the difference of two uniformly bounded positive processes) equal to zero at T (since $g_T^1 = g_T^2 = 0$). It thus follows that $\Delta_t = 0\ \forall t$.

A.3. *Proof of Main Theorem 1*

The proof is in several steps.

A.3.1. *Step 1: Market Maker's Updating*

First, we establish that if the market maker conjectures that the insider's trading strategy is linear in his per-period profit, that is, that

(41) $\qquad \theta_t = \beta_t(v - P_t),$

where β_t is an $\mathcal{F}_t^Y$-adapted process that measures the speed at which the insider decides to close the gap between the fundamental value v (known only to him) and the market price P_t, and where we define Σ_t as the conditional variance of the terminal payoff:

(42) $\qquad \Sigma_t = \mathrm{E}\left[(v - P_t)^2 | \mathcal{F}_t^Y\right],$

then the equilibrium price process which results from the market maker's break-even pricing rule given in equation (3) is such that price changes are conditionally linear in order flow.

LEMMA 2: *If the insider adopts a trading strategy of the form given in* (41), *then the stock price given by equation* (3) *starts at* P_0 *and has dynamics*

$$dP_t = \lambda_t \, dY_t, \tag{43}$$

where the price impact is a function of the conjectured trading rule:

$$\lambda_t = \frac{\beta_t \Sigma_t}{\sigma_t^2}. \tag{44}$$

Further, *the dynamics of the posterior variance are given by*

$$d\Sigma_t = -\lambda_t^2 \sigma_t^2 \, dt. \tag{45}$$

PROOF: This follows directly from an application of Theorems 12.6, 12.7 in LS (2001). We provide a simple 'heuristic' motivation of the result using standard Gaussian projection theorem below:

$$P_{t+dt} = \mathrm{E}\big[v|Y^t, Y_{t+dt}, \sigma^t, \sigma_{t+dt}\big] \tag{46}$$

$$= \mathrm{E}\big[v|Y^t, \sigma^t\big] + \frac{\mathrm{Cov}\big(v, Y_{t+dt} - Y_t|Y^t, \sigma^t\big)}{V\big(Y_{t+dt} - Y_t|Y^t, \sigma^t\big)} \times \big(Y_{t+dt} - Y_t - \mathrm{E}\big[Y_{t+dt} - Y_t|Y^t, \sigma^t\big]\big) \tag{47}$$

$$= P_t + \frac{\beta_t \Sigma_t \, dt}{\beta_t^2 \Sigma_t \, dt^2 + \sigma_t^2 \, dt}(Y_{t+dt} - Y_t) \tag{48}$$

$$\approx P_t + \frac{\beta_t \Sigma_t}{\sigma_t^2} \, dY_t. \tag{49}$$

The second equality uses the fact that the dynamics of σ_t is independent of the asset value distribution and of the innovation in order flow. The third equality uses the fact that the expected change in order flow is zero for the conjectured policy. The last line follows from going to the continuous-time limit (with $dt^2 \approx 0$). Similarly, by the projection theorem, we have

$$\mathrm{Var}\big[v|Y^t, Y_{t+dt}, \sigma^t, \sigma_{t+dt}\big] = \mathrm{Var}\big[v|Y^t, \sigma^t\big] - \left(\frac{\beta_t \Sigma_t}{\sigma_t}\right)^2 \mathrm{Var}\big[Y_{t+dt} - Y_t|Y^t, \sigma^t\big], \tag{50}$$

which gives

(51) $$\Sigma_{t+dt} = \Sigma_t - \lambda_t^2 \sigma_t^2 \, dt. \qquad \textit{Q.E.D.}$$

A.3.2. *Insider's Optimal Strategy*

Second, we establish that if price changes are linear in order flow with a specific choice of price impact process, namely:

(52) $$dP_t = \lambda_t \, dY_t,$$

(53) $$\lambda_t = \sqrt{\frac{\Sigma_t}{G_t}},$$

with G_t, Σ_t as defined in (11) and (45), then the optimal trading strategy of the insider is indeed of the form given in equation (41).

To establish this, we first need a preliminary result which establishes that the conjectured equilibrium price process converges at maturity to the liquidation value v.

LEMMA 3: *Suppose price dynamics are given by equations* (52), (45), (53), *and* (11); *then the price process P_t converges almost surely to v at time T.*

PROOF: The conjectured equilibrium price process is

(54) $$dP_t = \frac{(v - P_t)}{G_t} \sigma_t^2 \, dt + \sqrt{\frac{\Sigma_t}{G_t}} \sigma_t \, dZ_t,$$

(55) $$d\Sigma_t = -\frac{\Sigma_t}{G_t} \sigma_t^2 \, dt.$$

It is straightforward to solve the ODE for Σ_t and obtain equation (17). Consider the process $X(t) = P_t - v$:

(56) $$X(t) = e^{-\int_0^t (\sigma_u^2 / G_u) \, du} X_0 + \int_0^t e^{-\int_s^t (\sigma_u^2 / G_u) \, du} \sqrt{\frac{\Sigma_s}{G_s}} \sigma_s \, dZ_s$$

(57) $$:= I_1(t) + I_2(t),$$

where the second line defines the integrals I_1, I_2. Equation (12) implies that

$$\frac{\overline{\sigma}^2}{\underline{\sigma}^2} \log\left(\frac{T}{T-t}\right) \geq \int_0^t \frac{\sigma_u^2}{G_u} \, du \geq \frac{\underline{\sigma}^2}{\overline{\sigma}^2} \log\left(\frac{T}{T-t}\right).$$

It follows immediately from this inequality that

(58) $$\lim_{t \to T} I_1(t) = 0 \quad \text{a.s.}$$

Further, note that $I_2(t) = e^{-\int_0^t (\sigma_u^2/G_u)\,du} M_t$, where we define the Brownian martingale:

$$M_t = \int_0^t e^{\int_0^s (\sigma_u^2/G_u)\,du} \sqrt{\frac{\Sigma_s}{G_s}} \sigma_s \, dZ_s. \tag{59}$$

Note that the quadratic variation of M_t is equal to

$$\langle M \rangle_t = \int_0^t e^{\int_0^s (2\sigma_u^2/G_u)\,du} \frac{\Sigma_s}{G_s} \sigma_s^2 \, ds \tag{60}$$

$$= \Sigma_0 \left(e^{\int_0^t (\sigma_u^2/G_u)\,du} - 1 \right), \tag{61}$$

where we substituted Σ_s from equation (17) to obtain the second line.

Now, from Karatzas and Shreve (1991, Theorem 4.6, p. 174), we know there exists a standard Brownian motion B_t such that the continuous martingale can be seen as a time-changed Brownian motion, specifically $M_t = B_{\langle M \rangle_t}$. Using the strong law of large numbers for Brownian motion, which states that $\lim_{\tau\to\infty} B_\tau/\tau = 0$ a.s. (see Karatzas and Shreve (1991, p. 104)), we obtain

$$\lim_{t\to T} e^{-\int_0^t (\sigma_u^2/G_u)\,du} M_t = \lim_{t\to T} \frac{B_{\langle M \rangle_t}}{1 + \frac{\langle M \rangle_t}{\Sigma_0}} = \lim_{\tau\to\infty} \frac{B_\tau/\tau}{\frac{1}{\Sigma_0} + 1/\tau} = 0 \quad \text{a.s.} \tag{62}$$

This establishes that $\lim_{t\to T} I_2(t) = 0$ a.s. and completes the proof. *Q.E.D.*

We now establish another useful result about the limiting distribution of the standardized price process.

LEMMA 4: *The process $h_t = \frac{P_t - v}{\sqrt{\Sigma_t}}$ follows a time-changed Ornstein–Uhlenbeck process with the property that h_T has a normal distribution with $E[h_T] = 0$ and $\mathrm{E}[h_T^2] = 1$. It follows that P_t converges to v in L^2.*

PROOF: Simple calculations show that

$$dh_t = -\frac{1}{2} \frac{\sigma_t^2}{G_t} h_t \, dt + \frac{\sigma_t}{\sqrt{G_t}} \, dZ_t. \tag{63}$$

This is a time-changed Ornstein–Uhlenbeck process with stochastic time-change process $\tau_t = \int_0^t \frac{\sigma_s^2}{G_s} \, ds$, which is independent of the filtration generated by Z_t. Straightforward calculations show that $E[h_T] = 0$ and $E[h_T^2] = 1$ and that the limiting distribution of h_T is a standard normal. *Q.E.D.*

Next, we establish that market depth is a martingale.

LEMMA 5: *Market depth (which is the inverse of the price impact, i.e., Kyle's lambda) is a martingale that is orthogonal to the aggregate order flow. It follows that price impact (Kyle's lambda) is a submartingale.*

PROOF: Note that from its definition, the G_t process satisfies

$$d\sqrt{G_t} + \frac{\sigma_t^2}{2\sqrt{G_t}}\,dt = d\mathcal{M}_t, \tag{64}$$

where $\mathcal{M}_t = \mathrm{E}[\int_0^T \frac{\sigma_t^2}{2\sqrt{G_t}}\,dt|\sigma^t]$ is a bounded martingale (adapted to the filtration generated by the noise trader volatility process) by the law of iterated expectation and since from equation (12) it is straightforward to show that $\mathcal{M}_t \leq \frac{\overline{\sigma}^2}{\underline{\sigma}}\sqrt{T}\ \forall t$.

It follows, by definition of the process σ_t, that $d\mathcal{M}_t\,dZ_t = 0$.

From its definition in (53) and the definition for Σ_t and G_t above, we obtain

$$d\frac{1}{\lambda(t)} = \frac{1}{\sqrt{\Sigma_t}}\,d\sqrt{G_t} - \frac{\sqrt{G_t}}{2(\Sigma_t)^{3/2}}\,d\Sigma_t \tag{65}$$

$$= \frac{1}{\sqrt{\Sigma_t}}\,d\mathcal{M}_t. \tag{66}$$

It also follows that $d\frac{1}{\lambda_t}\,dY_t = 0$.

To prove that λ is a submartingale, we apply Jensen's inequality. We have $\frac{1}{\lambda_t} = \mathrm{E}_t[\frac{1}{\lambda_s}] \geq \frac{1}{\mathrm{E}_t[\lambda_s]}$. It follows that $\lambda_t \leq E_t[\lambda_s]$. *Q.E.D.*

We can also prove a useful inequality for the G_t function:

LEMMA 6: *If a bounded solution G_t exists, then it satisfies*

$$G_t = \mathrm{E}\left[\int_t^T \sigma_s^2\,ds - \int_t^T \Sigma_s d\left[\frac{1}{\lambda}\right]_s\right]$$
$$\leq \mathrm{E}\left[\int_t^T \sigma_s^2\,ds\right].$$

PROOF: Apply Itô's formula to $\sqrt{G_t}^2$:

$$d\sqrt{G_t}^2 = 2\sqrt{G_t}\,d\sqrt{G_t} + d[\sqrt{G}]_t$$
$$= -\sigma_t^2\,dt + 2\sqrt{G_t}\,d\mathcal{M}_t + d[\sqrt{G}]_t.$$

Thus, integrating and taking expectation, we get the desired result where we use the fact that $\sqrt{G_t} = \frac{\sqrt{\Sigma_t}}{\lambda_t}$ and that Σ_t is an absolutely continuous process, to write $d[\sqrt{G}]_t = \Sigma_t \, d[\frac{1}{\lambda}]_t$. *Q.E.D.*

We now can prove the main result for this step, namely a verification proof of optimality of the insider's trading strategy (16). Recall that the insider is optimizing the following value process:

$$J_t = \max_{\{\theta_s\}_{s \geq t} \in \mathcal{A}} \mathrm{E}\left[\int_t^T (v - P_s)\theta_s \, ds \middle| \mathcal{F}_t^Y, v\right], \tag{67}$$

where the set of admissible strategies $\mathcal{A}$ is defined as the set of processes θ_t such that $\mathrm{E}[\int_0^T |\theta_s|^2 \, ds] < \infty$.

LEMMA 7: *Suppose price dynamics are given by equations* (52), (45), (53), *and* (11), *and that volatility is uniformly bounded* ($\underline{\sigma} \leq \sigma_t \leq \overline{\sigma}$); *then the optimal value process is given by*

$$J_t = \frac{(v - P_t)^2 + \Sigma_t}{2\lambda_t}, \tag{68}$$

and the optimal strategy is given by

$$\theta_t^* = \frac{1}{\lambda_t} \frac{\sigma_t^2}{G_t}(v - P_t). \tag{69}$$

PROOF: Apply Itô's rule to the conjectured value function to get

$$\begin{aligned} dJ_t = {} & \frac{(v - P_t)^2 + \Sigma_t}{2} \, d\frac{1}{\lambda_t} + \frac{1}{\lambda_t}\left(-(v - P_t)\, dP_t + \frac{1}{2}\, dP_t^2\right) \\ & - (v - P_t)\, dP_t \, d\frac{1}{\lambda_t} + \frac{1}{2\lambda_t}\, d\Sigma_t. \end{aligned} \tag{70}$$

The insider takes the price impact process as given and assumes the price process follows:

$$dP_t = \lambda_t(\theta_t \, dt + \sigma_t \, dZ_t),$$

with the λ process as in equation (53) above. Using Lemma 5 and the Σ_t dynamics, and integrating the above, we obtain

$$\begin{aligned} & J_T - J_0 + \int_0^T (v - P_t)\theta_t \, dt \\ & \quad = \int_0^T (P_t - v)\sigma_t \, dZ_t + \int_0^T \frac{(v - P_t)^2 + \Sigma_t}{2\sqrt{\Sigma_t}} \, d\mathcal{M}_t. \end{aligned} \tag{71}$$

Now, since $J_T \geq 0$, it follows by taking expectation (and using the fact that the stochastic integrals are martingales, as established in Lemma 8 below), that

$$\mathrm{E}\left[\int_0^T (v - P_t)\theta_t \, dt\right] \leq J_0 \tag{72}$$

for any admissible policy $\{\theta_t\}$. Further, if there exists a trading strategy θ_t consistent with the updating equations (44), such that $\mathrm{E}[J_T] = 0$ then, the inequality holds with equality.

The candidate policy in equation (16) satisfies this.

Indeed, note that

$$J_T = \frac{(v - P_T)^2 + \Sigma_T}{2\lambda_T} = \frac{(v - P_T)^2}{2\lambda_T} + \frac{\sqrt{\Sigma_T G_T}}{2} = \frac{(v - P_T)^2}{2\lambda_T}.$$

In turn,

$$\mathrm{E}\left[\frac{(v - P_T)^2}{\lambda_T}\right] \leq \sqrt{\mathrm{E}\big[(v - P_T)^2 G_T\big]\mathrm{E}\left[\left(\frac{v - P_T}{\sqrt{\Sigma_T}}\right)^2\right]} = 0,$$

where the right-hand-side equality follows from Lemma 3 and Lemma 4.

We have therefore proved the optimality of the value function and of the proposed policy. *Q.E.D.*

LEMMA 8: *Since σ_t is uniformly bounded above and below, the stochastic integrals $J_1(t) = \int_0^t (v - P_s)\sigma_s \, dZ_s$ and $J_2(t) = \int_0^t \frac{(v-P_s)^2 + \Sigma_s}{2\sqrt{\Sigma_s}} \, d\mathcal{M}_s$ are martingales for any admissible strategy.*

PROOF: To prove that $J_1(t)$ is a martingale, it is sufficient to show that $\mathrm{E}[\int_0^T (v - P_t)^2 \sigma_t^2 \, dt] < \infty$. In turn, because σ_t is uniformly bounded, it is sufficient to show that P_t has finite variance for all t. Note that $P_t = P_0 + \int_0^t \lambda_s \theta_s \, ds + \int_0^t \sigma_s \lambda_s \, dZ_s$. Thus, for P_t to have finite variance, it is sufficient that $\mathrm{E}[(\int_0^t \lambda_s \theta_s \, ds)^2] < \infty$ and $\mathrm{E}[\int_0^t \sigma_s^2 \lambda_s^2 \, ds] < \infty$. Clearly, $\mathrm{E}[\int_0^t \sigma_s^2 \lambda_s^2 \, ds] = \Sigma_0 - \Sigma_t < \infty$. Further, using the Cauchy–Schwarz inequality, we have

$$\mathrm{E}\left[\left(\int_0^t \lambda_s \theta_s \, ds\right)^2\right] \leq \mathrm{E}\left[\int_0^t \lambda_s^2 \, ds \int_0^t \theta_s^2 \, ds\right].$$

The right-hand side is finite for any admissible trading strategy since $\int_0^t \lambda_s^2 \, ds \leq \frac{1}{\underline{\sigma}} \int_0^t \frac{\Sigma_s}{G_s} \sigma_s^2 \, ds = \frac{\Sigma_0 - \Sigma_t}{\underline{\sigma}} < \frac{\Sigma_0}{\underline{\sigma}}$.

Next, to show that $J_2(t)$ is a martingale, since $\mathcal{M}_t$ is a uniformly bounded martingale (from Lemma 5) and Σ_t is a decreasing process, it is sufficient to show that $\int_0^T \frac{(v-P_t)^2}{\sqrt{\Sigma_t}}\, d\mathcal{M}_t$ is a martingale. For this stochastic integral to have finite variance, it is sufficient that $\mathrm{E}[\frac{(v-P_t)^4}{\Sigma_t}] < \infty\ \forall t$ given that $\mathcal{M}_t < \frac{\overline{\sigma}^2\sqrt{T}}{\underline{\sigma}^2}\ \forall t$ as shown in Lemma 5, and thus its quadratic variation is bounded.

Note that $\mathrm{E}[\frac{(v-P_t)^4}{\Sigma_t}] < \Sigma_0 \mathrm{E}[h_t^4]$, with h defined in Lemma 4. It follows from the properties of h_t obtained in Lemma 4 that $\mathrm{E}[h_t^4] < \infty\ \forall t$. *Q.E.D.*

A.4. *Derivation of the Expression (20) for Slippage Costs*

Intuitively, the total losses incurred between 0 and T by noise traders can be computed path-wise as

$$\int_0^T (P_{t+dt} - v)\sigma_t\, dZ_t = \int_0^T (P_t + dP_t - v)\sigma_t\, dZ_t = \int_0^T \lambda_t \sigma_t^2\, dt + \int_0^T (P_t - v)\sigma_t\, dZ_t. \tag{73}$$

The first component is the pure execution or slippage cost due to the fact that, in Kyle's model, agents submit market orders at time t that get executed at date $t + dt$ at a price set by competitive market makers. The second component is a fundamental loss due to the fact that, based on the price they observe at t, noise traders purchase a security with fundamental value v that is unknown to them. Note that since prices are set efficiently by market makers, on average this second component is zero. Therefore, we obtain the result that the unconditional expected total losses incurred by noise traders are entirely driven by expected execution costs. Further, these are also equal to the total unconditional expected profits of the insider. However, note that, path-wise, neither quantity need be equal. To show that unconditional expected execution costs paid by noise traders are equal to the unconditional expected profits of the insider, note that the insider's unconditional expected profits are

$$\mathrm{E}^v\left[\int_0^T \theta_t(v - P_t)\, dt\right] = \mathrm{E}^v\left[\int_0^T \frac{\sigma_t^2}{\sqrt{\Sigma_t G_t}}(v - P_t)^2\, dt\right] = \mathrm{E}^v\left[\int_0^T \frac{\sigma_t^2}{\sqrt{\Sigma_t G_t}}\Sigma_t\, dt\right] = \mathrm{E}^v\left[\int_0^T \sigma_t^2 \lambda_t\, dt\right], \tag{74}$$

where the first equality follows from the definition of θ^* and the second from the law of iterated expectations. This is the same expression obtained for the unconditional expected execution costs paid by noise traders. By definition, this is also equal to $\mathrm{E}^v[J(0)] = \sqrt{\Sigma_0 G_0}$, where the expectation superscript emphasizes that it is taken over the unconditional distribution of v.

A.5. *Proof of Theorem 2: Deterministic Growth Rate of Noise Trading Volatility*

Suppose the drift m_t is deterministic. Guess that the solution is of the form $G_t = \sigma_t^2 B_t$, where B_t is the solution given in the lemma. Plugging into equation (11), we see that our guess is correct if B_t solves

$$\sqrt{B_t} = \int_t^T \frac{e^{\int_t^u m_s\,ds}}{2\sqrt{B_u}}\,du,$$

where we have used the fact that, for any martingale W_t, we have (for $u \geq t$)

$$\mathrm{E}_t[\sigma_u] = \sigma_t e^{\int_t^u m_s\,ds}. \tag{75}$$

The solution for B_t given in the lemma indeed satisfies this integral equation. Since uniqueness was established before, we have found the solution of equation (22).

All results in the theorem follow directly from Theorem 1 and the expression for G_t.

The only new result is the calculation of the expected trading rate of the insider:

$$\mathrm{E}[\theta_t | v, \mathcal{F}_0] = \mathrm{E}\left[\frac{(v - P_t)}{\sqrt{\Sigma_t}} \frac{\sigma_t^2}{\sqrt{G_t}}\right] \tag{76}$$

$$= \frac{(v - P_0)}{\sqrt{\Sigma_0}} e^{-\int_0^t (1/(2B_s))\,ds} \frac{\sigma_0 e^{\int_0^t m_s\,ds}}{\sqrt{B_t}}, \tag{77}$$

where we used the dynamics of h_t from Lemma 5 and the expression for G_t from equation (22). The result in the theorem then follows from standard manipulations (in particular, note that $e^{\int_0^t -(1/(2B_s))\,ds} = \sqrt{\frac{B_t}{B_0}} e^{\int_0^t m_s\,ds}$). *Q.E.D.*

A.6. *Proof of Theorem 3*

Consider a pair of functions $G^L(\cdot), G^H(\cdot)$ that solve the ODE system (29)–(30) subject to the boundary condition $G^L(0) = G^H(0) = 0$; then it is straightforward to show that if we define $G(t, \sigma_t) = \mathbf{1}_{\{\sigma_t = \sigma^L\}} G^L(T - t) + \mathbf{1}_{\{\sigma_t = \sigma^H\}} G^H(T - t)$, then $J(t) = \sqrt{G(t, \sigma_t)} + \int_0^t \frac{\sigma_u^2}{2\sqrt{G(u, \sigma_u)}}\,du$ is a pure jump martingale (i.e., $\mathrm{E}_t[dJ(t)] = 0$). It follows that $J(t) = \mathrm{E}_t[J(T)]$ and, using the definition of $J(t)$ and the boundary conditions from the ODEs, that $\sqrt{G(t, \sigma_t)} = \mathrm{E}_t[\int_t^T \frac{\sigma_u^2}{2\sqrt{G(u, \sigma_u)}}\,du]$. *Q.E.D.*

A.7. *Proof for the Time Change*

This result follows from the definition of the time-change which implies

$$\begin{aligned} d\tau_t &= -\frac{1}{\sigma_v^2} d\Sigma_t \\ &= \frac{\Sigma_t \sigma_t^2}{\sigma_v^2 G_t} \\ &= \frac{(\tau_T - \tau_t)\sigma_t^2}{G_t}. \end{aligned}$$

Substituting in the definition of the equilibrium price process in Theorem 1 establishes the result. Note that the properties of the equilibrium imply that $\tau_T = T$ a.s. and thus the time-change is indeed adapted. Of course, $\tau_0 = 0$. Importantly, we cannot obtain our equilibrium simply as a time-change of the standard Kyle–Back models when $\nu \neq 0$ since, in that case, price impact is a stochastic process in our equilibrium (i.e., is not constant), whereas price impact is constant in the Kyle–Back model.

REFERENCES

ADMATI, A., AND P. PFLEIDERER (1988): "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, 1 (1), 3–40. [1443,1450]

ANDERSEN, T. G. (1996): "Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility," *The Journal of Finance*, 51 (1), 169–204. [1444,1463]

BACK, K. (1992): "Insider Trading in Continuous Time," *Review of Financial Studies*, 5, 387–409. [1442,1444,1445,1447,1449,1450,1452]

BACK, K., AND S. BARUCH (2004): "Information in Securities Markets: Kyle Meets Glosten and Milgrom," *Econometrica*, 72 (2), 433–465. [1442,1450]

BACK, K., AND H. PEDERSEN (1998): "Long-Lived Information and Intraday Patterns," *Journal of Financial Markets*, 1, 385–402. [1442,1443,1450,1460]

BACK, K., C. H. CAO, AND G. A. WILLARD (2000): "Imperfect Competition Among Informed Traders," *The Journal of Finance*, 55 (5), 2117–2155. [1450]

BARUCH, S. (2002): "Insider Trading and Risk Aversion," *Journal of Financial Markets*, 5 (4), 451–464. [1442,1450]

BIAIS, B., L. GLOSTEN, AND C. SPATT (2005): "Market Microstructure: A Survey of Microfoundations, Empirical Results, and Policy Implications," *Journal of Financial Markets*, 8 (2), 217–264. [1461]

CALDENTEY, R., AND E. STACCHETTI (2010): "Insider Trading With a Random Deadline," *Econometrica*, 78 (1), 245–283. [1450]

CLARK, P. K. (1973): "A Subordinated Stochastic Process Model With Finite Variance for Speculative Prices," *Econometrica*, 41 (1), 135–155. [1442,1444,1462,1463]

COLLIN-DUFRESNE, P., AND V. FOS (2015): "Do Prices Reveal the Presence of Informed Trading?" *The Journal of Finance*, 70 (4), 1555–1582. [1442,1462]

——— (2016a): "Insider Trading With a Random Horizon," Working Paper, SFI, EPFL. [1458]

——— (2016b): "Supplement to 'Insider Trading, Stochastic Liquidity, and Equilibrium Prices'," *Econometrica Supplemental Material*, 84, http://dx.doi.org/10.3982/ECTA10789. [1454]

EPPS, T. W., AND M. L. EPPS (1976): "The Stochastic Dependence of Security Price Changes and Transaction Volumes: Implications for the Mixture-of-Distributions Hypothesis," *Econometrica*, 44 (2), 305–321. [1463]

FOSTER, F. D., AND S. VISWANATHAN (1990): "A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Securities Markets," *Review of Financial Studies*, 3 (4), 593–624. [1443]

——— (1993): "Variations in Trading Volume, Return Volatility, and Trading Costs: Evidence on Recent Price Formation Models," *The Journal of Finance*, 48 (1), 187–211. [1443]

——— (1995): "Can Speculative Trading Explain the Volume–Volatility Relation?" *Journal of Business & Economic Statistics*, 13 (4), 379–396. [1443]

——— (1996): "Strategic Trading When Agents Forecast the Forecasts of Others," *The Journal of Finance*, 51 (4), 1437–1478. [1450]

GALLANT, A. R., P. E. ROSSI, AND G. TAUCHEN (1992): "Stock Prices and Volume," *Review of Financial Studies*, 5 (2), 199–242. [1442,1463]

KARATZAS, I., AND S. E. SHREVE (1991): *Brownian Motion and Stochastic Calculus*. Berlin: Springer-Verlag. [1468]

KYLE, A. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53 (6), 1315–1335. [1441,1443,1444,1450,1460,1464]

LAMOUREUX, C. G., AND W. D. LASTRAPES (1990): "Heteroskedasticity in Stock Return Data: Volume versus GARCH Effects," *The Journal of Finance*, 45 (1), 221–229. [1442]

LEPELTIER, J.-P., AND J. SAN MARTIN (1997): "Existence for BSDE With Superlinear–Quadratic Coefficient," *Stochastics and Stochastic Reports*, 63 (3–4), 227–240. [1464,1465]

LIPTSER, R., AND A. SHIRYAEV (2001): *Statistics of Random Processes II: Applications* (Second Ed.). Berlin: Springer-Verlag. [1445]

MADHAVAN, A., M. RICHARDSON, AND M. ROOMANS (1997): "Why Do Security Prices Change? A Transaction-Level Analysis of NYSE Stocks," *Review of Financial Studies*, 10 (4), 1035–1064. [1443]

RICHARDSON, M., AND T. SMITH (1993): "A Test for Multivariate Normality in Stock Returns," *The Journal of Business*, 66 (2), 295–321. [1444,1463]

STOLL, H. R. (1978): "The Supply of Dealer Services in Securities Markets," *The Journal of Finance*, 33 (4), 1133–1151. [1462]

TAUCHEN, G. E., AND M. PITTS (1983): "The Price Variability–Volume Relationship on Speculative Markets," *Econometrica*, 51 (2), 485–505. [1444,1463]

WATANABE, M. (2008): "Price Volatility and Investor Behavior in an Overlapping Generations Model With Information Asymmetry," *The Journal of Finance*, 63 (1), 229–272. [1443]

SFI, CDM, EPFL, SFI-PCD EXTRA 209 (Bâtiment Extranef UNIL), CH-1015 Lausanne, Switzerland; pierre.collin-dufresne@epfl.ch

and

Dept. of Finance, Carroll School of Management, Boston College, Fulton Hall 338, 140 Commonwealth Avenue, Chestnut Hill, MA 02467, U.S.A.; fos@bc.edu.

Co-editor Lars Peter Hansen handled this manuscript.

Manuscript received May, 2012; final revision received February, 2016.