On the Correlation Structure of Microstructure Noise: A Financial Economic Approach

Francis X. Diebold*
University of Pennsylvania
and NBER

Georg Strasser@
Boston College

This Draft: April 24, 2012

Abstract

We introduce the financial economics of market microstructure into the financial econometrics of asset return volatility estimation. In particular, we use market microstructure theory to derive the cross-correlation function between latent returns and market microstructure noise, which feature prominently in the recent volatility literature. The cross-correlation at zero displacement is typically negative, and cross-correlations at nonzero displacements are positive and decay geometrically. If market makers are sufficiently risk averse, however, the cross-correlation pattern is inverted. We derive model-based volatility estimators, which we apply to stock and oil prices. Our results are useful for assessing the validity of the frequently-assumed independence of latent price and microstructure noise, for explaining observed cross-correlation patterns, for predicting as-yet undiscovered patterns, and for microstructure-based volatility estimation.

Acknowledgments: For comments and suggestions we thank Rob Engle, Peter Hansen, Charles Jones, Eugene Kandel, Gideon Saar, Frank Schorfheide, Enrique Sentana, and three anonymous referees. We are also grateful to participants at the NBER Conference on Market Microstructure, the Oxford-Man Institute Conference on the Financial Econometrics of Vast Data, and the Christmas Meeting of German Economists Abroad. We thank Ole E. Barndorff-Nielsen, Peter R. Hansen, Asger Lunde and Neil Shephard for sharing their data with us.

Key Words: Realized volatility, Market microstructure theory, High-frequency data, Financial econometrics

JEL Codes: G14, G20, D82, D83, C51

*fdiebold@sas.upenn.edu
@Georg.Strasser@bc.edu

1 Introduction

Recent years have seen substantial progress in asset return volatility measurement, with important applications to asset pricing, portfolio allocation and risk management. In particular, so-called realized variances and covariances ("realized volatilities"), based on increasingly-available high-frequency data, have emerged as central for several reasons.[1] They are, for example, largely model-free (in contrast to traditional model-based approaches such as GARCH or stochastic volatility), they are computationally trivial, and they are in principle highly accurate.

A tension arises, however, linked to the last of the above desiderata. Econometric theory suggests the desirability of sampling as often as possible to obtain highly accurate volatility estimates, but financial market reality suggests otherwise. In particular, market microstructure noise (MSN), such as bid-ask bounce associated with ultra-high-frequency sampling, may contaminate the observed price, potentially rendering naively-calculated realized volatilities unreliable.

Early work (e.g., Andersen, Bollerslev, Diebold, and Ebens, 2001a; Andersen, Bollerslev, Diebold, and Labys, 2001b, 2003, Barndorff-Nielsen and Shephard, 2002a,b) addressed the sampling issue by attempting to sample often, but not "too often," typically resulting in use of five- to thirty-minute returns. Much higher-frequency data are usually available, however, so reducing the sampling frequency to insure against MSN discards potentially valuable information.

To use all information, more recent work has emphasized MSN-robust realized volatilities that use returns sampled at very high frequencies. Examples include Zhang, Mykland, and Aït-Sahalia (2005), Bandi and Russell (2008), Aït-Sahalia, Mykland, and Zhang (2011), Hansen and Lunde (2006), and Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008, 2011b). That literature is almost entirely *statistical*, however, which is unfortunate because it makes important assumptions regarding the nature of the latent price, the MSN, and their interaction, and purely statistical thinking offers little guidance. A central example concerns the interaction (if any) between latent price and MSN. Some authors such as Bandi and Russell assume no correlation (perhaps erroneously), whereas in contrast Barndorff-Nielsen et al. (2008); Barndorff-Nielsen, Hansen, Lunde, and Shephard (2011a) allow for correlation (perhaps unnecessarily).

[1] Several surveys are now available, ranging from the comparatively theoretical treatments of Barndorff-Nielsen and Shephard (2007) and Andersen, Bollerslev, and Diebold (2010) to the applied perspective of Andersen, Bollerslev, Christoffersen, and Diebold (2006).

To improve this situation, we explicitly recognize that MSN results from the behavior of economic agents, and we push toward integration of the financial economics of market microstructure with the financial econometrics of volatility estimation. In particular, we explore the implications of microstructure theory for the relationship between latent price and MSN, characterizing the cross-correlation structure between latent price and MSN, contemporaneously and dynamically, in a variety of leading environments, including those of Roll (1984), Glosten and Milgrom (1985), Kyle (1985), Easley and O'Hara (1992), and Hasbrouck (2002).[2]

We proceed as follows. In Section 2 we introduce our general framework, which nests a variety of microstructure models. In Sections 3 and 4 we provide detailed analyses of models of private information, distinguishing two types of latent prices based on the implied level of market efficiency. In particular, we treat strong form efficiency in Section 3 and semi-strong form efficiency in Section 4. In Section 5 we discuss the relationship between price change frequency and sampling frequency. Based on this, we suggest several microstructure-founded estimators and apply them to stock and oil market data in Section 6. We conclude in Section 7.

2 The Framework

We begin in Section 2.1 by introducing a general framework relating latent prices, observed prices, and MSN in a wide range of market-making environments. We then provide, in Section 2.2, a generic (model-free) statistical result on the nature of correlation between latent price and MSN. Finally, in Section 2.3, we introduce market makers, or – more generally – learning market participants, who are central in the subsequent analyses.

2.1 Latent Prices, Observed Prices and Microstructure Noise

Let p_t^* denote the (logarithm of the) *strong form efficient price* of some asset in the calendar (or business) time period t. This price, strictly exogenously changing every T^{th}-period, could stem from sampling increments of standard Brownian motion every T periods, in which case the standard deviation σ would be proportional to T. At time t, p_t^* is known only to the

[2]For insightful surveys of the key models, see O'Hara (1995) and Hasbrouck (2007). For an interesting related perspective, see Engle and Sun (2007). Their approach and environment (conditional duration modeling), however, are very different from ours.

informed traders, and follows the process

$$p_t^* = \begin{cases} p_{t-1}^* + \sigma\varepsilon_t, & \forall t = \kappa T, \kappa \in \mathbb{Z} \\ p_{t-1}^*, & \text{otherwise} \end{cases} \tag{1}$$

$$\text{with } \varepsilon_t \overset{iid}{\sim} (0,1). \tag{2}$$

This price process is very restrictive. For simplicity of exposition we do not model jumps, time-varying volatility (σ_t), or time-varying sampling intervals (T_t), which are the subject of sophisticated models of market microstructure theory. In all its simplicity, however, this process is the discrete time analogue of the latent price process that estimators of integrated volatility (IV) are based on. As we show later in this paper, different assumptions about the nature of the latent price process will lead to different estimates of IV. In particular, the properties of the latent price relevant in many applications depend on the information set. In this paper we aim to bridge the gap between market microstructure theory and IV estimation by introducing for the first time a simple price determination framework founded on market microstructure theory to IV estimation.

Microstructure *noise* (MSN) is the difference between the observed market return and the latent return. Instead of ad-hoc assumptions about the properties of the *strong form noise*

$$\Delta u_t \equiv \Delta p_t - \Delta p_t^*, \tag{3}$$

which are common in the IV estimation literature, we add additional market microstructure that helps explain key properties of MSN.

Let q_t denote the direction of the trade in period t, where $q_t = +1$ denotes a buy, $q_t = -1$ a sell, and $q_t = 0$ a no-trade period. Define p_t^e as the expected efficient price directly before the trade occurs. The *semi-strong form efficient price*, which summarizes the knowledge of the market maker after the trade,[3] is in logarithmic terms

$$\tilde{p}_t^e = p_t^e + \lambda_t q_t, \tag{4}$$

where $\lambda_t \geq 0$ captures the response to asymmetric information revealed by the trade direction q_t. The admittedly stylized assumption that quantities do not matter for market maker

[3] This terminology is borrowed from the asset pricing literature. In contrast to the strong form efficient price, which incorporates all public and private information, the semi-strong form efficient price only incorporates all publicly available information (Fama, 1970).

learning obtains e.g. in a pooling equilibrium of informed with uninformed traders (Kelly and Steigerwald, 2004). It fits the observation that in recent years order-splitting into many small trades has become dominant. Because the estimators we derive rely (at most) on trade direction data, further model detail would not add to our results.

At the beginning of each trading round, additional information about p_t^* and ε_t might be revealed by information diffusion from other sources, e.g. other markets. With this information, summarized by ω_t, the market maker revises his price expectation for the next period according to

$$p_t^e = \tilde{p}_{t-1}^e + \omega_t. \tag{5}$$

In periods in which p_{t-1}^* becomes public information, (5) becomes $p_t^e = p_{t-1}^* + \tilde{\omega}_t$. Assuming that the price quotes in logarithmic terms are symmetric around the expected efficient price before the trade, the observed *transaction price* can be written as

$$p_t = p_t^e + s_t q_t, \tag{6}$$

where s_t is one-half of the spread. In particular, the bid price is $p_t^{bid} = p_t^e - s_t$, the ask price is $p_t^{ask} = p_t^e + s_t$, and the midprice is p_t^e. These prices and their relationships are illustrated by Figure 1. We assume throughout that market conditions are stable and that transaction prices p_t adjust sufficiently fast so that the noise process Δu_t is covariance stationary.

Figure 1: Timing of Information and Prices

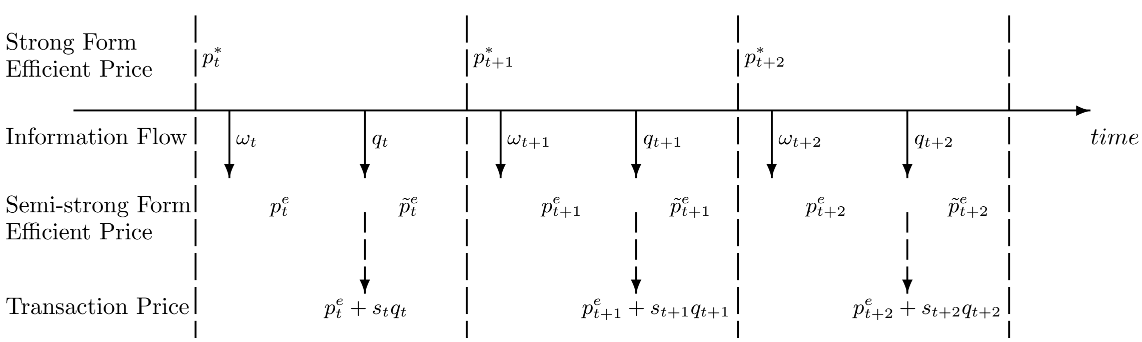

Our stylized setup covers three levels of information: full, intermediate (market maker), and public information. Of course, in reality market participants are more heterogeneous with respect to their information sets. Consider, for example, the difference between traders with to those without access to Nasdaq level II screens. The former traders cannot see the order book, whereas the latter can. We model for concreteness' sake the intermediate price as the market maker's price. It could, of course, also reflect some other information set, e.g. the one of traders with access to a semi-public market information source.

Strong form efficient returns in periods $t = \kappa T$ are therefore

$$\Delta p_t^* \equiv p_t^* - p_{t-1}^* = \sigma \varepsilon_t, \tag{7}$$

and zero in all other periods. *Semi-strong form efficient returns* are

$$\Delta \tilde{p}_t^e \equiv \tilde{p}_t^e - \tilde{p}_{t-1}^e = \lambda_t q_t + \omega_t, \tag{8}$$

and *semi-strong form noise* is accordingly

$$\Delta \tilde{u}_t \equiv \Delta p_t - \Delta \tilde{p}_t^e. \tag{9}$$

We use the term "*latent price*" as a general term comprising both types of efficient prices. The two latent prices defined here are conceptually very distinct and appeal to distinct audiences. For example, on the one hand, a pure theorist may want to understand the properties of the full-information price, and is thus interested in an estimate of the volatility of the strong form efficient return (7). One the other hand, a market maker may need a volatility measure to calculate his risk exposure, thus his relevant price for the asset is $\tilde{p}_t^e$, the price at which he keeps the asset on his accounts. It is the volatility of (8), and not of (7), that affects his balance sheet.

Semi-strong form noise (9) differs fundamentally in its cross-correlation properties from (3). It is therefore essential for a researcher to be clear what type of latent price the object of interest is, because each requires different procedures to remove MSN appropriately.

Observed *market returns* are

$$\Delta p_t \equiv p_t - p_{t-1} = \Delta p_t^e + s_t q_t - s_{t-1} q_{t-1}.$$

A convenient estimator of the variance of the strong form efficient return, σ^2, and therefore of the IV of the underlying continuous time process, is the realized volatility (RV) as in Andersen et al. (2001b). RV during the time interval $[0, \bar{T}]$ is defined as the sum of squared market returns over the interval, i.e. as

$$Var(\Delta p_t) = \sum_{t=1}^{\bar{T}} \Delta p_t^2.$$

In the presence of MSN, the RV is generally a biased estimate of σ^2. To see this, decompose the noise into two components, one uncorrelated and one correlated with the latent price, so that $\Delta u_t = \Delta u_t^u + \Delta u_t^c$. The uncorrelated component, Δu_t^u, reflects for example the bid-ask bounce in a market populated with uninformed traders only. The correlated component, Δu_t^c, reflects for example the effect of asymmetric information. RV can now be decomposed – here shown for the strong form efficient price – as

$$\begin{aligned} Var(\Delta p_t) &= Var(\Delta p_t^* + \Delta u_t^u + \Delta u_t^c) \\ &= \sigma^2 + Var(\Delta u_t^u) + Var(\Delta u_t^c) + 2Cov(\Delta p_t^*, \Delta u_t^c). \end{aligned}$$

The bias of RV can stem from any of the last three terms, which are all nonzero in general. IV estimation under the independent noise assumption accounts for the second and third positive terms, but ignores the last term, which is typically negative (Hansen and Lunde, 2006). Correcting the estimates for independent noise only, always reduces the volatility estimate. But because such a correction ignores the last term, which is the second channel through which asymmetric information affects the IV estimate, the overall reduction might be too much. Further, serial correlation of noise, or equivalently a cross-correlation between noise and latent returns at nonzero displacement, requires the use of robust estimators for both the variance and the covariance terms. In this paper we determine what correlation and serial correlation market microstructure theory predicts, and how market microstructure theory can be useful for improving IV estimates.

2.2 Statistical Characterization of Return/Noise Correlations

We focus in this paper on the *cross-correlation* between latent returns and noise contemporaneously and at all displacements. Throughout, we refer to this quantity simply as the "cross-correlation".

Under very general conditions the contemporaneous cross-correlation for the price processes given by (1)–(6) is positive only if the market return, Δp_t, is more volatile than the latent return. More precisely, for strong form efficient returns

$$Corr(\Delta p_t^*, \Delta u_t) > 0 \Leftrightarrow E(\Delta p_t \Delta p_t^*) > Var(\Delta p_t^*) \Leftrightarrow Corr(\Delta p_t, \Delta p_t^*) > \sqrt{\frac{Var(\Delta p_t^*)}{Var(\Delta p_t)}}. \quad (10)$$

Cross-correlations at displacements $\tau \geq 1$ are positive if and only if the current transaction price responds stronger in the direction of a latent price change τ periods ago than the current latent price itself. More precisely, for strong form efficient returns

$$Corr(\Delta p^*_{t-\tau}, \Delta u_t) > 0 \Leftrightarrow E(\Delta p_t \Delta p^*_{t-\tau}) > 0. \tag{11}$$

The conditions for semi-strong form efficient returns are analogous (see Diebold and Strasser, 2010). Whereas the price processes as defined in the previous subsection suffice to mechanically derive expressions for their cross-correlation, this reduced form setup alone does not give much guidance about sign and time pattern of these cross-correlations. In the financial economic environments that will concern us the properties of prices are determined by the market microstructure. Hence we introduce it now in some detail.

2.3 Introducing Markets and Market Makers

Whereas the strong form efficient price (1) is an exogenous stochastic process, the semi-strong form efficient price (4) and the transaction price (6) are an outcome of the market participants' optimizing behavior. As such the latter are not time series of unknown properties generated by a black box. Instead, key properties of the data generator – the financial market – are often observable and allow inferring properties of these price series. This is what we do in this paper.

Generally speaking, the transaction price depends on the information available about the strong form efficient price and the market participants' response to this information. Three features of the *information* process matter in particular: First, information content, second, the diffusion speed of information into public knowledge, and third, the duration of its validity. The *price updating* rule determines how, and how quickly, transaction prices respond to new information. Of particular importance is whether the market maker can quote prices dependent on the direction of trade, i.e. whether he is free to charge any spread, because direction-dependent quotes allow prices to react instantaneously.

We focus here on a stylized limit-order market, populated by informed and uninformed traders. Market makers are the counterparty of all trades. Each trading round they quote price p^e_t and spread s_t for one unit of the asset. Thereafter, as shown in Figure 2, informed traders screen the market with probability α for profitable trading opportunities. They buy if $p^*_t > p^{ask}_t$, sell if $p^*_t < p^{bid}_t$, and refuse to trade otherwise. In periods of no informed trade, uninformed traders trade instead with probability β, buying and selling with equal probability.

When trading with an informed trader the market maker always loses. His expected loss

Figure 2: Sequence of Informed and Uninformed Trading Decisions

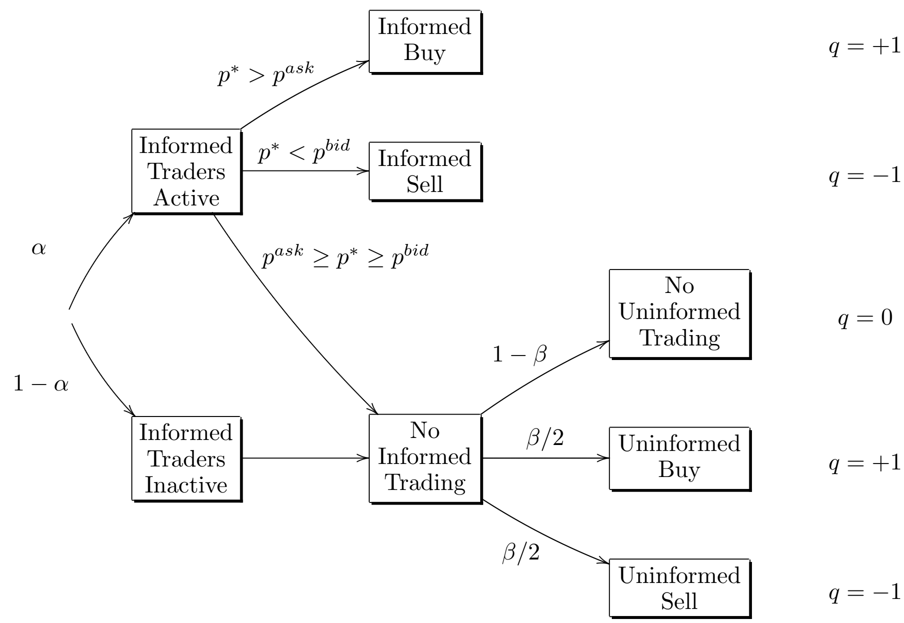

is

$$L_n\left[p_t, F(\cdot;\underline{p}^*, \overline{p}^*)\right] \quad = \quad -\int_{\underline{p}}^{\overline{p}} \left|(p_t - p_t^*)E(q_t\,|p_t^*, p_t, s_t)\right|^n f\left(p_t^*\right)dp_t^*, \tag{12}$$

where $E(q_t\,|p_t^e + s_t < p_t^*) = \alpha$, $E(q_t\,|p_t^e - s_t > p_t^*) = -\alpha$, $E(q_t\,|p_t^e - s_t \leq p_t^* \leq p_t^e + s_t) = 0$, and n reflects the risk aversion of the market maker. $F(\cdot)$ and $f(\cdot)$ denote the cdf and pdf with support $\left[\underline{p}, \overline{p}\right]$ of the market maker's belief about the latent price. Similar to Aghion, Bolton, Harris, and Jullien (1991), the market maker faces a tradeoff between avoiding losses today and learning quickly.[4]

Because price quotes are only for limited quantities, and the market maker can in principle update his price quote after every trade, his risk exposure is usually small. Accordingly, we assume risk neutrality ($n = 1$) throughout the paper, and relegate the implications of risk aversion to Section 3.3.2. As shorthand notation for the probability of a trade we define

$$\phi_t = E(q_t^2) = E\left[Prob(|q_t| = 1)\right] = \beta + (1-\beta)\alpha\left[1 - F(p_t^e + s_t) + F(p_t^e - s_t)\right].$$

[4]Diebold and Strasser (2010) describe the market setup and maker problem in more detail.

Note that the model can be recast in tick-time by setting $\phi_t = 1 \ \forall t$. We add the following assumption, which simplifies the model without affecting its basic behavior.

Assumption 1 *Ex ante, a buy and a sell is equally likely, so that $E(q_t) = 0$. There is no "momentum" in uninformed trading, and thus trades are serially uncorrelated beyond the time of a strong form efficient price change, i.e. $E(q_{\kappa T+\tau_1}|q_{\kappa T-\tau_2}) = 0 \ \ \forall \kappa, \tau_1 \in \mathbb{N}_0, \forall \tau_2 \in \mathbb{N}$.*

In the following Sections 3 and 4 we look at specializations of this general market maker problem and examine the effect of various model setups on the cross-correlation function. For both strong form and semi-strong form efficient returns we first examine the multiperiod case, where private information is not revealed until after many periods. We then specialize to the one-period case, a case where private information becomes public, and worthless, after only one period, where we specifically address the effect of risk-aversion.

3 Return-Noise Correlations in Financial Economic Environments I: Strong Form Efficient Prices

Here we characterize cross-correlations in an environment of strong form efficient prices. We calculate the cross-correlations between strong form efficient returns (7) and the corresponding noise (3) in various market settings. To study the effect of one efficient price change in isolation, suppose for now that there is a change in the strong form efficient price at a commonly known time at which the previous change becomes public knowledge. To fix ideas, let this change occur also every T periods.

3.1 The General Multi-Period Case

The cross-correlations, as shown in Web Appendix A.0.1, follow directly from the price and noise processes. The contemporaneous cross-covariance is

$$Cov(\Delta p_t^*, \Delta u_t) = \frac{\sigma}{T}\left[s_0 E(q_0\varepsilon_0) - \sigma + E(\omega_0\varepsilon_0)\right]. \tag{13}$$

For cross-covariance at higher displacements $\tau \in [1; T-1]$ we get

$$Cov(\Delta p_{t-\tau}^*, \Delta u_t) = \frac{\sigma}{T}\left[(\lambda_{\tau-1} - s_{\tau-1})E(q_{\tau-1}\varepsilon_0) + s_\tau E(q_\tau\varepsilon_0) + E(\omega_\tau\varepsilon_0)\right], \tag{14}$$

for cross-covariance at displacement T, which is when private information becomes public,

$$Cov(\Delta p^*_{t-T}, \Delta u_t) = \frac{\sigma}{T}\left[\sigma - s_{T-1}E(q_{T-1}\varepsilon_0) - \sum_{i=0}^{T-2}\lambda_i E(q_i\varepsilon_0) - \sum_{i=0}^{T-1}E(\omega_i\varepsilon_0)\right], \tag{15}$$

and for all higher order displacements $\tau > T$

$$Cov(\Delta p^*_{t-\tau}, \Delta u_t) = 0. \tag{16}$$

Combining (13) with the noise variance derived in the Web Appendix gives the contemporaneous cross-correlation

$$Corr(\Delta p^*_t, \Delta u_t) = \frac{s_0 E(q_0\varepsilon_0) - \sigma + E(\omega_0\varepsilon_0)}{\sqrt{T\,Var(\Delta u_t)}}. \tag{17}$$

All other cross-correlations can be obtained analogously.

The term $E(q_\tau\varepsilon_0)$ enters the expressions for the cross-covariance (13)–(15) linearly but the denominator of the cross-correlation under a square root. Because this term decreases in the share of uninformed trades, the contemporaneous cross-correlation is the smaller, the less informed traders are active. In absence of both informed traders ($E(q_\tau\varepsilon_0) = 0$) and of extra information ($E(\omega_0\varepsilon_0) = 0$), the market microstructure reduces to a bid-ask bounce, as in Roll (1984). Even in this case, shown in the first row of Table 1, the latent price and noise are not independent. The contemporaneous cross-correlation (17) is negative, the cross-correlations at displacement T is positive and all other cross-correlations are zero.

Because of order splitting, effective spreads have become very small for liquid assets. If no extra information is available and the spread sufficiently small, then the contemporaneous cross-correlation is negative even in presence of informed traders, because p_t does not react sufficiently to Δp^*_t. It is strictly larger than negative one, because the delayed response of Δp_t to $\Delta p^*_{t-\tau}$ generates cyclical noise with – absent other market microstructure effects – up to twice the variance of Δp^*_t. Likewise, if the spread roughly matches the adverse selection coefficient, by (14) the cross-correlations at displacements one up to $T-1$ are positive, which reflects that the more the market maker learns, the closer p_t gets to p^*_t, and the more noise shrinks to zero. If, additionally, the adverse selection coefficient λ and extra information ω in all periods are sufficiently small, i.e. if some private information persists until period T, then by (15) the cross-correlation at displacement T is positive as well.

In general, however, the sign of the cross-correlations depends on the behavior of market makers and traders. We now turn to models that allow us to introduce these explicitly.

3.2 Special Multi-Period Cases of Informed Trading

The market maker does not observe the strong form efficient price, p_t^*, directly, but only signals which allow him to narrow down the range of the current p_t^* level. He observes in particular the response of traders to his previous price quote and uses this signal to revise his quote. Because in this section p_t^* by assumption does not change after the initial jump for T periods, the market maker can use the entire sequence of signals to learn p_t^* over time. The market maker has an incentive to find out p_t^*, because he loses in every trade with an informed trader. His optimization task is to quote prices that minimize his losses by learning about p_t^* as quickly as possible.

He learns over time "by experimentation" about the informed traders' private information by setting prices and observing the resulting trades (Aghion et al., 1991; Aghion, Espinosa, and Jullien, 1993). We will see that rational behavior of market participants and the market setup pins down the cross-correlation sign pattern. Only the absolute value of the cross-correlation differs depending on how market participants interact.

The recursive problem of the market maker is hard to solve, and in particular there are in general no closed form policy functions p_t^{bid} and p_t^{ask}. Therefore we follow the market microstructure literature by discussing interesting polar cases, which can be solved because $f(p_t^*)$ is degenerate. In particular, we limit our discussion to the midprice under a constant spread.

3.2.1 No Strategic Traders

Consider first a market in which the market maker observes only a noisy signal of whether p_t^* has changed, but in which traders do not behave strategically. The market maker has to learn both about the quality of the signal and about the latent price. A useful illustration is the stylized model of Easley and O'Hara (1992). As in our general setup in Section 2.3 informed traders are active with probability α. In this model, the strong form efficient price is not a martingale. The latent price can assume one of two possible levels, namely $p_t^* = \underline{p}^*$ or $p_t^* = \overline{p}^* > \underline{p}^*$. These levels, as well as the probability γ of $p_t^* = \underline{p}^*$, are publicly known, but the actual realization of p_t^* is not.[5]

[5] The case of signal certainty, which implies the absence of any uninformed traders, is trivial here: Because p_t^* can assume only one of two price levels, the first trade reveals the true strong form efficient price. Until the first trade occurs, the expected efficient price is $\gamma \underline{p}^* + (1-\gamma)\overline{p}^*$.

The direction-of-trade signal, q_t, is thereby noisy in two ways. Not only does the market maker not know if a specific trade originates from informed traders, thereby being informative; the market maker does not even know if there *are* any informed traders. He learns by updating in a Bayesian manner his belief about the probabilities that nobody observed a signal, that informed traders observed $p_t^* = \overline{p}^*$, or that they observed $p_t^* = \underline{p}^*$, using his information set of all previous quotes and trades. Even no-trade intervals contain information about p_t^*, because they lower the probability that informed traders are active.[6]

Denote $\beta_{\tau,\{\overline{p}^*\}} \in [0,1]$ the belief at time $t+\tau$ that a high latent price has been observed, $\beta_{\tau,\{\underline{p}^*\}}$ the belief that a low latent price has been observed and $\beta_{\tau,\{\}}$ the belief that nobody has observed any signal, all conditional on the market maker's information set. The market maker sets the bid price, for example, under perfect competition to

$$\begin{aligned} p_\tau^{bid} - \underline{p}^* &= \beta_{\tau,\{\underline{p}^*\}}(1-\beta_{\tau,\{\}})\underline{p}^* + \beta_{\tau,\{\overline{p}^*\}}(1-\beta_{\tau,\{\}})\overline{p}^* + \beta_{\tau,\{\}}\frac{\underline{p}^* + \overline{p}^*}{2} - \underline{p}^* \\ &= \left(\beta_{\tau,\{\overline{p}^*\}} + \frac{\beta_{\tau,\{\}}}{2}\right)\left(\overline{p}^* - \underline{p}^*\right). \end{aligned}$$

A sufficiently large τ allows invoking a law of large numbers for the observations included in the market maker believes. Easley and O'Hara (1992) show for the case that traders observed a low latent price that $\beta_{\tau,\{\overline{p}^*\}} = exp(-r_1\tau)$ and $\beta_{\tau,\{\}} = exp(-r_2\tau)$ for some $r_1, r_2 > 0$. For large τ the bid price p_t^{bid} converges exponentially to $\underline{p}^*$ almost surely at the learning rate $r = min(r_1, r_2)$. They derived this for market makers sampling in calendar time. Market makers sampling tick-by-tick have the same correlation pattern, but a lower learning rate, because they miss the no-trade periods, which reveal information as well. An analogous result applies to the convergence of the ask price to $\underline{p}^*$.

Overall, transaction prices converge to the strong form efficient price in clock time at exponential rates for large τ. The following proposition summarizes the cross-correlations in Easley and O'Hara (1992)-type models. It considers only the dominant exponential learning pattern, and ignores lower order terms which disappear at faster rates as τ gets large.

Proposition 1 ***(Cross-correlations in the Easley-O'Hara model)***

The contemporaneous cross-correlation in the Easley and O'Hara (1992) model is

$$Corr\left(\Delta p_t^*, \Delta u_t\right) = -\frac{1 + e^{-r(T-1)}}{2\sqrt{K}} < 0,$$

[6] A variation of this setup is the model of Diamond and Verrecchia (1987), where short selling constraints cause periods of no trading to be a noisy signal of a low latent price.

and the cross-correlations at sufficiently large nonzero displacements follow

$$Corr\left(\Delta p^*_{t-\tau}, \Delta u_t\right) = \frac{e^r - 1}{2\sqrt{K}} e^{-r\tau} > 0,\ \forall \tau \in [1, T-1]$$

$$Corr\left(\Delta p^*_{t-T}, \Delta u_t\right) = \frac{e^{-r(T-1)}}{2\sqrt{K}} > 0,$$

where $K = K(r, T)$.

Proof: The proofs to all propositions are collected in Web Appendix A.

As before, the contemporaneous correlation is negative, and approaches its minimum for small r and small T. Furthermore, the cross-correlation of the strong form efficient price decays geometrically to zero until $\tau = T$:

$$Corr\left(\Delta p^*_{t-\tau}, \Delta u_t\right) = e^{-r(\tau-1)} Corr\left(\Delta p^*_{t-1}, \Delta u_t\right) \quad \forall \tau \in [1, T-1].$$

We graph this cross-correlation function in the first row of Figure 3. The cross-correlation pattern in the upper left panel is for a learning rate of $r = 0.5$, and in the upper right panel for a faster learning rate of $r = 2$. Often, optimal learning stops before p^*_t is reached (Aghion et al., 1991), e.g. if the spread is large or if market maker risk aversion is small. In that case the cross-correlations cut off at some $\tau < T$.

This decay pattern is not unique to the Easley and O'Hara (1992)-model. Glosten and Milgrom (1985) show more generally that if learning is costless, the expectations of market makers and traders necessarily converge as the number of trades increases. Because of the uncertainty of whether a trade reflects information or just noise, the market maker faced with a noisy signal adjusts only *partially.* Therefore, whereas the cross-correlations under a noisy signal have the same signs as under signal certainty, their absolute values are all dampened toward zero.

3.2.2 Strategic Traders

Because the market maker cannot distinguish informed from uninformed trades, informed traders can act strategically. Informed traders aim to make the signals about p^*_t conveyed by their orders as noisy as possible, while still executing the desired trades. By mimicing uninformed traders they keep the market maker unaware about the change in p^*_t. Because the market maker observes the order flow and uses it to detect informed trading, the informed

Figure 3: Cross-Correlation Functions ρ_τ of the Strong Form Efficient Price

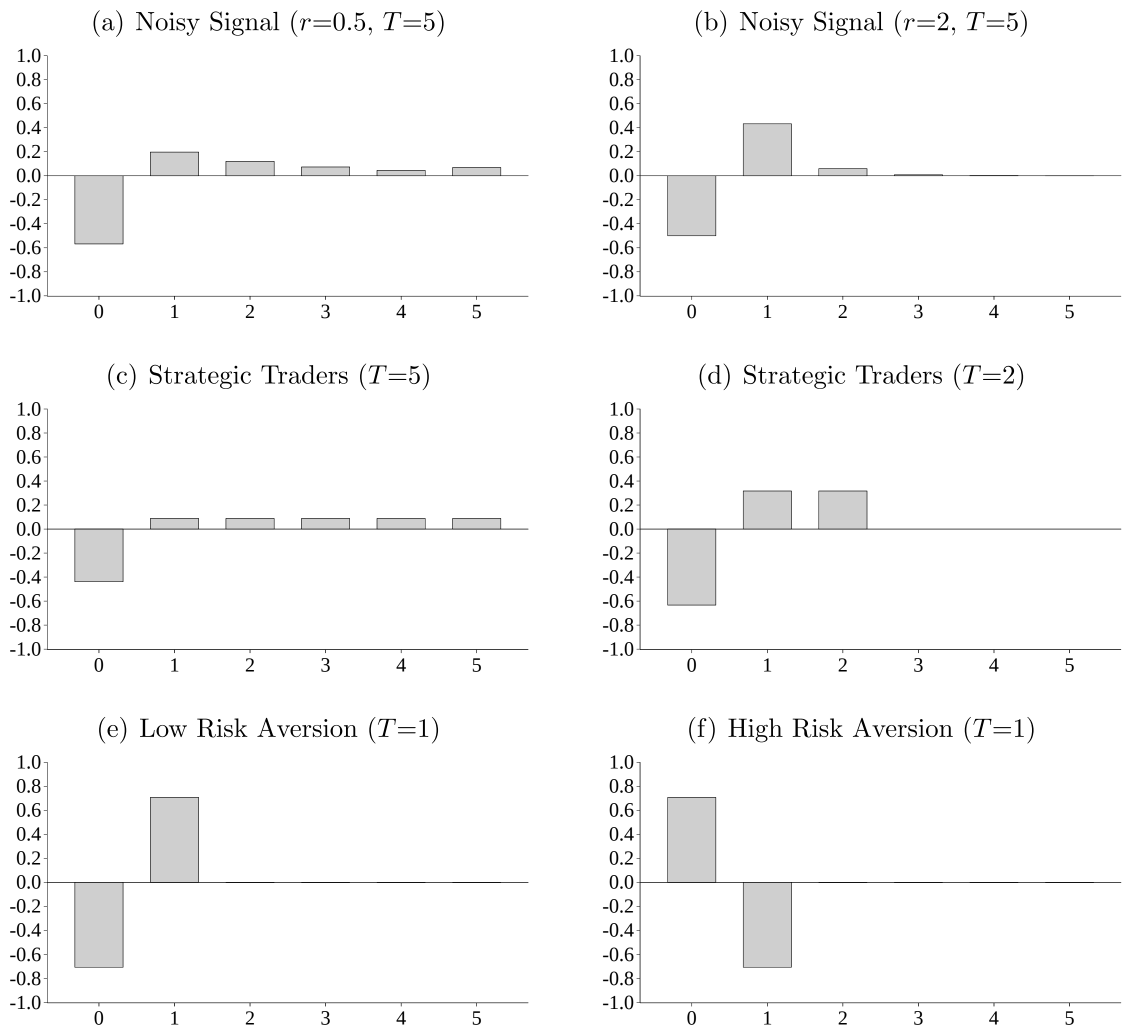

traders strategically stretch their orders over a long time period such that detecting an abnormal trading pattern is difficult. The market maker will, of course, notice the imbalance in trades over time. By sequentially updating his belief about p_t^* based on the history of trades he still learns about p_t^*, but very slowly, because of the strategic behavior of traders.

Markets of this type have been described in Kyle (1985) and Easley and O'Hara (1987). In the following we discuss the cross-correlation function implied by the Kyle (1985) model. The strategic behavior described by Kyle (1985) requires that exactly one trader is informed, or that all informed traders coordinate trading in a monopolistic manner. Here, the market maker does not maximize a particular objective function, he merely ensures market efficiency, i.e. sets the transaction price such that it equals the expected strong form efficient price, p_t^e, given the observed aggregate trading volume from informed and uninformed traders. The only optimizing agent in this model is a risk neutral, informed trader who optimally spreads his orders over the day to minimize the unfavorable price reaction of the market maker. Doing so, he maximizes his expected total daily profit using his private information and taking the price setting rule of the market maker as given. Effectively, the informed trader trades most when the sensitivity of prices to trading quantity is small.

Kyle (1985) assumes a linear reaction function of the market maker, which implies $\lambda_t = \lambda$ $\forall t \in [1, T]$, and a linear reaction function for the informed trader, which implies $q_t = q$ $\forall t \in [0, T-1]$. Under these assumptions he shows that in expectation the transaction price approaches the latent price linearly, not exponentially. The reason for this difference to the previous subsection is that there the market maker updates his beliefs in a Bayesian manner, whereas here the market maker's actions are constrained to market clearing. The other feature of strategic trading is that just before p_t^* becomes public the transaction price reflects all information.

More specifically, from the continuous auction equilibrium in Kyle (1985) the price change at time t is

$$dp^e(t) = \frac{p^* - p^e(t)}{T-t}dt + \sigma dz, \ t \in [0, T].$$

The innovation term dz is white noise with $dz \sim N(0,1)$ and reflects the price impact of uninformed traders. This stochastic differential equation has the solution[7]

$$p^e(t) = \frac{t}{T}p^* + \frac{T-t}{T}p^e(0) + (T-t)\int_0^t \frac{\sigma}{T-s}dB_s,$$

[7]The third term reflects uninformed trading. It has an expected value of zero, and the impact of this random component increases during the early trading day and decreases lateron – its contribution to $p^e(t)$ is therefore hump-shaped over time.

where $dB_s \equiv dz$. The increments of the expected price over a discrete interval of time follow therefore

$$\Delta p^e_\tau = \frac{\Delta p^*_0}{T} + (T-\tau)\int_{\tau-1}^{\tau} \frac{\sigma}{T-s} dB_s - \int_0^{\tau-1} \frac{\sigma}{T-s} dB_s. \tag{18}$$

This implies the following cross-correlations:

Proposition 2 ***(Cross-correlations in the Kyle model)***

The contemporaneous cross-correlation in Kyle (1985) is

$$Corr\left(\Delta p^*_t, \Delta u_t\right) = -\sqrt{\frac{T}{T^2+1}},$$

the cross-correlations at displacements $\tau \in [1;T]$ *are*

$$Corr\left(\Delta p^*_{t-\tau}, \Delta u_t\right) = \sqrt{\frac{1}{T(T^2+1)}},$$

and all higher order cross-correlations are zero.

The cross-covariance at nonzero displacements is a positive constant. It is positive because of market maker learning. It is constant because of the strategic behavior of traders, which spread new information equally over time. This maximizes the time it takes the market maker to include the entire strong form efficient price change in his quotes. The more periods, the more pronounced is the negative contemporaneous cross-correlation, and the smaller are the cross-correlations at nonzero displacements. We plot the cross-correlation function given by Proposition 2 in the second row of Figure 3. We show the cross-correlation function a Kyle (1985)-type model under modestly frequent changes in the latent price ($T = 5$) in the left panel, and for more frequent changes ($T = 2$) in the right panel.

Table 1 compares the cross-correlation patterns of standard multiperiod market microstructure models: The Roll (1984) model in row 1, the Glosten and Milgrom (1985) model in row 2, the Easley and O'Hara (1992) model in row 3, and the Kyle (1985) in row 4, which includes oscillating, linearly decaying and exponentially decaying patterns.

3.3 One-Period Case

In this section we return to the general latent price process, and consider the extreme case that p^*_t automatically becomes public information at the end of each period, i.e. $\omega_t =$

Table 1: Cross-Correlations between Δp_t^* and MSN in Multi-period Models

	p_t^* martingale	signal	traders strat.	ρ_0	ρ_τ $\tau \in [1, T-1]$	ρ_T	ρ_τ $\tau > T$
Roll	yes	none	n.a.	$\rho_0 < 0$	0	$-\rho_0$	0
G-M	yes	certain/ noisy	no	$\rho_0 < 0$	$\rho_{\tau-1} > \rho_\tau > 0$	$\rho_T > 0$	0
E-O	no	noisy	no	$-\frac{1+e^{-r(T-1)}}{2\sqrt{K(r,T)}}$	$\frac{-e^{-r\tau}+e^{-r(\tau-1)}}{2\sqrt{K(r,T)}}$	$\frac{e^{-r(T-1)}}{2\sqrt{K(r,T)}}$	0
Kyle	yes	noisy	yes	$-\sqrt{\frac{T}{T^2+1}}$	$\sqrt{\frac{1}{T(T^2+1)}}$	$\sqrt{\frac{1}{T(T^2+1)}}$	0

$p_{t-1}^* - \tilde{p}_{t-1}^e$ and $T = 1$. This allows us to investigate the impact of risk aversion for the cross-correlation pattern. p_{t-1}^* is thus known when the market maker decides on p_t, which removes any incentive for informed traders to behave strategically. They therefore react immediately, which implies that $E(q_{t-\tau}\varepsilon_t) = 0 \;\; \forall \tau \neq 0$ and that all trades are serially uncorrelated, i.e. $E(q_t|q_{t-1}) = 0$. For the market maker all periods are identical, and therefore the spread and reaction parameters are both constant over time, i.e. $s_t = s$ and $\lambda_t = \lambda \;\; \forall t$.

The cross-correlation function inherits its shape from (13)–(16). At displacement one it has the opposite sign and same absolute value as contemporaneously, and it is zero at displacements larger than one. In order to pin down the value of the contemporaneous cross-correlation, we now turn to specific models.

3.3.1 No Market Maker Information

We start with our baseline assumption that the market maker at time t has no information whatsoever about Δp_t^*. Plugging $T = 1$, $s_t = s$, and $\lambda_t = \lambda$, and thus $\phi_t = \phi$, into the general multiperiod results of Section 3.1 gives

Proposition 3 ***(Strong form cross-correlation, one period model)***

$$Corr(\Delta p_t^*, \Delta u_t) = \frac{1}{\sqrt{2}} \frac{sE(q_t\varepsilon_t) - \sigma}{\sqrt{\phi s^2 + \sigma^2 - 2s\sigma E(q_t\varepsilon_t)}}, \tag{19}$$

$$Corr(\Delta p_{t-1}^*, \Delta u_t) = -Corr(\Delta p_t^*, \Delta u_t).$$

If there is trading in every period ($\beta = 1$, and thus $\phi = 1$), then the cross-correlation (19) is bounded from above and below by

Proposition 4 ***(Bounds of contemporaneous cross-correlation)***

$$-\frac{1}{\sqrt{2}} \leq Corr(\Delta p_t^*, \Delta u_t) \leq 0.$$

The cross-correlation reaches the lower bound for zero spread. Thus for midprices, or extremely small spreads due to order splitting, the cross-correlation is highest. For transaction prices the contemporaneous cross-correlation is less pronounced. The contemporaneous cross-correlation for midprices is negative, because p_t^e does not react instantaneously to the change in the strong form efficient price in the same period. This is an instance of the price stickiness that Bandi and Russell (2006) show to generate "mechanically" a negative contemporaneous cross-correlation. It differs from negative unity because transaction prices move in adjustment to the strong form efficient return one period earlier.

Table 2: Cross-Correlations between Latent Prices and MSN in One-period Models

latent price	s	λ	loss function	ρ_0	ρ_1	ρ_τ $\tau > 1$
p_t^*	0	any	any	$-\frac{1}{\sqrt{2}}$	$\frac{1}{\sqrt{2}}$	0
	≥ 0	any	any	$-\frac{1}{\sqrt{2}} \leq \rho_0 < 0$	$-\rho_0$	0
	≥ 0	any	high $n +$ extra info	$\rho_0 > 0$	$-\rho_0$	0
$\tilde{p}_t^e$	≥ 0	λ^{opt}	quadratic	$-\frac{1}{\sqrt{2}} \leq \rho_0 \leq \frac{1}{\sqrt{2}}$	$-\rho_0$	0
	$\in [0, \lambda[$	$> \frac{\lambda^{opt}}{2}$	any	$\rho_0 < 0$	$\rho_1 > 0$	0
	$\in [0, \lambda[$	$< \frac{\lambda^{opt}}{2}$	any	$\rho_0 > 0$	$\rho_1 > 0$	0
	λ	any	any	0	0	0
	$\geq \lambda$	$> \frac{\lambda^{opt}}{2}$	any	$\rho_0 > 0$	$\rho_1 < 0$	0
	$\geq \lambda$	$< \frac{\lambda^{opt}}{2}$	any	$\rho_0 < 0$	$\rho_1 < 0$	0

We summarize these results in the upper two rows of Table 2. Compared to the multi-period case in Table 1 the absolute value of the cross-correlation at lag one is large, because all information is revealed. Cross-correlations at any displacement beyond one are, in contrast, necessarily all zero.

3.3.2 Incomplete Market Maker Information and Risk Aversion

Throughout this paper we assume a risk-neutral market maker. In this subsection we lift this assumption, which can be justified in times of market turbulence. If extreme events occur, strong form efficient prices become highly correlated across assets, or, to stay with our maintained example, stocks. Although the market maker is bound by his quote only up to a fixed quantity on an individual stock, the total exposure of a market maker that has quotes outstanding in many markets might be non-trivial.

Without information about Δp_t^* risk aversion does not change the market maker behavior. Extra information, however, e.g. about the direction of the change in the latent price, $\{\text{sgn}(\varepsilon_t)\}$, can under risk aversion invert the cross-correlation pattern. Knowing $\{\text{sgn}(\varepsilon_t)\}$ the market maker adjusts his quotes before informed traders can take advantage of the latent price change. The market maker updates his prior about p_t^*, summarized by the distribution $p_t^* \sim f(p_{t-1}^*, \sigma^2)$, with the signal $\{\text{sgn}(\varepsilon_t)\}$. For convenience of exposition we use

Assumption 2 *The probability density function of ε_t is symmetric around its zero mean, monotonically increasing on $]-\infty; 0]$ and monotonically decreasing on $[0; \infty[$.*

The updated belief $\tilde{f}(\cdot)$ differs from $f(\cdot)$ in that it is truncated from below or above at $p_t^* = p_{t-1}^*$ when $\text{sgn}(\varepsilon_t) > 0$ or $\text{sgn}(\varepsilon_t) < 0$, respectively. After observing signal and p_{t-1}^*, the market maker quotes a bid and an ask price for the following period, taking the spread s as given:

$$p_t = p_{t-1}^* + sq_t + R(\{\text{sgn}(\varepsilon_t)\}). \tag{20}$$

This equation resembles (6), with $\omega_t = -\tilde{p}_{t-1}^e + p_{t-1}^* + R(\{\text{sgn}(\varepsilon_t)\})$. The market maker response $R(\cdot)$ to the extra information depends in particular on the market maker's risk aversion, n.

An approximation[8] to the problem of choosing $p_t^e(n)$ based on loss function (12) is

$$p^e(n) = \underset{x \in [\underline{p}^*, \overline{p}^*]}{\text{argmax}} - \int_{\underline{p}^*}^{x} (x - p^*)^n f(p^*) dp^* - \int_{x}^{\overline{p}^*} (p^* - x)^n f(p^*) dp^*. \tag{21}$$

The higher the risk aversion n, the more sensitive is the expected loss, $L_n \left[p_t, F(\cdot, \underline{p}^*, \overline{p}^*)\right]$, to the support of p_t^*, that is, to $\underline{p}^*$ and $\overline{p}^*$. For some values of n, explicit solutions to (21)

[8] This approximation is exact for $s = 0$ or, more generally, for

$$\int_{p^e(n)-s}^{p^e(n)} (p^e(n) - p^*)^n f(p^*) dp^* + \int_{p^e(n)}^{p^e(n)+s} (p^* - p^e(n))^n f(p^*) dp^* = 0.$$

are available. A well-known result is that the optimal choice for a risk neutral market maker ($n = 1$) is to set p_t^e equal to the median of $f(\cdot)$, and for a modestly risk averse market maker ($n = 2$) to the mean. An extremely risk averse ($n \to \infty$) market maker follows the most robust pricing role possible: He minimizes his expected loss at the price in the middle of the support of $f(\cdot)$, i.e. $p_t = \frac{\underline{p}^* + \overline{p}^*}{2}$. We summarize this in

Proposition 5 ***(Optimal Midprice)*** *The optimal midprice, $p^e(n)$, monotonically shifts from the median to the midpoint of the support of p_t^* with increasing risk aversion. In particular,*

$$
\begin{aligned}
p^e(1) &= \text{Median}(p_t^*) \\
p^e(2) &= \text{E}(p_t^*) \\
p^e(\infty) &= \text{Midsupport}(p_t^*).
\end{aligned}
$$

Figure 4, which plots the transaction price as a function of risk aversion n, illustrates this increasing sensitivity. For a right-skewed distribution $f(\cdot)$ with infinite support, namely the halfnormal distribution, $p^e(n)$ increases in n, starting from the median for $n = 1$, monotonically without bound. If, in contrast, $f(\cdot)$ has finite support, then $p^e(n)$ increases from the median monotonically toward a finite asymptote $p^e(\infty)$. This is shown in the right panel of Figure 4 for the right-triangular distribution defined on $[0, 1]$. For left-skewed distributions the result is analogous. This has implications for the possible cross-correlations:

Proposition 6 ***(Cross-correlation under market maker information)*** *If the distribution of the expected latent price with ex-ante support $[\underline{p}_t^*, \overline{p}_t^*]$ satisfies*

$$
\left[\frac{\underline{p}_t^* + \overline{p}_t^*}{2} - p_{t-1}^*\right] \text{sgn}(\varepsilon_t) > s + \frac{\sigma}{E(|\varepsilon_t|)}, \tag{22}
$$

then $\exists n_0 > 1$ such that $\forall n > n_0$ it holds that $Corr(\Delta p_t^*, \Delta u_t) > 0$.

Condition (22) holds, for example, for normally distributed, but not for tent distributed Δp_t^*. This is reflected in Figure 4, where the price in the left panel quickly reaches the cutoff $\frac{\sigma}{E(|\varepsilon|)}$, plotted as dashed line, whereas in the right panel it never does.

Comparing these results in the third row of Table 2 with the other models, it appears that even though the contemporaneous cross-correlation can be positive for high risk aversion levels, the usual case is that it is negative. For the halfnormal distribution, for example, we need a rather high risk aversion of $n \geq 8$. Nevertheless, changes in risk aversion of the

Figure 4: Optimal Mid-Price for Right-Skewed Expected Latent Price Distributions

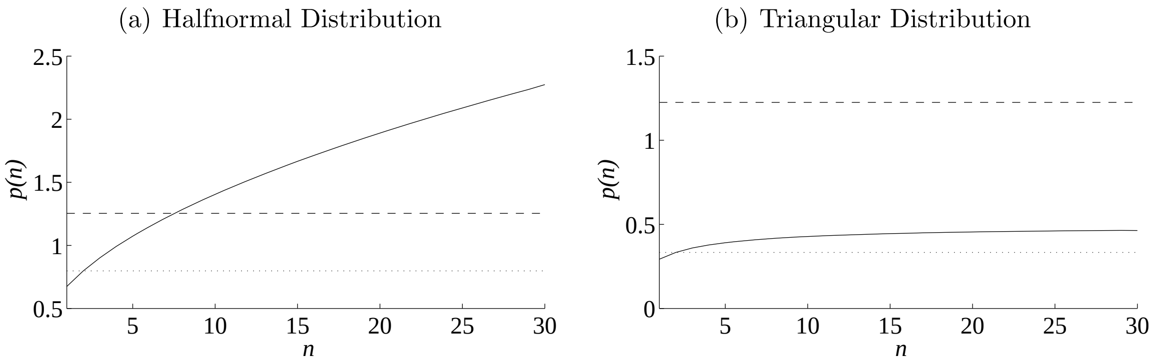

market maker have a distinctive impact on the cross-correlation. Hansen and Lunde (2006) note as their "Fact IV" that "the properties of the noise have changed over time." Because they base this observation on a comparison of year 2000 with year 2004 it is possible that the underlying cause is a change in risk aversion.

The link between properties of noise and risk aversion offers itself as a way to estimate the time path of risk aversion from the cross-correlation pattern of transaction prices. In stable periods with low risk aversion the contemporaneous cross-correlation is negative, but as uncertainty shoots up, contemporaneous cross-correlation shoots up with it. In periods of crisis this can lead to the extreme case of an inverted cross-correlation pattern that we have described in this section. The lower row of Figure 3 illustrates this inversion: it shows the typical cross-correlation pattern of strong form efficient prices in a one-period model with modest risk aversion on the left, and under higher risk aversion on the right.

In summary we have shown in this section that many market properties leave their mark on the cross-correlation pattern: The displacement beyond which correlation is zero gives an indication of the frequency of information events. The larger the correlation is in absolute value terms the fewer uninformed trades occur in the market. If contemporaneous strong form cross-correlation is positive, then market makers are very risk averse and have access to extra information. If the cross-correlations at nonzero displacements decay quickly, then market makers learn fast. If they do not decay at all, then informed traders act strategically.

4 Return-Noise Correlations in Financial Economic Environments II: Semi-Strong Efficient Prices

Now we base the cross-correlation calculation on another latent price, the semi-strong form efficient price, $\tilde{p}_t^e$. Equivalently this setup can be seen as an endogenous latent price process, determined by an exogenous trading process q_t, because then the strong-form efficient price remains unobserved and enters the model only via the informed trades. It is closely related to the "generalized Roll model" in Hasbrouck (2007). To keep the terms manageable, we assume no extra information here, i.e. $\omega_t = 0 \; \forall t$.

4.1 Multi-Period Case

Simple calculations (see the Web Appendix A.0.2) give for the contemporaneous covariance of semi-strong form efficient prices

$$\begin{aligned} Cov(\Delta\tilde{p}_t^e, \Delta\tilde{u}_t) &= \frac{1}{T}\{-\phi_0\lambda_0(\lambda_0 - s_0) + \sigma(\lambda_{-1} - s_{-1})E(q_{-1}\varepsilon_{-T}) \\ &- \sum_{i=-1}^{-T}(\lambda_{-1} - s_{-1})\,\lambda_i E(q_i q_{-1}) \\ &+ \sum_{i=1}^{T-1}(-\phi_i\lambda_i(\lambda_i - s_i) + \lambda_i(\lambda_{i-1} - s_{i-1})E(q_i q_{i-1}))\Big\}, \end{aligned} \tag{23}$$

for covariance at higher displacements $\tau \in [1, T-1]$

$$\begin{aligned} Cov(\Delta\tilde{p}_{t-\tau}^e, \Delta\tilde{u}_t) &= \frac{1}{T}\{-\lambda_0(\lambda_\tau - s_\tau)E(q_0 q_\tau) \\ &+ \lambda_0(\lambda_{\tau-1} - s_{\tau-1})E(q_0 q_{\tau-1}) + \lambda_{T-\tau}\,(\lambda_{T-1} - s_{T-1})\,E(q_{T-\tau}q_{T-1}) \\ &+ \sum_{i=\tau+1}^{T-1}[\lambda_{i-\tau}(-\lambda_i + s_i)E(q_{i-\tau}q_i) + \lambda_{i-\tau}(\lambda_{i-1} - s_{i-1})E(q_{i-\tau}q_{i-1})]\Big\}, \end{aligned} \tag{24}$$

for covariance at displacement T

$$Cov(\Delta\tilde{p}_{t-T}^e, \Delta\tilde{u}_t) = \frac{1}{T}\lambda_0\,(\lambda_{T-1} - s_{T-1})\,E(q_0 q_{T-1}), \tag{25}$$

and for all higher order displacements $\tau > T$

$$Cov(\Delta\tilde{p}_{t-\tau}^e, \Delta\tilde{u}_t) = 0.$$

The cross-correlations for semi-strong form efficient prices stem from a gap between the spread, s_t, and the adverse selection parameter, λ_t. Such a gap can result from processing costs ($s_t > \lambda_t$), from legal restrictions ($s_t < \lambda_t$), or merely from suboptimal behavior of the market maker. Noisy signals or strategic behavior do not affect the semi-strong cross-correlations, as for example in Easley and O'Hara (1992), where prices are semi-strong form efficient by definition. Under semi-strong market efficiency ($s_t = \lambda_t \ \ \forall t$) the cross-correlation function is zero for all displacements.

The Kyle (1985) model assumptions $\lambda_t = \lambda$ and $s_t = s \ \forall t$ give with (24)

$$Cov(\Delta\tilde{p}^e_{t-\tau}, \Delta\tilde{u}_t) = \frac{\lambda(\lambda - s)}{T}\left\{E(q_{T-\tau}q_{T-1}) + \sum_{i=\tau}^{T-1}\left[E(q_{i-\tau}q_{i-1}) - E(q_{i-\tau}q_i)\right]\right\}.$$

If $\lambda = 0$, then this cross-correlation is flat at zero. Likewise, if additionally $E(q_{i-\tau}q_i)$ is a positive constant between the time of the latent price change and its public announcement, the cross-correlation is flat and proportional to $\frac{\lambda(\lambda-s)}{T}$. If $E(q_iq_j) > E(q_{i-\tau}q_j) > 0 \ \forall i \leq j$, $\forall \tau > 0$, the cross-correlation decreases in τ.

4.2 One-Period Case

The simpler case of markets in which all information is revealed after one period without any extra information, i.e.

$$\Delta\tilde{p}^e_t = \lambda(q_t - q_{t-1}) + \sigma\varepsilon_{t-1}, \tag{26}$$

$$\Delta\tilde{u}_t = (s - \lambda)(q_t - q_{t-1}). \tag{27}$$

offers itself again for illustration of these cross-correlation effects. Unlike their strong form counterpart the semi-strong form efficient prices are not a martingale. We see in the following proposition that in contrast to the strong form correlations, the absolute value of semi-strong form cross-correlation at displacement zero and one usually differs even in one-period models.

Proposition 7 ***(Semi-strong form cross correlation, one-period model)***

The contemporaneous cross-correlation is

$$Corr(\Delta\tilde{p}^e_t, \Delta\tilde{u}_t) = \frac{2\phi\lambda - \sigma E(q_t\varepsilon_t)}{\sqrt{\sigma^2 - 2\sigma\lambda E(q_t\varepsilon_t) + 2\phi\lambda^2}}\frac{\operatorname{sgn}(s-\lambda)}{\sqrt{2\phi}}.$$

The cross-correlation at displacement one equals

$$Corr(\Delta\tilde{p}^e_{t-1}, \Delta\tilde{u}_t) = \frac{-\phi\lambda}{\sqrt{\sigma^2 - 2\sigma\lambda E(q_t\varepsilon_t) + 2\phi\lambda^2}} \frac{\text{sgn}(s-\lambda)}{\sqrt{2\phi}}.$$

All cross-correlations at higher displacements are zero.

Bounds on the contemporaneous cross-correlation can be obtained by assuming a specific market marker loss function and then solving for the market maker's optimal λ. For example, suppose the market maker has a quadratic loss function, then

$$\lambda^{opt} = \underset{\lambda}{\text{argmin}}\, E\left[(\tilde{p}^e_t - p^*_t)^2\right],$$

which becomes

$$\lambda^{opt} = \underset{\lambda}{\text{argmin}}\, \phi\lambda^2 - 2\sigma\lambda E\left(q_t\varepsilon_t\right),$$

and therefore $\lambda^{opt} = \frac{\sigma}{\phi}E\left(q_t\varepsilon_t\right) > 0$. At λ^{opt} we have

$$Corr(\Delta\tilde{p}^e_t, \Delta\tilde{u}_t) = E(q_t\varepsilon_t)\frac{\text{sgn}(s-\lambda^{opt})}{\sqrt{2\phi}},$$

$$Corr(\Delta\tilde{p}^e_{t-1}, \Delta\tilde{u}_t) = -E(q_t\varepsilon_t)\frac{\text{sgn}(s-\lambda^{opt})}{\sqrt{2\phi}},$$

and because $0 \leq E\left(q_t\varepsilon_{t-\tau}\right) < 1,\ \forall t, \tau$[9]

$$\left|Corr(\Delta\tilde{p}^e_t, \Delta\tilde{u}_t)\right| = \left|Corr(\Delta\tilde{p}^e_{t-1}, \Delta\tilde{u}_t)\right| \leq \frac{1}{\sqrt{2\phi}}.$$

Under a quadratic market maker loss function and an uninterrupted flow of trades ($\phi = 1$), the absolute value of cross-correlations is bounded from above by $\frac{1}{\sqrt{2}}$.

The contemporaneous cross-correlation is positive as in Diebold (2006) for $s > \lambda > \frac{\sigma}{2\phi}E(q_t\varepsilon_t) = \frac{\lambda_{opt}}{2}$ and for $s < \lambda < \frac{\lambda_{opt}}{2}$. Proposition 7 shows that the size of the spread matters only relative to the adverse selection parameter. The cross-correlation at displacement one, for example, is negative if and only if the spread exceeds the adverse selection cost. For these parameters again an inverted (compared to Hansen and Lunde (2006)) cross-correlation function obtains as in the lower right panel of Figure 3. Either parametrization reflects a plausible market situation. A large spread scenario without violating the market maker's

[9]Note that by Jensen's inequality $0 \leq E\left(q_t\varepsilon_{t-\tau}\right) < E\left(|\varepsilon_{t-\tau}|\right) < \sqrt{E\left(|\varepsilon_{t-\tau}|^2\right)} = 1$.

zero-profit condition can be the result of high risk aversion. By the same reasoning as in Section 3.3.2, there exists a risk aversion level n_0 such that all $n > n_0$ generate a spread $s > \lambda$. Whereas the spread is likely to exceed the trade response, because the spread must cover the order processing cost, also the small spread scenario could obtain in some markets from competition or regulatory constraints.

We summarize the results in the lower four rows of Table 2. Unlike for the strong-form efficient prices, positive contemporaneous cross-correlation for semi-strong form efficient prices obtains even in situations where the market maker does not observe a signal.

Summing up, what sign of contemporaneous cross-correlation does market microstructure theory predict? Positive contemporaneous cross-correlations occur for (1) strong form efficient prices under sufficiently high risk aversion if a signal is observed, and (2) semi-strong form efficient prices for several parameterizations. Bandi and Russell (2006) and Diebold (2006) rightly wonder whether a negative cross-correlation is inevitable. We have seen that for latent price processes different from Brownian motion a positive cross-correlation is not unlikely. For strong form efficient prices a positive cross-correlation is possible, but a negative cross-correlation appears most realistic. Markets in which Bandi and Russell (2008) find no "obvious evidence of a significant, negative correlation," are likely subject to an extraordinary microstructure effect such as high risk aversion.

5 The Relationship between Price Change Frequency and Sampling Frequency

In this section we discuss the implications that the frequency of price changes in financial markets has for the choice of sampling frequency. We begin with a discussion of the effects of incompletely observed latent price changes, turn then to the effect of sampling frequency on return-noise correlations, and finally examine the implications of trade frequency for econometric theory.

5.1 Infrequent Latent Price Disclosure

For clarity of exposition in most of this paper we discuss models, where p^*_{t-1} becomes public information just before it changes. In general, however, its exact value might never become public. In this case, because $Corr(p^*_t, \Delta p^*_{t-\tau}) > 0 \ \forall \tau > 0$, the past $p^*_{t-\tau}$ contain unrevealed information about p^*_t. As $p^*_{t-\tau}$ is not precisely known itself, potentially the entire history of observed transaction prices contains information about the current p_t^*.

More specifically, suppose that exact values of the κ most recent latent prices are not fully revealed and therefore partly private information. This changes the market maker's problem in two ways: First, informed trades now convey the signal $\{\text{sgn}(p_t^* - p_t)\}$, distinct from the signal $\{\text{sgn}(\varepsilon_t)\}$. Second, the larger κ, the more spread out is ceteris paribus the distribution of the market maker's belief about p_t^*.

This signal conveyed by a trade mixes information on Δp_t^* with the κ previous latent price changes, Δp_{t-iT}^*, $i \in [1, \kappa]$. The contemporaneous cross-correlation is dampened toward zero, because the covariance between latent prices and noise does not offset the higher noise variance. A potentially wider spread dampens the cross-correlation further.

By (11) the signs of the cross-correlations at displacements $\tau > 0$ remain unchanged as long as learning induces p_t to move in the same direction as $p_{t-\tau}^*$. They are the closer to zero, the less informative the signal conveyed by the current trade is about past latent returns. That is, the closer to zero the cross-correlation $Corr(p_t, \Delta p_{t-\tau}^*)$ is, and ultimately, the more often p_t^* changes during the period.

Overall, slowly decaying private information keeps the cross-correlation sign pattern unchanged, but dampens its absolute values toward zero.

5.2 Sampling Frequency and Return-Noise Correlations

We have so far assumed that p_t, $\tilde{p}_t^e$ and p_t^e are all updated at the same frequency and chose this as our sampling frequency. Sampling at faster or slower rates will affect the shape of cross-correlation functions. Because for example the reaction speed of the market maker is generally unknown, econometric sampling may proceed at faster or slower rates. This has immediate implications for the shape of empirically estimated (sample) cross-correlation functions. Clearly, the cross-correlations are the smaller in absolute value, the more variation from other periods increases the variance of Δu_t without increasing its covariance with $\Delta p_{t-\tau}^*$.

Consider first the effects of sampling "too fast", in particular more frequently than trades occur. Suppose we sample m times during an interval of no changes in market prices, and for that matter, latent prices. Recording each time the most recent observed price, all returns but the first in each interval are zero and thus the cross-correlation function becomes a spread-out version of the cross-correlation functions derived in the previous sections: after each dampened non-zero cross-correlation follow $m - 1$ zero cross-correlations. Zeros in the middle of a cross-correlation function thus indicate overly fast sampling.

A variant of sampling "too fast" is sampling faster than information evolves. That is, sampling at trading frequency, i.e. the frequency of p_t, although the market maker updates p_t^e only infrequently, for example only every m-th trade. A single change of p_{im}^e ($i \in \mathbb{N}$) now reflects the information about Δp_0^* conveyed by trading activity between $(i-1)m$ and im. Δp_m is thus more correlated with Δp_0^* than under period-by-period updating. But because the quote is fixed during $(i-1)m+1$ and im, the trades in the interim period jointly provide less information than under period-by-period updating. Because further the variance of noise increases due to the delayed accumulated market maker response, the cross-correlation function oscillates between dampened values.

Now consider the effects of sampling "too slowly". Suppose, for example, that we sample in the one-period model of Section 4.2 only every m-th tick, where $\hat{t}$ indexes the m-tick blocks. Then (26) becomes

$$\Delta \tilde{p}_{\hat{t}}^e = \sum_{i=(\hat{t}-1)m+1}^{\hat{t}m} \Delta \tilde{p}_i^e = \lambda(q_{\hat{t}m} - q_{(\hat{t}-1)m}) + \sigma \sum_{i=(\hat{t}-1)m}^{\hat{t}m-1} \varepsilon_i,$$

and the variance increases to $Var(\Delta \tilde{p}_{\hat{t}}^e) = m\sigma^2 - 2\sigma\lambda E(q_t \varepsilon_t) + 2\phi\lambda^2$. Assuming that the statistical properties of the interim periods are the same as the properties of the sampled periods, the expressions for noise (27), its variance $Var(\Delta u_{\hat{t}})$, and the covariance $Cov(\Delta \tilde{p}_{\hat{t}}^e, \Delta u_{\hat{t}})$ remain unchanged. Increasing the sampling interval averages the initial transaction price reaction with later price changes, thereby dampening the entire cross-correlation pattern toward zero:

$$\left|Corr(\Delta \tilde{p}_{\hat{t}}^e, \Delta \tilde{u}_{\hat{t}})\right| = \left| \frac{2\phi\lambda - \sigma E(q_t \varepsilon_t)}{\sqrt{2\phi}\sqrt{m\sigma^2 - 2\sigma\lambda E(q_t\varepsilon_t) + 2\phi\lambda^2}} \right| < \left|Corr(\Delta \tilde{p}_t^e, \Delta u_t)\right|.$$

Hansen and Lunde (2006) find a negative contemporaneous cross-correlation between returns and noise, which diminishes as more ticks are combined into one transaction price sample. Our results show that this can stem from two different sources: Either from the averaging effect across latent price changes just described, or from cross-correlations at nonzero displacements offsetting the contemporaneous correlation for the same latent price change. This ambiguity can be resolved by evaluating the entire cross-correlation function, which shows the importance of not limiting noise analysis to the contemporaneous cross-correlation.

Standard RV is unbiased if sampling frequency is sufficiently low so that microstructure effects are averaged out. Applying "noise-corrected" RV estimators to data at lower frequencies results in biased estimates, because at lower frequencies slow moving features of the price process are removed, not microstructure noise. Thus they should only be applied to data sampled at frequencies at which microstructure effects can conceivably exist, e.g. above 1/100 seconds.

The upshot is that sampling frequency does not change the sign pattern of cross-correlations but can severely impact their absolute values. Our results suggest that sampling at a rate detached from the updating frequency of prices and information, in particular sampling too fast or too slow, mutes complications as well as information originating from dependent noise, and effectively changes the properties of the data. Sampling frequency should therefore be chosen based on the price updating frequency of the market.

5.3 Sampling Frequency and Asymptotic Theory

The previous section has shown that the microstructure of a market implies a natural sampling frequency. In practice, sampling frequency is also central for econometric theory. Infill asymptotic theory, for example, requires the number of sampling intervals during a fixed time span to go to infinity. Sampling at an infinite frequency is impossible in real financial markets, but as trading keeps becoming faster and faster we can view it as the trading frequency limit in the infinite future. Can econometric theory gain anything from examining the developments in financial markets?

Consider the Zhou (1996)-estimator as an example. Its consistency hinges on the ratio of the lag length measured by the number of sample periods to sampling frequency going to zero as sampling becomes infinitely frequent. That is, under infill asymptotics, the time span that the lag window spans must asymptotically shrink to zero. It is commonly argued that this assumption is "inappropriate" for financial markets (e.g. Hansen and Lunde, 2006, p.139). Effectively, the question comes down to whether MSN decays according to a tick-time or a calendar-time schedule. Linking econometrics to market structure, we argue in the following that tick-time dependence is reasonable in many cases.

When deriving the limiting behavior of IV estimators, econometric theory commonly assumes that the properties of transaction prices are invariant to the sampling frequency. This might be correct in many instances, but just as often it is not. In the case of financial markets, the maximum feasible sampling frequency is dictated by the trading frequency. As the trading frequency in a given market changes, other features of that market change as well. Therefore asymptotic theory must account for the possibility that price behavior changes as feasible sampling frequency increases.

To verify the relevance of this possibility, let us revisit the economics of financial markets. The analogue of shrinking the interval length in infill asymptotics is a higher trading frequency in financial markets, which implies a higher feasible sampling frequency. In the following three examples, we examine how a higher feasible sampling frequency affects noise persistence. We consider a slow and a fast market: The slow market is rather illiquid, so that a trade is observed only once during a five-minute interval. The fast market is more liquid, and trades are observed once every minute. The latent price process is the same in both markets. In fact, both slow and fast market might be the very same market at different points in time. The latent price moves more between two trades in the slow market, which means that there the IV over the shortest possible sampling interval is higher.

Consider first a bid-ask bounce. Bid-ask bounces are purely mechanic, and directly linked to observed trades. In the slow market, the possible rebounce occurs five minutes after the original trade, whereas in the fast market it occurs after only one minute. Thus the market microstructure noise (MSN) is autocorrelated for five minutes in the slow market, but only for one minute in the fast market.

Next, consider asymmetric information. If learning of market participants is automated and limited to information extracted from trade signals, then the amount of learning grows in the number of trade signals observed, not in the time that has passed. For a specific example, suppose the market maker needs ten trades to include half of the latent price change into his price quote. This will take 50 minutes in the slow, but only ten minutes in the fast market. MSN persistence measured in calendar time is thus much shorter in the faster market.

Our third example shows that this applies only to tick-dependent MSN, i.e. to situations where private information is revealed by trades only and where the speed of information processing is not a binding constraint. Some properties of MSN, however, might be invariant to sampling frequency. For example, the time that strategic informed traders allocate to fully reveal their information might be exogenous to the trading frequency. Instead, its optimal value might be a function of the speed of information diffusion outside the market, e.g. due to reporting delays, which are fixed in calendar time. Thus the autocorrelation of MSN generated by strategic informed traders is the same in calendar time in the slow and the fast market; it does not shrink as sampling frequency increases.

Overall, the autocorrelation of MSN to due a bid-ask bounce and asymmetric information without strategic traders shrinks in calendar time as the feasible sampling frequency increases. The autocorrelation of MSN due to strategic traders does not.

This has an important implication for the asymptotic theory of IV estimators of the Zhou (1996)- and Hansen and Lunde (2006)-type. When private information is revealed by trades only, the necessary lag length is fixed in terms of ticks, not calendar time. Therefore, the ratio of lag length to sampling frequency approaches zero when sampling infinitely fast. In these cases the estimators are consistent. They must be modified to ensure consistency when relevant information transmission occurs outside of the financial market, e.g. by subsampling (Barndorff-Nielsen et al., 2011b) or kernel-based downweighting of higher-order autocovariances (Barndorff-Nielsen et al., 2008).

6 Practical Implications and Empirical Application

We have already drawn some econometric implications insofar as we have shown that market microstructure models predict rich cross-correlation patterns between latent prices and market microstructure noise (MSN), which have yet to be investigated empirically. Here we go farther, sketching some specific aspects of such empirics, including strategies for using microstructural information to obtain improved "structural" volatility estimators, and comparative aspects of structural and non-structural volatility estimators. We apply our methodology to the stock and the oil futures market.

6.1 Structural Volatility Estimation via Microstructural Restrictions

In the introduction we highlighted the key issue of estimation of integrated volatility (IV) using high-frequency data, the potential problems of the first-generation estimator (simple realized volatility – RV) in the presence of MSN, and subsequent attempts to "correct" for MSN.

In an important development, Barndorff-Nielsen et al. (2008) suggest making RV robust to serial correlation via realized kernel estimation methods, which are asymptotically justified under very general conditions. That asymptotic generality is, however, not necessarily helpful in finite samples. Indeed the frequently unsatisfactory finite-sample performance of nonparametric HAC estimators leads Bandi and Russell (2011) to suggest sophisticated alternative statistical approaches.

Here we explore aspects of a different approach that *specializes* the estimator in accordance with the implications of market microstructure theory. We follow the idea of Aït-Sahalia, Mykland, and Zhang (2005) of modeling MSN explicitly in a fully parametric framework, which makes sampling as often as possible optimal. No claim is made about optimality; instead we show the practical relevance of tailoring the estimator to the market at hand.

Consider strong form noise given by (3), so that $\Delta p_t = \Delta p_t^* + \Delta u_t$. Then we have, absent insider information, using the notation $\gamma_i \equiv E(\Delta p_t \Delta p_{t-i})$ and $RV \equiv \gamma_0$, that the variance of strong form efficient returns (7) is

$$\sigma^2 = RV + 2\sum_{i=1}^{k} \gamma_i - 2E(u_t \Delta u_{t-k}) - 2E(\Delta p_t^* u_{t+k}). \tag{28}$$

Proof: See Web Appendix B.1.

If MSN is asymptotically uncorrelated, i.e. if $\lim_{k\to\infty} E(u_t \Delta u_{t-k}) = 0$ and $\lim_{k\to\infty} E(\Delta p_t^* u_{t+k}) = 0$, then Equation (28) simplifies to

$$\sigma^2 = RV + 2\sum_{i=1}^{\infty} \gamma_i. \tag{29}$$

This is equivalent to the constant realized kernel estimator discussed in Hansen and Lunde (2006). Without insight in the market microstructure all higher order autocovariances are potentially important. Empirically most will be noisy estimates of zero (Barndorff-Nielsen et al., 2008). Without insights in what patterns in transaction prices are caused by MSN, a noise correction like (29) will remove all. But actual transaction prices consist not only of a martingale strong form efficient price plus MSN, but also of other disturbances of unknown form. These other disturbances might not be part of any microstructure model. In fact, their existence might not even be known. Lacking better knowledge by any market participant, these must be considered risk, and therefore be part of the volatility estimate of the latent price. A noise correction as Equation (29) "corrects" price features that are not MSN, but an essential part of the volatility of the latent price process.

The key point we stress in this paper is that it is indispensable to sort out the market microstructure before choosing a noise correction. This applies no matter whether MSN is dependent on the latent price or not.

In the following we consider ten potential sources of MSN, five of which are independent, and five are dependent on the latent price. We start with a discussion of two examples of parsimonious noise-robust estimators for realized volatility, both of which are special cases of (29), before providing an overview of estimators in Table 3.

Consider first a "bid-ask bounce estimator", based on a one-period model without extra information and constant spread. From (3), (5) and (6) we obtain $\Delta u_t = \sigma(\varepsilon_{t-1} - \varepsilon_t) +$

$s(q_t - q_{t-1})$, and this implies a variance of strong form efficient returns of

$$E\left[(\Delta p_t^*)^2\right] = E\left[(\Delta p_t - \Delta u_t)^2\right] = E\left(\Delta p_t^2\right) + 2s\left[\sigma E(q_t\varepsilon_t) - \phi s\right].$$

Simple calculations reveal that the last term equals twice the first-order autocorrelation of market returns, so that, even if $E(q_t\varepsilon_t) \neq 0$, an unbiased estimator for $IV = \sigma^2$ is[10]

$$\widehat{IV} = RV + 2\gamma_1. \tag{30}$$

It is interesting to note the resemblance to estimators of Roll (1984), based on standard asymptotic theory, and Zhou (1996), based on infill asymptotic theory.

As another example, consider an estimator for a market with nonstrategic incompletely informed traders. Absent any exogenous noise, the transaction price follows an MA(∞) process in the innovations of the latent price:

$$\Delta p_t = (\beta + \sigma)\varepsilon_t + \beta(\alpha - 1)\sum_{i=0}^{\infty} \alpha^i \varepsilon_{t-i-1} \tag{31}$$

This parsimonious form of Δp_t accommodates very persistent cross-correlations, similar to the idea behind the examples in Oomen (2006). If our knowledge of the market is this comprehensive, we can obtain an unbiased estimate for IV from (31) in a GMM framework using three moments.[11] More specifically, a standard exponential learning model (e.g. Easley and O'Hara, 1992) imposes $\alpha = e^{-r}$ and $\beta = -\sigma$, so that

$$\Delta p_t = 0 \cdot \sigma\varepsilon_t + \sigma\left(e^r - 1\right)\sum_{i=1}^{\infty} e^{-ri}\varepsilon_{t-i} = \sum_{i=1}^{\infty}\left[-e^{-ri} + e^{-r(i-1)}\right]\sigma\varepsilon_{t-i}.$$

The resulting estimate of IV is a scaled version of standard RV

$$\widehat{IV} = \frac{e^{\hat{r}} + 1}{e^{\hat{r}} - 1} \cdot RV = \frac{RV + \gamma_1}{RV - \gamma_1} \cdot RV, \tag{32}$$

[10] Hasbrouck (1993) and recently Hansen, Large, and Lunde (2008) show how to embed (30) into general moving average (MA)-based estimators. Such general MA-estimators are warranted if the researcher has only limited information about the microstructure of the market or has interest different from IV estimation, such as forecasting the latent price process. If, however, the microstructure is known and interest centers on estimating IV, as we assume here, then our estimators may be more appealing.

[11] The proof, which we sketch here, is straightforward. Recast the price process (1) and (2) in continuous time, so that $\Delta p_t = \Delta p_t^* + \Delta u_t/\sqrt{m}$, with m denoting the number of subintervals, tm equal to one unit of calendar time, and the scale of t suitably redefined. Then, following the considerations of Section 5.3, under standard assumptions r is invariant to m and local infill asymptotic theory can be applied.

where the scaling factor requires a consistent estimate of only one additional parameter, the market maker's learning rate, r. It is interesting to note the resemblance to the estimator of Hansen et al. (2008), which is also a scaled variant of RV. In contrast to our approach, they do not exploit (that is, condition on) a specific market microstructure, but attempt to achieve robustness to a wide range of possible microstructures.

Estimator (32) offers a structural interpretation to estimates of noise and IV. The learning model predicts that the MSN at all lags decreases with the learning rate. Slow learning implies a very persistent cross-correlation between noise and latent returns, and hence persistent autocorrelation of noise, so that fluctuations in MSN tend to dominate the IV.

Figure 5 provides some perspective. It is based on the noise-to-IV ratios reported by Hansen and Lunde (2006), which are (unfortunately) derived under the assumption of independent noise. The ratio of noise to IV shrinks with the number of price-changing quotes per day. If the number of times that the market maker changes his price quote during a trading day is indicative of his speed of learning, then MSN indeed decreases as the learning rate of the market maker increases. This supports the multiperiod learning model.

Figure 5: Ratio of Noise to Integrated Variance, as a Function of Quotes per Day

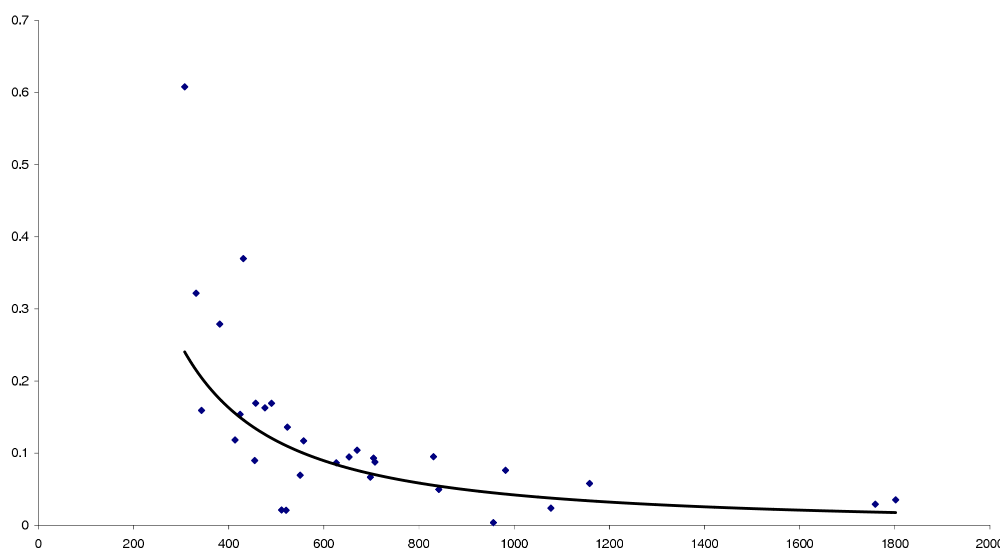

Notes: The vertical axis measures the noise-to-signal ratio as 100 times noise divided by IV under the assumption of independent noise. The horizontal axis gives the number of quotes per day with a price change. Data are for 30 NYSE and NASDAQ equities in 2000, obtained from Hansen and Lunde (2006) Tables 1 and 3. The solid line is a fitted trend.

Furthermore, the recent decline in noise-induced bias of RV (Hansen and Lunde's (2006) fact III) suggests that the learning rate r has increased. Adding to this Meddahi's (2002) finding that the standard deviation of the bias is large relative to the IV suggests that the learning rate itself may have fluctuated considerably around its increasing trend.

Our example uses a MA process with only two free coefficients, but the large sample sizes typical with high frequency data can accommodate much richer specifications. Empirical work in market microstructure tends to favor extreme parameterizations, ranging from the very parsimonious as in the regressions of Glosten and Harris (1988), to the profligate as in the vector autoregressions of Hasbrouck (1996). For RV noise correction the most useful parameterizations may be intermediate, imposing a general correlation pattern but avoiding highly situation-specific assumptions.

Dynamic market microstructure models imply much richer noise structure than the two polar cases of immediate and slow decay that we just discussed. These restrictions can be exploited to construct tailored volatility estimators. In Table 3 we do so by suggesting parsimonious estimators for a variety of market microstructures. Whereas these estimators inevitably also remove price features that are empirically indistinguishable from modelled MSN, their parsimony ensures that this miscorrection is kept to a minimum.

The table is structured as follows: The left column gives the estimator, the middle column an example of independent MSN, to which this estimator applies, and the right column an example of dependent MSN. Interestingly, for many market microstructures that generate dependent noise there is a corresponding market structure with independent noise to which the same estimator fits. The Web Appendix B shows that all these estimators are unbiased. They are consistent under the conditions discussed in Section 5.3 or under subsampling (Barndorff-Nielsen et al., 2011b).

We only discuss the dependent noise cases here, because these are – as we have shown in this paper – the ones of relevance in actual financial markets. The first row of Table 3 shows that the Zhou (1996)-estimator is the most parsimonious way to deal with a market in which the only MSN stems from the bid-ask bounce, even if trades are driven by private information. The geometric decay of MSN over time under learning is covered by rows two and three, for various exogenous noise processes. The decay becomes linear if traders act strategically, reflected in row four. These three learning estimators specialize to the Zhou (1996)-estimator with $\gamma_2 = 0$ for nonstrategic, and with $\gamma_1 = \gamma_2$ or $S = 1$ for strategic informed traders. Likewise, the noise process we discussed earlier in Equation (32), $u_t = \alpha u_{t-1} - \sigma \varepsilon_t$, is a special case of the nonstrategic incomplete informed trader case, with $\beta = -\sigma$ and $v_t = 0$. Finally, the estimator for strategic informed traders collapse to the estimator for linear independent noise decay if $\gamma_2 = 2\gamma_1$.

The contribution of the delayed price responses to the learning RV estimators in rows two and three can be expressed by any pair of autocovariances, γ_i, γ_{i+1}, $i \geq 2$. Whereas in the table we show the most parsimonious expression, replacing the last term by an average stabilizes the estimates. For example, in the nonstrategic incompletely informed trader case, we can use $\gamma_0 + 2\gamma_1 \frac{1}{S} \sum_{i=1}^{S} \frac{\gamma_i}{\gamma_i - \gamma_{i+1}}$, for any $S \geq 1$.

With strategic informed traders choosing the correct length of the private information period is critical for unbiased results, as noted already by Kelly and Steigerwald (2004). In our setup in row four S can be estimated by $\hat{S} = \sqrt{\left(\frac{3\gamma_1 - \gamma_2}{2(\gamma_2 - \gamma_1)}\right)^2 + \frac{2}{\gamma_2 - \gamma_1} \sum_{i=1}^{\infty} \gamma_i} - \frac{3\gamma_1 - \gamma_2}{2(\gamma_2 - \gamma_1)}$.

The MSN in the upper four rows of Table 3 is asymptotically uncorrelated, so IV can be expressed by Equation (29). This equation does not hold in market maker inventory models, as the ones in the bottom row. There, MSN follows a unit root process with $\Delta u_t = \alpha q_t$ so that $E(u_t \Delta u_{t-i}) = \alpha^2 \; \forall i$. In this case autocovariances alone are not sufficient, and $I\hat{V}$ must be based on the general Equation (28).

A common argument for using estimators that, contrary to Equation (29), downweight autocovariances at non-zero displacements is that it rules out the possibility of a negative volatility estimate. Starting the analysis with such an estimator, however, strips the researcher of the chance to falsify his assumptions on the market microstructure. After all, a negative variance estimate first and foremost indicates that the estimator is misspecified for the microstructure of the market under analysis, and that it should be refined. We therefore suggest starting with microstructure-inspired estimators as the ones in Table 3, and resort to microstructure-free estimators if the market microstructure appears to obey to none of the common models.

Barndorff-Nielsen, Hansen, Lunde, and Shephard (2009) notice that their IV estimator's ability of detecting properties of volatility depends crucially on the bandwidth: "The 'strength' of this 'microscope' is controlled by the bandwidth parameter, and the realized kernel gradually looses its ability to detect volatility at the local level as ... [the bandwidth] is increased." (Barndorff-Nielsen et al., 2009, p.C27) In effect, there is a tradeoff between the loss of local volatility information and the MSN bias. Utilizing prior knowledge about the market microstructure, Table 3 allows an informed bandwidth choice instead of having to rely exclusively on statistical arguments.

Table 3: Noise-Robust Estimators for Realized Volatility

IV	Independent – Strong-Form Noise –	Dependent
$\gamma_0 + 2\gamma_1$	Measurement Error/ Discrete Data $u_t = v_t$ Bid-Ask Bounce $u_t = \alpha q_t,\ q_t \in \{-1, +1\}$ iid	Dependent Measurement Error $u_t = \alpha\varepsilon_t + v_t$ Bid-Ask Bounce from Informed Traders $u_t = \alpha q_t,\ q_t = \begin{cases} -1 & \text{if } \varepsilon_t < 0 \\ +1 & \text{if } \varepsilon_t > 0 \end{cases}$
$\gamma_0 + 2\frac{\gamma_1^2}{\gamma_1 - \gamma_2}$	Autoregressive Noise $u_t = \alpha u_{t-1} + v_t$	Nonstrategic Incompletely Informed Traders $u_t = \alpha u_{t-1} + \beta(\varepsilon_t + v_t)$
$\gamma_0 + 2\gamma_1 + 2\frac{\gamma_2^2}{\gamma_2 - \gamma_3}$	Autoregressive Noise with Measurement Error $u_t = \sum_{i=0}^{\infty} \alpha^i v_{t-i} + w_t$	Autoregressive Noise with One-Period Private Information $u_t = \sum_{i=0}^{\infty} \alpha^i v_{t-i} + \beta\varepsilon_t + w_t$ Nonstrategic Informed Traders $u_t = \alpha \sum_{i=0}^{\infty} \beta^i \varepsilon_{t-i} + v_t$
$\gamma_0 + S(S+1)\gamma_1$	Linear Noise Decay over S Periods $u_t = \alpha \sum_{i=t}^{t-S} \frac{i-t+S}{S} v_i$	
$\gamma_0 + S(3-S)\gamma_1$ $+S(S-1)\gamma_2$		Strategic Informed Traders with S-Period Private Information $u_t = \alpha \sum_{i=t}^{t-S} \frac{i-t+S}{S} (\varepsilon_i + v_i)$
$\gamma_0 - {\gamma_0^*}^2$	Market Maker Inventory from Noise Trading $u_t = \alpha \sum_{i=0}^{\infty} q_{t-i},$ $q_t \in \{-1, +1\}$ iid	
$\frac{\pi}{\pi-2}\left(\gamma_0 - {\gamma_0^*}^2\right)$		Market Maker Inventory from Informed Trading $u_t = \alpha \sum_{i=0}^{\infty} q_{t-i},\ \varepsilon_t \overset{iid}{\sim} N(0,1)$ $q_t = \begin{cases} -1 & \text{if } \varepsilon_t < 0 \\ +1 & \text{if } \varepsilon_t > 0 \end{cases}$

The estimators are based on the observable moments $\gamma_i \equiv E(\Delta p_t \Delta p_{t-i})$ and $\gamma_i^* \equiv E(\Delta p_t q_{t-i})$. They are based on the assumption that the latent price changes every period ($T = 1$), and remains unobserved for one or more periods, depending on the noise specification. The probability of no latent price change has measure zero. The two white noise processes are $v_t \overset{iid}{\sim} (0, \eta_v^2)$ and $w_t \overset{iid}{\sim} (0, \eta_w^2)$, where $E(v_s w_t) = 0\ \forall s, t$.

6.2 On Structural vs. Non-Structural Volatility Estimators

Here we emphasize that the more the econometrician knows about the price process of relevance, the more the noise correction can be tailored to it by exploiting microstructure theory. This is important, because as discussed in Section 2, the price process of interest may differ across users of volatility estimates. Many users are likely to be interested in price processes different from (1), which has implications for appropriate volatility estimation. The variance of strong form efficient returns, $E(\Delta p_t^*) = \frac{\sigma^2}{T}$, the price under full information, differs both conceptually and numerically from the variance of semi-strong form efficient returns,

$$E\left[(\Delta \tilde{p}_t^e)^2\right] = \frac{1}{T}\left\{\sigma^2 + \sum_{i=0}^{T-1}\phi_i\lambda_i^2 + E\left[\left(\sum_{i=-1}^{-T}\lambda_i q_i\right)^2\right] - 2\sigma\sum_{i=-1}^{-T}\lambda_i E(q_i\varepsilon_{-T})\right\}, \tag{33}$$

which is the volatility that affects the balance sheet of the market maker. It might therefore be more applicable to studies of market maker behavior than $E\left[(\Delta p_t^*)^2\right]$. To take a simple example, consider again one-period private information, $T = 1$, in which case strong form volatility is σ^2 and semi-strong volatility (33) simplifies to

$$E\left[(\Delta \tilde{p}_t^e)^2\right] = \sigma^2 + 2\phi\lambda^2 - 2\sigma\lambda E(q_t\varepsilon_t) \neq \sigma^2. \tag{34}$$

The RV estimator of Zhou (1996) is

$$RV_{AC(1)} = E(\Delta p_t^2) + E(\Delta p_{t-1}\Delta p_t) + E(\Delta p_t \Delta p_{t+1}),$$

which is equivalent to Equation (30). For $T = 1$ it is

$$E\left(RV_{AC(1)}\right) = E\left\{[s(q_t - q_{t-1}) + \sigma\varepsilon_{t-1}] \times [\sigma(\varepsilon_t + \varepsilon_{t-1} + \varepsilon_{t-2}) + s(q_{t+1} - q_{t-2})]\right\} = \sigma^2.$$

Hence although $RV_{AC(1)}$ is unbiased for σ^2, it is in general biased with ambiguous direction for $Var(\Delta \tilde{p}_t^e)$ in (34). The same applies to a noise-robust estimator with a large, potentially infinite, lag window, which removes *any* microstructure and other correlation effect. For these estimators to work, the latent return process *of interest* must follow a martingale difference sequence. Semi-strong form prices do not; they are serially correlated and inevitably $RV_{AC(1)}$ is biased relative to $Var(\Delta \tilde{p}_t^e)$.

What could an estimator of semi-strong form volatility look like? Consider, for example, a market where the strong form efficient price become public after two periods. From (4)–(9), we obtain $\Delta p_t = \Delta \tilde{p}_t^e + \Delta u_t$ with noise given by (27). It follows that

$$E\left[(\Delta \tilde{p}_t^e)^2\right] = E\left(\Delta p_t^2\right) + 2s\phi(\lambda - s).$$

As $\tilde{p}_t^e$ is generated by a more complex process that p_t^*, we need additional market data. Using the autocorrelations of prices, additional market information such as an estimate of the spread and of the trade frequency ϕ, an unbiased estimator for IV of the semi-strong form efficient price is

$$\widehat{IV} = RV + 2\gamma_1 - 2\frac{\gamma_1\gamma_2}{\gamma_2 + s^2\phi}. \tag{35}$$

The obvious difference to the estimators in Table 3 emphasizes the importance of carefully defining the latent price series of interest.[12] This is where market microstructure theory can contribute new insights to IV estimation. By providing distinctive but flexible relationships between MSN and latent returns, and using additional market information, the agnostic statistical noise estimate can be decomposed into its various MSN and fundamental components.

6.3 Empirical Application

Estimates of volatility are important in many areas. They are, for example, central to risk management or serve as input to policy making. In this section we use market microstructure-based estimators in two real-world applications. We first compare our estimators' properties with the standard RV and a statistical IV estimator in a well-known stock market dataset.

[12]To avoid confusion we adhere in this paper to the convention that the strong form efficient price follows a martingale. Therefore we introduced $\tilde{p}_t^e$ as another latent price series of interest. But there is no guarantee that a price with martingale properties exists in a given market. For example, the latent price could itself be the result of learning about random-walk fundamentals, in which case p_t^* has the properties of the semi-strong form efficient price $\tilde{p}_t^e$. Specifically, let fundamentals follow $\chi_t = \chi_{t-1} + \varepsilon_t$ with $\varepsilon_t \overset{iid}{\sim} (0, \sigma^2)$. Then the latent price process, known only to the best informed market participants, is

$$\Delta p_t^* = \sigma \sum_{i=1}^{T} \left[-e^{-r_1 i} + e^{-r_1(i-1)}\right] \varepsilon_{t-i}.$$

If market makers are well informed ($p_t^e = p_t^*$) and the bid-ask bounce follows Equation (6), then mechanically calculating $RV_{AC(T)}$ gives the variance of the fundamental, not the variance of the strong form efficient price. Obviously, a purely statistical noise correction cannot distinguish between cross-correlation caused by fundamentals and cross-correlation caused by MSN.

After this, we turn to a current policy debate centering on the volatility in the oil futures market.

6.3.1 Alcoa Stock

As a first application, we compare microstructure-based estimates with statistical estimates of IV of Alcoa Inc. (AA) stock.[13] We use prices on the New York Stock Exchange for the year 2004 from Barndorff-Nielsen et al. (2009). All overnight returns and days with less than five hours of trading were removed from this dataset, which means that the IV-estimates apply only to the price process *within* trading days. They do not capture the overall riskiness of the stock, because price changes between trading days are excluded.

The estimators in Tables 4 to 7 are for daily IV, i.e. $E[(\Delta p_t^*)^2]$, averaged across the year. *RV Standard* is simple realized volatility, $E(\Delta p_t^2)$, *RV Bid-ask* is the bid-ask estimator (30), *RV Learning - Restricted* the learning estimator (32), and *RV Learning - Nonstr. Noisy* stands for the estimator for nonstrategic, incompletely informed traders. All microstructure-based estimators are defined by Table 3. *RV Statistical* is the consistent flat-top kernel estimator $RV_{ACNW(30)} = \gamma_0 + 2\sum_{i=1}^{30}\gamma_i + 2\sum_{i=1}^{30}\frac{30-i}{30}\gamma_{30+i}$ of Hansen and Lunde (2006). It serves as benchmark, as a statistical estimator that removes all deviations of the transaction price from a martingale, which might be different from the IV of the – in our terminology – true latent price process, i.e. RV corrected for market-microstructure-induced noise only.

All estimators except the standard estimator allow for correlation between noise and latent price. We do not implement the inventory estimators here, because the dataset does not contain signed trades. In this section, we refer to the difference between *RV Standard* and *RV Statistical* as "noise", in contrast to deviations due to market microstructure effects, which we call MSN.

Table 4 reveals that under 1/second calendar time sampling (CTS) sampling both the restricted learning and bid-ask estimators explain one third of noise in transaction prices (in the second column). Learning appears to be very fast ($\hat{r} > 3$), which implies that γ_1 is small compared to RV. As a result the learning and bid-ask volatility estimates are very similar. More flexible learning estimators capture more of the noise. *RV Learning - Nonstrategic*, in particular, captures more than 90% percent of what *RV Statistical* removes as noise. This means that for Alcoa under CTS indeed most of the noise correction embedded in RV

[13] Gatheral and Oomen (2010) compare 19 IV estimators on simulated data and conclude that a realized kernel and a maximum likelihood based estimator perform best in practice. However, they ignore microstructure noise for the most part. Patton (2011) compares four statistical IV estimators of IBM stock prices under time-varying volatility. Absent jumps, they perform better than standard RV sampled 1/5 minutes.

Statistical is most likely justified – it is MSN stemming from nonstrategic informed traders. Similarly, *RV Learning - Nonstr. Noisy* and *RV Learning - Strategic* capture between two-thirds and all of noise.

Table 4: Comparison of Realized Volatility Estimators (CTS at 1/second)

RV	price	mid	bid	ask
Standard	2.493	1.605	2.733	2.685
Statistical - $ACNW(30)$	2.146	2.141	2.255	2.257
Bid-ask - $AC(1)$	2.377	1.547	2.524	2.475
Learning - Restricted	2.379	1.548	2.532	2.483
Learning - Nonstr. Noisy	2.268	1.437	2.603	2.395
Learning - Nonstrategic	2.171	1.435	2.429	2.237
Learning - Strategic	2.361	2.363	2.368	2.409

Under tick time sampling (TTS) all microstructure-based estimators estimate IV substantially lower than *RV Standard* and *RV Statistical*. But if microstructure-based estimators remove the most common MSN types at this sampling frequency, then what does *RV Statistical* add back in? What positive cross-correlation between the latent price and noise different from learning can justify the higher estimate? And this point we have to leave this for further research, but also as a warning against a noise correction without a microstructure interpretation in mind.

Table 5: Comparison of Realized Volatility Estimators (TTS at 1/tick)

RV	price	mid	bid	ask
Standard	2.494	1.605	2.733	2.685
Statistical - $ACNW(30)$	2.386	2.511	2.506	2.534
Bid-ask - $AC(1)$	1.813	1.603	2.313	2.238
Learning - Restricted	1.895	1.603	2.343	2.272
Learning - Nonstr. Noisy	1.938	1.602	2.424	2.194
Learning - Nonstrategic	1.816	1.677	2.208	2.096
Learning - Strategic	2.164	2.290	2.284	2.304

At lower sampling frequencies the microstructure-based estimators are less tightly linked to the model setup under which we derived them. Whereas *RV Standard* and *RV Statistical* almost coincide that these frequencies, the learning estimators suggest a downward correction under CTS (Table 6) and upward correction under TTS (Table 7).

Table 6: Comparison of Realized Volatility Estimators (CTS at 1/10 seconds)

RV	price	mid	bid	ask
Standard	2.149	1.585	2.328	2.244
Statistical - $ACNW(30)$	2.155	2.158	2.160	2.169
Bid-ask - $AC(1)$	1.970	1.764	2.160	2.102
Learning - Restricted	1.977	1.774	2.166	2.107
Learning - Nonstr. Noisy	2.094	1.848	2.190	2.141
Learning - Nonstrategic	1.983	1.994	2.168	2.132
Learning - Strategic	2.211	2.204	2.206	2.214

Table 7: Comparison of Realized Volatility Estimators (TTS at 1/10 ticks)

RV	price	mid	bid	ask
Standard	2.117	1.835	2.216	2.162
Statistical - $ACNW(30)$	2.364	2.232	2.227	2.242
Bid-ask - $AC(1)$	2.493	2.273	2.307	2.323
Learning - Restricted	2.530	2.333	2.309	2.330
Learning - Nonstr. Noisy	2.494	2.206	2.243	2.272
Learning - Nonstrategic	2.525	2.429	2.400	2.424
Learning - Strategic	2.482	2.354	2.361	2.357

Examining the structural parameter estimates (not tabulated) provides additional guidance about which microstructure effects are at work at a given frequency. For example, under TTS and transaction prices, the restricted learning estimator fits the data at sampling intervals below 20 ticks (and beyond 130), whereas the strategic learning estimator at intervals up to about 130 ticks. Under CTS, restricted learning fits at frequencies of 1/30 seconds and slower (and is thus not reliable for the frequencies reported in the tables), and strategic learning at frequencies of 1/30 seconds and faster.

The IV estimates based on mid quotes are smaller than the other estimates at sampling frequencies of 1/1 second or 1/1 tick, in Tables 4 and 5 respectively. This calls for caution.

Figure 6: RV Estimators and Market Maker Learning

(a) CTS (b) TTS

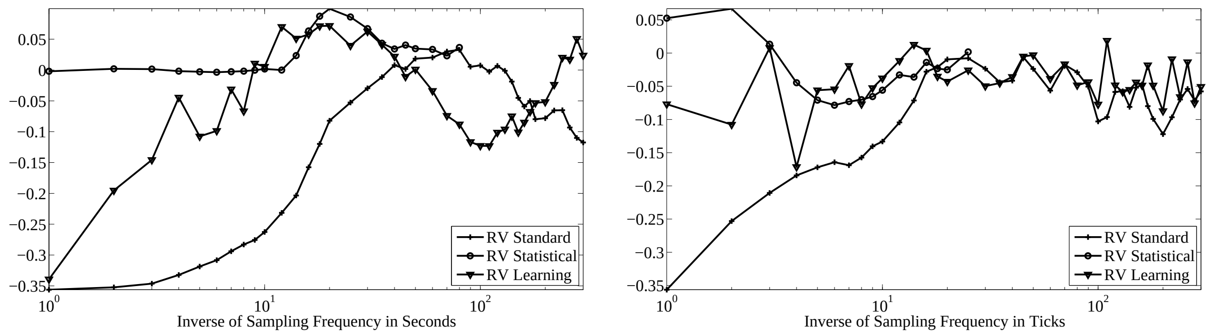

In the microstructure model setup we discussed, midprices are the least noisy among the four prices. This would even be true if the spread was time-varying, as long as it was independent of future changes in midprices, e.g. $E(\Delta p_t^e \Delta p_{t-1}^{bid}) = E(\Delta p_t^e \Delta p_{t-1}^{ask}) = 0$. If this was the case, the noise correction would push estimates towards midprice-based estimates, and when lowering sampling rates all IV estimates would converge to these midprice-based values. Tables 6 and 7 reveal that the opposite is the case: At lower sampling frequencies the midprice IV estimates reach the IV estimates of the other three price series. Because none of our microstructure-based IV estimators acceptably corrects the midprice estimates, we conclude that the midprices are subject to a microstructure effect that we did not take into account in deriving the estimators. A likely explanation is an asymmetrically moving spread, where a change in the bid price, say, is followed by an analogous change in the ask price in a later period, thus temporarily widening the spread. The temporary uncertainty that the wider spread represents is justified, because over longer horizons the latent price is indeed that volatile. The unraveling of uncertainty can be seen as an instance of market maker learning, so there is reason to hope that a learning estimator such as *RV Learning Nonstrategic* improves the estimate. This is indeed the case. Figure 6 shows the deviation of IV estimates based on midprices from the estimates based on transaction prices, expressed by the ratio $\frac{RV_{mid} - RV_{trans}}{RV_{trans}}$. Under TTS, shown on the right panel, the learning estimator does well despite its misspecification. It also improves the estimate under CTS, except at very high frequencies, which is shown on the left panel. Estimators with wide lag windows, such as *RV Statistical* and *RV Learning - Strategic* with estimated learning period are robust to this kind of time-varying spread. However, as said, they remove this part of MSN jointly with non-MSN components.

Figure 7: Volatility Signature Plots for Transaction Prices

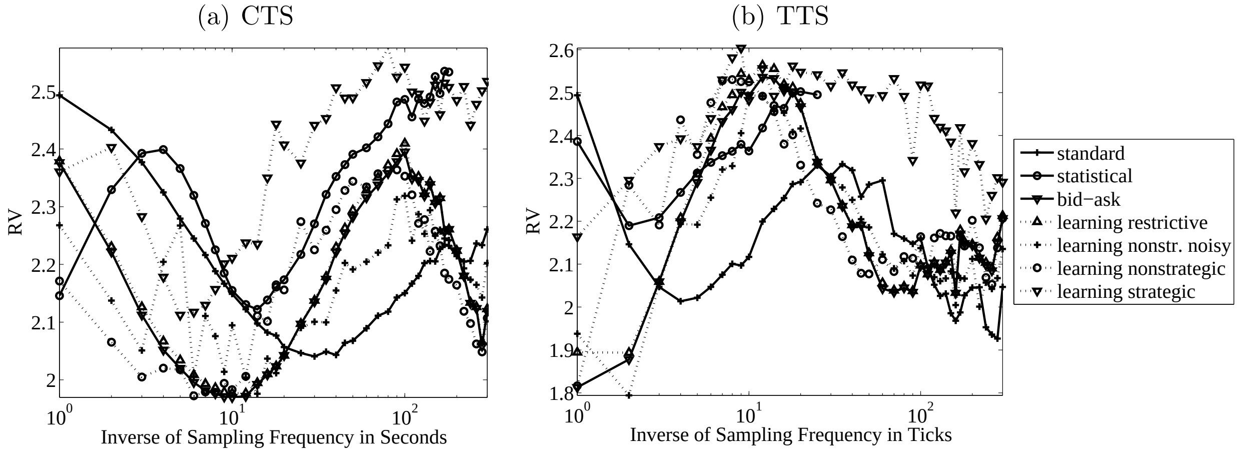

The volatility signature plots (Andersen, Bollerslev, Diebold, and Labys, 2000) in Figure 7 graph average daily realized volatility as a function of the underlying sampling frequency. One might argue against the use of parsimonious, but microstructure-based, estimators on the practical ground that they do not fully stabilize as the sampling frequency approaches its limit, i.e. 1/tick. The volatility signature plots reveal, however, that for the given data *RV Statistical* is not stable either – it moves in a range of 2.2 – 2.5 for sampling frequencies above 1/100 seconds or ticks. Most microstructure-based estimators are just as stable.

6.3.2 Crude Oil Futures

In this subsection we apply our estimators to Light Sweet Crude Oil futures traded on the NYMEX in Chicago (symbol CL). Our dataset consists of tick-by-tick transaction data from Tick Data, Inc., covering the period from January 2^{nd}, 1987 until September 24^{th}, 2010. It contains trades both within and outside of the main trading hours, which are Monday through Friday from 9:00 a.m. until 2:30 p.m. Eastern Time. The oil future is a standardized contract. One contract covers 1000 barrels with a fixed expiration date, on which oil has to be physically delivered at Cushing, OK. 66% of trades in our sample are for one contract, and less than 10% are for more than ten contracts.

Physical delivery is the exception, however, as most market participant roll their positions over to a new contract. We replicate this rollover by constructing a single time series of oil futures prices from the set of futures of different maturities simultaneously traded at a each point in time. We switch from one contract to the contract maturing next as soon as the daily volume of the latter exceeds the current contract's volume. In the following analysis, we use TTS and exclude contract rollover and overnight returns.

Figure 8: Volatility Signature Plots of Oil Futures by Day of Week

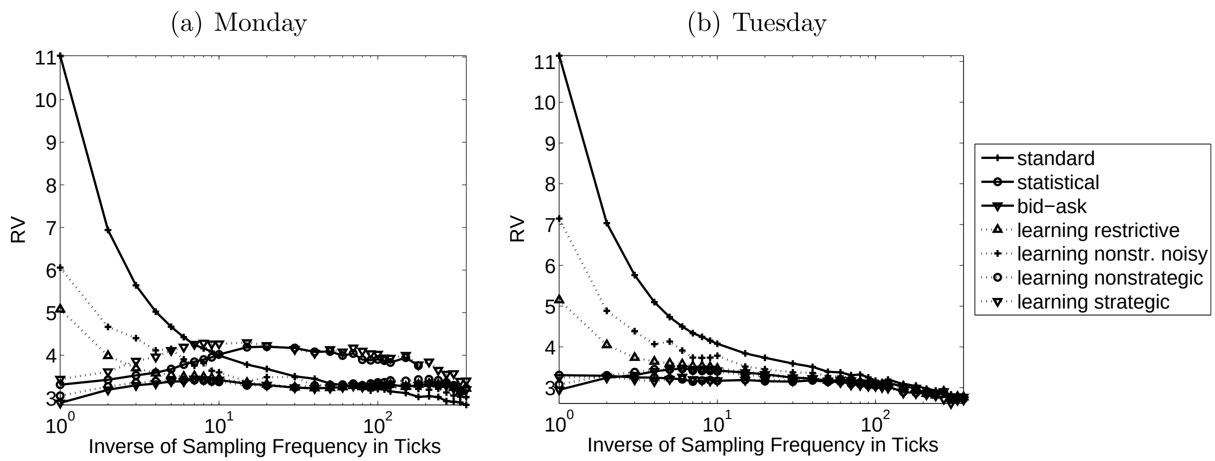

Figure 9: Volatility Signature Plots of Oil Futures by Year

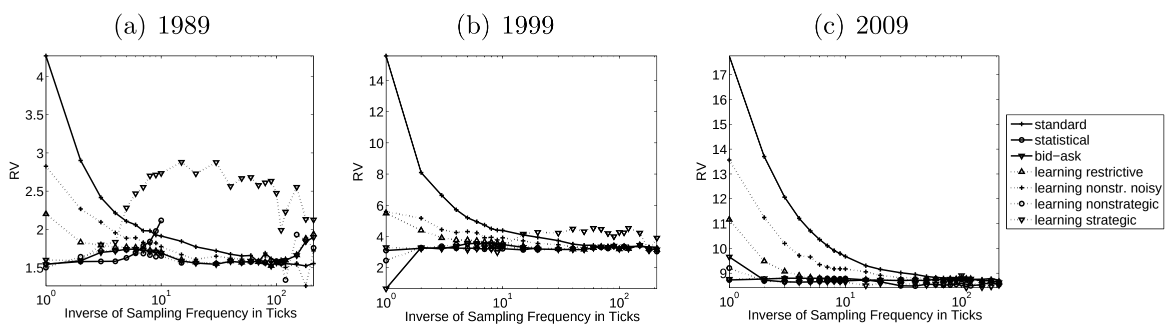

Comparing the IV estimates from the estimators discussed in this paper, the volatility signature plots in Figures 8 and 9 reveal that when sampling at a rate of 1/10 ticks or slower all estimators coincide. At higher sampling frequencies *RV Standard* diverges, which vividly depicts the MSN in oil futures data. The two most restrictive learning estimators, *RV Learning - Restrictive* and *RV Learning - Nonstrategic Noisy*, do not stabilize either at higher frequencies, suggesting that MSN in oil futures is more complex than this. In contrast, the estimators *RV Learning - Nonstrategic*, *RV Learning - Strategic*, and *RV Statistical* stabilize as sampling frequency reaches 1/1 tick. This convergence pattern in ticks did not change over the years despite a decline in the time between ticks from more than one per minute in 1989 to less than one per second in 2009.

Figure 10: Integrated Variance of Oil Futures by Calendar Month, 1987-2010

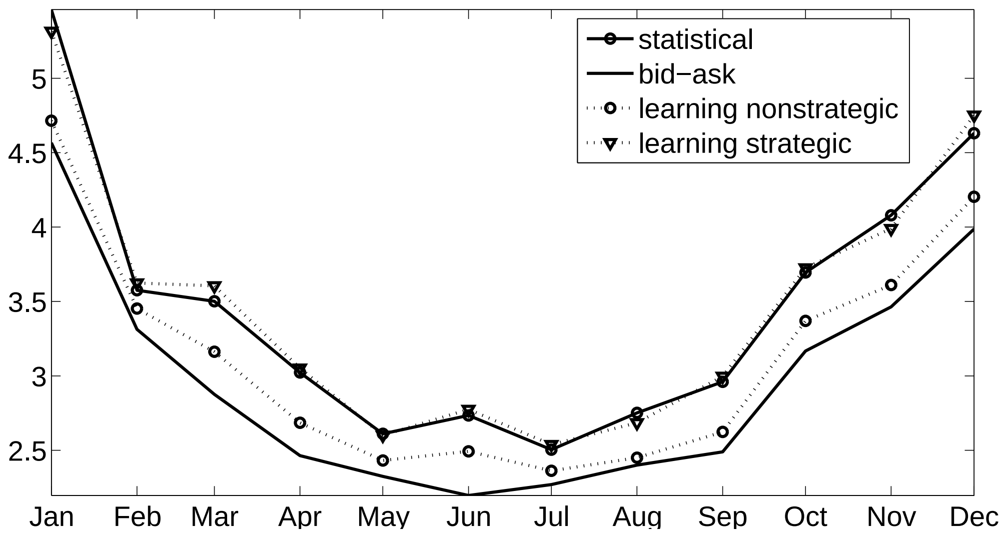

In the remaining analysis we sample at the highest possible frequency, i.e. 1/1 tick, and use accordingly only the four estimators that we identified in the volatility signature plot to converge with oil futures data. Prices of oil futures follow a pronounced seasonal volatility pattern. Figure 10 shows that volatility during 1987–2010 is particularly high in January, and reaches its low around July. There is no Monday effect. Instead, the volatility peaks on Wednesdays - where it is about 20% higher than on Mondays.

The fluctuations of the IV estimates over the years summarize the recent history of oil prices. In Figure 11 the average daily volatility of oil futures first spikes in 1990, when the world was faced with the Gulf War. After four calm years, 1992 to 1995, it plateaued at an intermediate level from 1996 until 2007, despite the steep increase in oil prices. The financial crisis pushed the volatility of oil futures to unprecedented levels in 2008 and 2009. As of 2010, the volatility is back to the plateau level from before the financial crisis. Given the seasonal pattern in average daily realized volatility, the 2010 value has to be adjusted upwards by a factor of about 1.5, because our dataset ends just before the Fall 2010. Even then, however, there is no clear evidence of excess volatility in oil prices at the volatility level of 2010. Based on the data available, regulation of derivatives in the oil market has to be justified with the volatility during the crisis years 2008/2009 – not with the most recent data –, or with the destabilizing effects of speculators on the market microstructure. For example, unlike in

Figure 11: Integrated Variance of Oil Futures by Calendar Year, 1987-2010

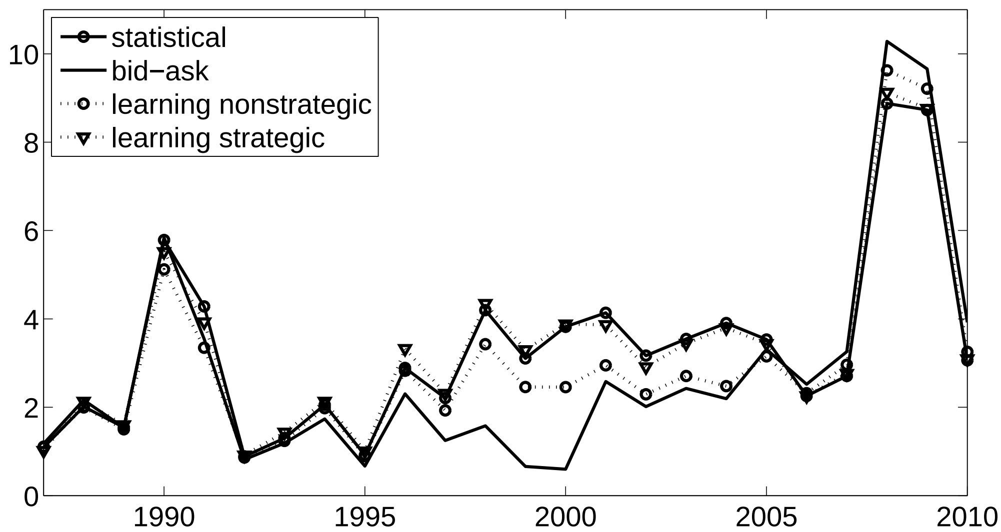

previous years, in 2009 the IV estimate correcting for strategic learning is smaller than the one for nonstrategic learning (Figure 9). This indicates that strategic trading, maybe part of speculative trading schemes, increased high-frequency market volatility.

The four estimators show a similar volatility path over time. Numerically, however, they differ considerably. *RV Bid-ask* and *RV Learning - Nonstrategic* estimate IV to be lower from the mid 1990s to the mid 2000s, and higher during 2008/2009 than the other two estimators. The switch between *RV Learning - Nonstrategic* and *RV Statistical* around 2006 precedes the financial crisis; it suggests that around that time the market structure changed. Noise different from learning first increased volatility, but dampened it during the financial crisis. What type of MSN can explain this change is an interesting question for further research. For example, a cross-correlation link between the strong form efficient price and noise in addition to learning, e.g. debt-financed trading, might have been muted during the crisis.

7 Concluding Remarks

The recent realized volatility literature provides statistical insights into market microstructure noise (MSN) and its effects. In this paper we have provided complementary *economic* insights, treating MSN not simply as a nuisance, but rather as the result of financial economic decisions, which we seek to understand. In that regard, we derived the predictions of economic theory regarding correlation between MSN and two types of latent price, characterizing and contrasting the entire cross-correlation functions in a variety of market environments, with a variety of results.

Some results are generic. For example, cross-correlations between strong form efficient price and MSN at displacements greater than zero have sign opposite to that of the contemporaneous correlation.

Some results are not generic but nevertheless quite robust to model choice. For example, all models predict negative contemporaneous correlation between latent price and MSN, so long as the risk aversion of market makers is not too high.

Finally, some results are highly model-specific. For example, the cross-correlation patterns and absolute magnitudes depend critically on the frequency of latent price changes, the presence of bid/ask bounce, the timing of information and actions, and the degree of market maker risk aversion.

We hope that the results of this paper will promote the use of theory in disciplining data. As our empirical applications suggest, a standard learning model goes quite far in explaining, and controlling for, MSN. We have also shown, that attention to market microstructure theory enables us to assess the validity of the independence assumption, to offer explanations of empirically observed cross-correlation patterns, to predict the existence of as-yet undiscovered patterns, and to make informed suggestions for improving volatility estimation methods. And conversely, of course, additional work along our lines may help promote the use of data in disciplining theory, by helping to sift the comparative merits of various competing theoretical microstructure models. Further improvements are possible as additional market information is now widely available at high frequency, such as, for example, trade size or trade origin data. This will allow to refine noise correction further to get an even better microstructure-founded volatility estimate.

Other novel uses of our results may also be possible. For example, the rate of decay of cross-correlations might be used to assess the extent to which strategic traders are active in the market, and the sign and size of the contemporaneous correlation might be used to assess the degree of market maker risk aversion. Indeed market maker risk aversion might be time-varying, with associated time-varying cross-correlation structure between latent price and MSN. During crises, for example, market makers may be more risk averse, as borrowing and hedging possibilities are reduced. If so, the "normal pattern" of negative contemporaneous cross-correlation and positive higher-order cross-correlations might switch to a "crisis pattern" of positive contemporaneous cross-correlation and negative higher-order cross-correlations. Such possibilities await future empirical exploration.

References

Aghion, P., Bolton, P., Harris, C., Jullien, B., 1991. Optimal learning by experimentation. Review of Economic Studies 58 (4), 621–654.

Aghion, P., Espinosa, M. P., Jullien, B., 1993. Dynamic duopoly with learning through market experimentation. Economic Theory 3, 517–539.

Aït-Sahalia, Y., Mykland, P. A., Zhang, L., 2005. How often to sample a continuous-time process in the presence of market microstructure noise. Review of Financial Studies 18 (2), 351–416.

Aït-Sahalia, Y., Mykland, P. A., Zhang, L., 2011. Ultra high frequency volatility estimation with dependent microstructure noise. Journal of Econometrics 160 (1), 160–175.

Andersen, T. G., Bollerslev, T., Christoffersen, P. F., Diebold, F. X., 2006. Volatility and correlation forecasting. In: Elliot, G., Granger, C. W. J., Timmermann, A. (Eds.), Handbook of Economic Forecasting. Amsterdam: North-Holland, pp. 777–878.

Andersen, T. G., Bollerslev, T., Diebold, F. X., 2010. Parametric and nonparametric volatility measurement. In: Hansen, L. P., Aït-Sahalia, Y. (Eds.), Handbook of Financial Econometrics. Amsterdam: North-Holland, pp. 67–137.

Andersen, T. G., Bollerslev, T., Diebold, F. X., Ebens, H., 2001a. The distribution of realized stock return volatility. Journal of Financial Economics 61 (1), 43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X., Labys, P., 2000. Great realizations. Risk 13 (March), 105–108.

Andersen, T. G., Bollerslev, T., Diebold, F. X., Labys, P., 2001b. The distribution of realized exchange rate volatility. Journal of the American Statistical Association 96, 42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. X., Labys, P., 2003. Modeling and forecasting realized volatility. Econometrica 71 (2), 579–625.

Bandi, F. M., Russell, J. R., 2006. Realized variance and market microstructure noise: Comment. Journal of Business & Economic Statistics 24 (2), 167–173.

Bandi, F. M., Russell, J. R., 2008. Microstructure noise, realized variance, and optimal sampling. Review of Economic Studies 75 (2), 339–369.

Bandi, F. M., Russell, J. R., 2011. Market microstructure noise, integrated variance estimators, and the accuracy of asymptotic approximations. Journal of Econometrics 160 (1), 145–159.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., Shephard, N., 2008. Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. Econometrica 76 (6), 1481–1536.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., Shephard, N., 2009. Realized kernels in practice: Trades and quotes. Econometrics Journal 12 (3), C1–C32.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., Shephard, N., 2011a. Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. Journal of Econometrics 162 (2), 149–169.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., Shephard, N., 2011b. Subsampling realised kernels. Journal of Econometrics 160 (1), 204–219.

Barndorff-Nielsen, O. E., Shephard, N., 2002a. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. Journal of the Royal Statistical Society B 64 (2), 253–280.

Barndorff-Nielsen, O. E., Shephard, N., 2002b. Estimating quadratic variation using realized variance. Journal of Applied Econometrics 17, 457–477.

Barndorff-Nielsen, O. E., Shephard, N., 2007. Variation, jumps, market frictions and high frequency data in financial econometrics. In: Blundell, R., Torsten, P., Newey, W. K. (Eds.), Advances in Economics and Econometrics, Theory and Applications, Ninth World Congress of Econometric Society. Cambridge: Cambridge University Press, pp. 328–372.

Diamond, D. W., Verrecchia, R. E., 1987. Constraints on short-selling and asset price adjustment to private information. Journal of Financial Economics 18, 277–311.

Diebold, F. X., 2006. Realized variance and market microstructure noise: Comment. Journal of Business & Economic Statistics 24 (2), 181–183.

Diebold, F. X., Strasser, G., 2010. On the correlation structure of microstructure noise: A financial economic approach. Working Paper 16469, NBER.

Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 69–90.

Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. Journal of Finance 47 (2), 576–605.

Engle, R. F., Sun, Z., 2007. When is noise not noise – A microstructure estimate of realized volatility, manuscript, Stern School, New York University.

Fama, E. F., 1970. Efficient capital markets: A review of theory and empirical work. Journal of Finance 25 (2), 383–417.

Gatheral, J., Oomen, R. C. A., 2010. Zero-intelligence realized variance estimation. Finance and Stochastics 14 (2), 249–283.

Glosten, L. R., Harris, L. E., 1988. Estimating the components of the bid/ask spread. Journal of Financial Economics 21, 123–142.

Glosten, L. R., Milgrom, P. R., 1985. Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.

Hansen, P. R., Large, J., Lunde, A., 2008. Moving average-based estimators of integrated variance. Econometric Reviews 27 (1), 79–111.

Hansen, P. R., Lunde, A., 2006. Realized variance and market microstructure noise. Journal of Business & Economic Statistics 24 (2), 127–161.

Hasbrouck, J., 1993. Assessing the quality of a security market: A new approach to transaction-cost measurement. Review of Financial Studies 6 (1), 191–212.

Hasbrouck, J., 1996. Modeling market microstructure time series. In: Maddala, G. S., Rao, C. R. (Eds.), Statistical Methods in Finance. Amsterdam: North Holland, pp. 647–692.

Hasbrouck, J., 2002. Stalking the efficient price in empirical microstructure specifications. Journal of Financial Markets 5, 329–339.

Hasbrouck, J., 2007. Empirical Market Microstructure. New York: Oxford University Press.

Kelly, D. L., Steigerwald, D. G., 2004. Private information and high-frequency stochastic volatility. Studies in Nonlinear Dynamics and Econometrics 8 (1), 1–28.

Kyle, A. S., 1985. Continuous auctions and insider trading. Econometrica 53 (6), 1315–1336.

Meddahi, N., 2002. A theoretical comparison between integrated and realized volatility. Journal of Applied Econometrics 17 (5), 479–508.

O'Hara, M., 1995. Market Microstructure Theory. Cambridge: Cambridge University Press.

Oomen, R. C. A., 2006. Realized variance and market microstructure noise: Comment. Journal of Business & Economic Statistics 24 (2), 195–202.

Patton, A. J., 2011. Data-based ranking of realised volatility estimators. Journal of Econometrics 161 (2), 284–303.

Roll, R., 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. Journal of Finance 39 (4), 1127–1139.

Zhang, L., Mykland, P. A., Aït-Sahalia, Y., 2005. A tale of two time scales: Determining integrated volatility with noisy high frequency data. Journal of the American Statistical Association 100 (472), 1394–1411.

Zhou, B., 1996. High-frequency data and volatility in foreign-exchange rates. Journal of Business & Economic Statistics 14 (1), 45–52.

Web Appendices

A Proofs of Propositions

A.0 General Multiperiod Cases

A.0.1 Multiperiod Case - Strong Form Efficient Prices

We calculate in the following cross-correlations between the strong form efficient returns (7)

$$\Delta p_t^* = \Delta p_{\kappa T}^* = \begin{cases} \sigma\varepsilon_{\kappa T} & \forall \kappa \in \mathbb{Z} \\ 0 & \forall \kappa \notin \mathbb{Z} \end{cases} \tag{36}$$

and the corresponding noise (3)

$$\Delta u_t = \Delta p_t - \Delta p_t^* = \Delta p_t^e + s_t q_t - s_{t-1} q_{t-1} - \Delta p_t^*. \tag{37}$$

As a shorthand notation we use in the following $p_x \equiv p_{\kappa T + x}\ \forall \kappa, x \in \mathbb{Z}$. In the period of a change in the strong form efficient price, the expectation about this price changes by

$$\begin{aligned} \Delta p_0^e &= p_0^e - p_{-1}^e \\ &= \sigma\varepsilon_{-T} - \sum_{t=2}^{T} (\lambda_{-t} q_{-t} + \omega_{-t}) - \omega_{-1} + \omega_0, \end{aligned}$$

and in all other periods by

$$\Delta p_t^e = \lambda_{t-1} q_{t-1} + \omega_t.$$

From (37) we get for $t = \kappa T$

$$\Delta u_0 = \sigma(\varepsilon_{-T} - \varepsilon_0) + s_0 q_0 - s_{-1} q_{-1} - \sum_{t=2}^{T} (\lambda_{-t} q_{-t} + \omega_{-t}) - \omega_{-1} + \omega_0 \tag{38}$$

and $\forall t \neq \kappa T$

$$\Delta u_t = \lambda_{t-1} q_{t-1} + s_t q_t - s_{t-1} q_{t-1} + \omega_t, \tag{39}$$

where the first term reflects information-revealing trades, the second and third term reflect the bid-ask bounce, and the last term new non-trade information.

This immediately gives to the cross-covariances of Section 3.1.

In order to obtain the cross-correlations, we need $Var(\Delta p_t^*)$ and $Var(\Delta u_t)$. The strong form efficient price has unconditional variance

$$Var(\Delta p_t^*) = \frac{1}{T} Var(\sigma\varepsilon_0) = \frac{\sigma^2}{T}, \tag{40}$$

and the corresponding noise has unconditional variance

$$\begin{aligned}
Var(\Delta u_t) &= \frac{1}{T}\sum_{i=0}^{T-1} Var(\Delta u_i) \\
&= \frac{1}{T}\Bigg\{ 2\sigma^2 + \sum_{i=0}^{T-1}\left(\phi_i s_i^2 + \phi_{i-1}s_{i-1}^2\right) - 2\sigma s_{T-1}E(q_{T-1}\varepsilon_0) - 2\sigma \sum_{i=0}^{T-2}\lambda_i E(q_i\varepsilon_0) \\
&- 2\sigma s_0 E(q_0\varepsilon_0) + 2s_{T-1}\sum_{i=0}^{T-2}\lambda_i E(q_i q_{T-1}) + E\left[\left(\sum_{i=0}^{T-2}\lambda_i q_i\right)^2\right] + E(\omega_0^2) \\
&+ E\left[\sum_{i=0}^{T-1}\omega_i^2\right] + 2s_0E(q_0\omega_0) - 2\sigma\sum_{i=0}^{T} E(\varepsilon_i\omega_i) + 2s_{T-1}\sum_{i=0}^{T-1}E(q_{T-1}\omega_i) \\
&+ \sum_{i=1}^{T-1}\left(\phi_{i-1}\lambda_{i-1}\left(\lambda_{i-1} - 2s_{i-1}\right) + 2(\lambda_{i-1} - s_{i-1})(s_i E(q_{i-1}q_i) + E(q_{i-1}\omega_i)\right) \\
&+ E(\omega_i^2) + 2s_i E(q_i\omega_i)\Big)\Bigg\},
\end{aligned} \tag{41}$$

where

$$\phi_t = E(q^2) = E\left[Prob(q = +1 \vee q = -1)\right] = \beta + (1-\beta)\alpha\left[1 - F(p_t^e + s_t) + F(p_t^e - s_t)\right].$$

A.0.2 Multiperiod Case - Semi-Strong Form Efficient Prices

In the period of a change in the strong form efficient price, in which also the previous strong form efficient price becomes public information, the semi-strong form efficient return is

$$\Delta\tilde{p}_0^e = \lambda_0 q_0 + \sigma\varepsilon_{-T} - \sum_{i=1}^{T}\left(\lambda_{-i}q_{-i} + \omega_{-i}\right) + \tilde{\omega}_0,$$

where the first term reflects the market maker's guess about the new strong form efficient return based on a trade, the second term internalizes the new information about the previous return, and as a countermove the sum undoes the now obsolete guesses about the previous return. In all other periods the semi-strong form efficient price changes by

$$\Delta \tilde{p}_t^e = \lambda_t q_t + \omega_t.$$

From (9) we get for $\forall t$

$$\Delta \tilde{u}_t = (s_t - \lambda_t) q_t - (s_{t-1} - \lambda_{t-1}) q_{t-1}, \tag{42}$$

where the first two terms reflect information-revealing trades, and the second two terms reflect the bid-ask bounce.

Using Assumption 1 this immediately leads to the cross-covariances in Section 4.1.

A.1 Proof of Proposition 1

Proposition (Cross-correlations in the Easley-O'Hara model)

The contemporaneous cross-correlation in the Easley and O'Hara (1992) model is

$$Corr\left(\Delta p_t^*, \Delta u_t\right) = -\frac{1 + e^{-r(T-1)}}{2\sqrt{K}} < 0,$$

and the cross-correlations at sufficiently large nonzero displacements follow

$$Corr\left(\Delta p_{t-\tau}^*, \Delta u_t\right) = \frac{e^r - 1}{2\sqrt{K}} e^{-r\tau} > 0, \ \forall \tau \in [1, T-1]$$

$$Corr\left(\Delta p_{t-T}^*, \Delta u_t\right) = \frac{e^{-r(T-1)}}{2\sqrt{K}} > 0,$$

where $K = K(r, T)$.

Proof:

Following the setup in Easley and O'Hara (1992), suppose the strong form efficient price process switches between a high state $\overline{p}$, a neutral state, and a low state $\underline{p}$, where χ is the probability of a non-neutral state and γ the probability of a high state given that the state is non-neutral.

$$p_t^* = \begin{cases} \overline{p} & \text{with probability } \chi\gamma \\ \gamma\underline{p} + (1-\gamma)\overline{p} & \text{w.p. } 1-\chi \\ \underline{p} & \text{w.p. } \chi(1-\gamma) \end{cases}$$

Therefore, for $t = \kappa T$, $\kappa \in \mathbb{Z}$

$$\Delta p_t^* = \begin{cases} \overline{p} - \underline{p} & \text{w.p. } \chi^2\gamma(1-\gamma) \\ \gamma(\overline{p} - \underline{p}) & \text{w.p. } 2\chi(1-\chi)(1-\gamma) \\ 0 & \text{w.p. } (1-\chi)^2 + \chi^2(\gamma^2 + (1-\gamma)^2) \\ \gamma(\underline{p} - \overline{p}) & \text{w.p. } 2\chi(1-\chi)\gamma \\ \underline{p} - \overline{p} & \text{w.p. } \chi^2\gamma(1-\gamma) \end{cases}$$

and $\Delta p_t^* = 0$ otherwise. Prices have the properties

$$\begin{aligned} E[(\Delta p_t^*)^2] &= \frac{(\overline{p} - \underline{p})^2}{T}\left[2\gamma^2\chi(1-\chi)(1-\gamma) + \chi^2\gamma(1-\gamma) + 2\gamma^2\chi(1-\chi)\gamma + \chi^2\gamma(1-\gamma)\right] \\ &= \frac{(\overline{p} - \underline{p})^2}{T} 2\chi\gamma(\gamma + \chi - 2\chi\gamma) \equiv \frac{\sigma^2}{T}, \end{aligned}$$

$$\begin{aligned} E(\Delta p_t^* \Delta p_t) &= 0, \\ E(\Delta p_{t-\tau}^* \Delta p_t^*) &= 0. \end{aligned}$$

For ease of exposition let us focus on the case $\gamma = 1/2$ and $\chi = 1$, i.e. latent prices are high and low with equal probability. Using the result from Easley and O'Hara (1992) that transaction prices converge to the strong form efficient price at an exponential rate we get

$$\Delta p_0 = \frac{\overline{p} - \underline{p}}{2}\left(e^{-r(T-1)} - 1\right) \operatorname{sgn}\left(p_{-T}^* - \frac{\overline{p} + \underline{p}}{2}\right)$$

$$\Delta p_\tau = \frac{\overline{p} - \underline{p}}{2}\left(e^{-r(\tau-1)} - e^{-r\tau}\right) \operatorname{sgn}\left(p_0 - \frac{\overline{p} + \underline{p}}{2}\right)$$

$$\Delta u_0 = \frac{\overline{p} - \underline{p}}{2}\left(e^{-r(T-1)} - 1\right) \operatorname{sgn}\left(p_{-T}^* - \frac{\overline{p} + \underline{p}}{2}\right) - \Delta p_0^*$$

$$\Delta u_\tau = \frac{\overline{p} - \underline{p}}{2}\left(e^{-r(\tau-1)} - e^{-r\tau}\right) \operatorname{sgn}\left(p_0 - \frac{\overline{p} + \underline{p}}{2}\right)$$

The contemporaneous cross-covariance ($\tau = 0$) is

$$\begin{aligned} Cov\left(\Delta p_t^*, \Delta u_t\right) &= \frac{1}{T} E\left(\Delta p_0^* \Delta u_0\right) \\ &= -\frac{\sigma^2}{2T}\left[1 + e^{-r(T-1)}\right]. \end{aligned}$$

The second term inside the brackets is an artifact of p_t^* not following a martingale. In the period of the efficient price change it is optimal for the market maker to set p_t to the unconditional mean of p_t^*, thereby offsetting the effect of all previous learning, which the efficient price change rendered obsolete.

The cross-covariance for $\tau \in [1; T-1]$ is

$$\begin{aligned} Cov\left(\Delta p_{t-\tau}^*, \Delta u_t\right) &= \frac{1}{T} E\left(\Delta p_0^* \Delta u_\tau\right) \\ &= \frac{\sigma^2}{2T}\left(-e^{-r\tau} + e^{-r(\tau-1)}\right), \end{aligned}$$

and for $\tau = T$ we have

$$\begin{aligned} Cov\left(\Delta p_{t-T}^*, \Delta u_t\right) &= \frac{1}{T} E\left(\Delta p_{-T}^* \Delta u_0\right) \\ &= \frac{\sigma^2}{2T} e^{-r(T-1)}. \end{aligned}$$

The variance of the noise is

$$\begin{aligned} Var(\Delta u_t) &= \frac{1}{T}\left[\frac{(\overline{p}-\underline{p})^2}{T}\left(e^{-r(T-1)} - 1\right)^2 + \sigma^2 + 2\frac{(\overline{p}-\underline{p})^2}{4}\left(e^{-r(T-1)} - 1\right)\right. \\ &+ \left.\sum_{\tau=1}^{T-1} \frac{(\overline{p}-\underline{p})^2}{4}\left(-e^{-r\tau} + e^{-r(\tau-1)}\right)^2\right] \\ &= \frac{\sigma^2}{T}\left[\frac{1}{2}e^{-2r(T-1)} + \frac{1}{2} + \frac{1}{2}(-e^r + 1)^2 \frac{\left(e^{-2r}\right)^{T-1} - 1}{e^{-2r} - 1}\right]. \end{aligned}$$

Denoting the term in brackets by $K = K(r, T)$ we get for the contemporaneous cross-correlation

$$Corr\left(\Delta p_t^*, \Delta u_t\right) = -\frac{1 + e^{-r(T-1)}}{2\sqrt{K}},$$

for the cross-correlation at displacements $\tau \in [1; T-1]$

$$Corr\left(\Delta p^*_{t-\tau}, \Delta u_t\right) = \frac{-e^{-r\tau} + e^{-r(\tau-1)}}{2\sqrt{K}},$$

and for the cross-correlation at displacement T

$$Corr\left(\Delta p^*_{t-T}, \Delta u_t\right) = \frac{e^{-r(T-1)}}{2\sqrt{K}}.$$

Q.E.D.

A.2 Proof of Proposition 2

Proposition (Cross-correlations in the Kyle model)

The contemporaneous cross-correlation in Kyle (1985) is

$$Corr\left(\Delta p^*_t, \Delta u_t\right) = -\sqrt{\frac{T}{T^2+1}},$$

the cross-correlations at displacements $\tau \in [1; T]$ are

$$Corr\left(\Delta p^*_{t-\tau}, \Delta u_t\right) = \sqrt{\frac{1}{T(T^2+1)}},$$

and all higher order cross-correlations are zero.

Proof:

In order to present a closed-form solution we use continuous time, $t \in [0, T]$, but note that Kyle (1985) discussed the discrete time case as well. The discussion is based on the assumption of Kyle (1985) that the reaction functions for quantity demanded and prices are linear, i.e. that $\lambda_t = \lambda$, and $s_t = s$. Nonlinear solutions might nevertheless exist as well.

We assume semi-strong market efficiency, and so $s = \lambda$. We get from (13)

$$Cov\left(\Delta p^*_t, \Delta u_t\right) = -\frac{\sigma}{T}\left(\lambda E(q\varepsilon_0) - \sigma\right) < 0.$$

From (14) the cross-covariance function at nonzero displacements

$$Cov\left(\Delta p^*_{t-\tau}, \Delta u_t\right) = \frac{\sigma}{T}\lambda E(q\varepsilon_0) > 0$$

is constant $\forall t \in [1, T-1]$, and zero $\forall t \geq T$.

More specifically, we derive based on (18) for the noise (assuming zero spread)

$$\Delta u_0 = \frac{\Delta p^*_{-T}}{T} - \int_0^{T-1} \frac{\sigma}{T-s} dB_s - \Delta p^*_0$$

and for $\tau \in [1, T-1]$

$$\Delta u_\tau = \frac{\Delta p^*_0}{T} + (T-\tau) \int_{\tau-1}^{\tau} \frac{\sigma}{T-s} dB_s - \int_0^{\tau-1} \frac{\sigma}{T-s} dB_s.$$

The variance of the noise is therefore

$$\begin{aligned} Var(\Delta u_t) &= \frac{1}{T}\left[E(\Delta u_0^2) + \sum_{t=1}^{T-1} E(\Delta u_t^2)\right] \\ &= \frac{\sigma^2}{T}\left[\frac{T+1}{T} + \frac{T-1}{T} + \frac{(T-1)^2}{T}\right] \\ &= \frac{\sigma^2}{T^2}\left(T^2+1\right). \end{aligned}$$

The covariances are simply, at displacement zero

$$Cov(\Delta p^*_t, \Delta u_t) = \frac{1}{T} Cov(\Delta p^*_0, -\Delta p^*_0) = \frac{-\sigma^2}{T},$$

and at higher order displacements

$$Cov(\Delta p^*_{t-\tau}, \Delta u_t) = \frac{1}{T} Cov(\Delta p^*_0, \frac{\Delta p^*_0}{T}) = \frac{\sigma^2}{T^2},$$

which leads directly to the cross-correlations given by Proposition 2. *Q.E.D.*

A.3 Proof of Proposition 3

Proposition (Strong form cross-correlation, one period model)

$$Corr(\Delta p^*_t, \Delta u_t) = \frac{1}{\sqrt{2}} \frac{sE(q_t \varepsilon_t) - \sigma}{\sqrt{\phi s^2 + \sigma^2 - 2s\sigma E(q_t \varepsilon_t)}}$$

$$Corr(\Delta p^*_{t-1}, \Delta u_t) = -Corr(\Delta p^*_t, \Delta u_t)$$

Proof:

With $T = 1$, no extra information, $\lambda_t = \lambda$, $s_t = s$, and thus $\phi_t = \phi \ \forall t$, the variance term (41) simplifies to $Var(\Delta u_t) = 2(\sigma^2 + \phi s^2) - 4s\sigma E(q_t \varepsilon_t)$. Plugging this into (17) gives the desired result. *Q.E.D.*

A.4 Proof of Proposition 4

Proposition (Bounds of contemporaneous cross-correlation)

$$-\frac{1}{\sqrt{2}} \leq Corr(\Delta p_t^*, \Delta u_t) \leq 0$$

Proof:

Negativity can be seen as follows. Uninformed traders trade randomly ($E(q_t|\varepsilon_t) = 0$), thus for them we have $sE(q_t^u \varepsilon_t) = 0$. In contrast, informed traders buy ($q_t = +1$) only when $\sigma\varepsilon_t > s$ and sell ($q_t = -1$) only when $\sigma\varepsilon_t < -s$. Thus in a market of only informed traders $\sigma q_t^i \varepsilon_t > s \geq 0 \ \forall t$. Therefore we can write

$$1 = E(q_t^{i\,2} \varepsilon_t^2) > E\left(\frac{s}{\sigma} q_t^i \varepsilon_t\right) > E\left(\frac{s^2}{\sigma^2}\right) > 0,$$

so in particular $\sigma > sE(q_t^i \varepsilon_t) > 0$. Combining informed and uninformed trades we have

$$\sigma \geq sE(q_t \varepsilon_t) > 0,$$

which implies that the contemporaneous cross-correlation (19) is negative.

Further, (19) is bounded from below by $-1/\sqrt{2}$, which we prove by contradiction. Suppose this was not the case, then from (19)

$$sE\left(q_t \varepsilon_t\right) - \sigma < -\sqrt{\phi s^2 + \sigma^2 - 2s\sigma E(q_t \varepsilon_t)}.$$

Squaring both sides and simplifying gives the condition

$$\left[E\left(q_t \varepsilon_t\right)\right]^2 > \phi, \tag{43}$$

but by Jensen's inequality and $\beta = 1$

$$\left[E\left(q_t \varepsilon_t\right)\right]^2 \leq E\left(q_t^2 \varepsilon_t^2\right) = 1,$$

which contradicts (43). *Q.E.D.*

A.5 Proof of Proposition 5

A.5.1 Optimal Midprice

Proposition (Optimal Midprice) The optimal midprice, $p(n)$, monotonically shifts from the median to the midpoint of the support of p_t^* with increasing risk aversion. In particular,

$$\begin{aligned} p(1) &= \text{Median}(p_t^*) \\ p(2) &= \text{E}(p_t^*) \\ p(\infty) &= \text{Midsupport}(p_t^*). \end{aligned}$$

Proof:

The first two equations in the proposition[14] are the well-known result that the median is the best predictor under linear (absolute) loss, whereas the mean is the best predictor under squared loss. The third equation is obtained by first noting that for any density $f(\cdot)$, which has all moments, we can apply Leibnitz's rule. Thus we obtain for (21) the first order condition

$$\int_{\underline{p}}^{p(n)} (p(n) - p^*)^{n-1} f(p^*)dp^* - \int_{p(n)}^{\overline{p}} (p^* - p(n))^{n-1} f(p^*)dp^* = 0. \tag{44}$$

We suppress here the asterisk from $\underline{p}^*$ and $\overline{p}^*$ and replace $p_t^e(n)$ by $p(n)$ to simplify notation. Rewriting (44) as a metric

$$\begin{aligned} &\lim_{n\to\infty} \left(\int_{\underline{p}}^{p(n)} (p(n) - p^*)^{n-1} f(p^*)dp^* \right)^{1/(n-1)} \\ &= \lim_{n\to\infty} \left(\int_{p(n)}^{\overline{p}} (p^* - p(n))^{n-1} f(p^*)dp^* \right)^{1/(n-1)}, \end{aligned}$$

[14] We assume $n \geq 1$ throughout, because this implies realistic market maker preferences. However, (21) can be solved for any $n \geq 0$. In particular, $p(0)$ is the mode of $f(\cdot)$ when $s = 0$, or the highest density (connected) region when $s > 0$. For $n \notin \{1, 2, \infty\}$ no explicit solution exists, and for $n > 25$ even obtaining numerical solutions creates difficulty for non-trivial distribution functions $f(\cdot)$.

which after taking the limit degenerates to the sup norm

$$\sup_{p^* \in [\underline{p}, p(\infty)]} (p(\infty) - p^*) = \sup_{p^* \in [p(\infty), \overline{p}]} (p^* - p(\infty)),$$

gives

$$p(\infty) = \frac{\underline{p} + \overline{p}}{2}. \tag{45}$$

Thus, by monotonicity (45) solves (44) for $n \to \infty$. *Q.E.D.*

A.5.2 Effect of Risk Aversion on Optimal Price

Here we show that high risk aversion pushes the optimal price toward the midpoint of the support. In other words, if $f(\cdot)$ is without loss of generality right-skewed, then $p(n)$ is increasing in n, $\forall n \geq 1$. First, note that $p(n)$, $p(n) \in [\underline{p}, \overline{p}]$, is continuous. If $\underline{p}$ or $\overline{p}$ are infinite, we replace these bounds with a function of n, thereby making the domain of p compact. As $f(\cdot)$ and all components of the integral are continuous functions, the theorem of the maximum gives continuity of $p(n)$.

Next, to evaluate how the optimal price $p(n)$ responds to changes in risk aversion n, take the total differential of (44) and rearrange to obtain

$$\frac{dp(n)}{dn} = \frac{1}{n-1} \times \left\{ -\int_{\underline{p}}^{p(n)} (p(n) - p^*)^{n-1} \ln(p(n) - p^*) f(p^*) dp^* \right.$$
$$\left. + \int_{p(n)}^{\overline{p}} (p^* - p(n))^{n-1} \ln(p^* - p(n)) f(p^*) dp^* \right\} \Big/$$
$$\left\{ \int_{\underline{p}}^{p(n)} (p(n) - p^*)^{n-2} f(p^*) dp^* + \int_{p(n)}^{\overline{p}} (p^* - p(n))^{n-2} f(p^*) dp^* \right\}. \tag{46}$$

In the following argument we use that $f(\cdot)$ is monotone and assume without loss of generality that $f(\cdot)$ is monotonically decreasing. This means $f(\cdot)$ is right-skewed on $[\underline{p}, \overline{p}]$, which occurs if the market maker has some information that the strong form efficient price has increased. Under this assumption (46) is positive. To see this, note first that both terms in the denominator are positive. To economize notation we replace $p \equiv p(n)$, $d \equiv p(n) - \underline{p}$

and $x \equiv p^*$. The numerator can be broken up into three parts:

$$\begin{aligned}
& - \int_{\underline{p}}^{p} (p-x)^{n-1} \ln(p-x) f(x)dx + \int_{p}^{\overline{p}} (x-p)^{n-1} \ln(x-p) f(x)dx \\
= & -\int_{p-d}^{p-1} (p-x)^{n-1} \ln(p-x) f(x)dx + \int_{p+1}^{p+d} (x-p)^{n-1} \ln(x-p) f(x)dx \\
- & \int_{p-1}^{p} (p-x)^{n-1} \ln(p-x) f(x)dx + \int_{p}^{p+1} (x-p)^{n-1} \ln(x-p) f(x)dx \\
+ & \int_{p+d}^{\overline{p}} (x-p)^{n-1} \ln(x-p) f(x)dx.
\end{aligned} \tag{47}$$

The first term, which exists only for $d > 1$, gives

$$\begin{aligned}
& - \int_{p-d}^{p-1} (p-x)^{n-1} \ln(p-x) f(x)dx + \int_{p+1}^{p+d} (x-p)^{n-1} \ln(x-p) f(x)dx \\
= & -\int_{p+1}^{p+d} (x-p)^{n-1} \ln(x-p) f(2p-x)dx \\
& +\int_{p+1}^{p+d} (x-p)^{n-1} \ln(x-p) f(x)dx \\
= & \int_{p+1}^{p+d} (x-p)^{n-1} \ln(x-p) \left[-f(2p-x) + f(x)\right] dx \\
\geq & \int_{p+1}^{p+d} (x-p)^{n-1} \ln(d) \left[-f(2p-x) + f(x)\right] dx \\
= & -\int_{p-d}^{p-1} (p-x)^{n-1} \ln(d) f(x)dx + \int_{p+1}^{p+d} (x-p)^{n-1} \ln(d) f(x)dx.
\end{aligned} \tag{48}$$

The second term is for $d \geq 1$

$$
\begin{aligned}
& - \int\limits_{p-1}^{p} (p-x)^{n-1} \ln(p-x) f(x) dx + \int\limits_{p}^{p+1} (x-p)^{n-1} \ln(x-p) f(x) dx \\
= & - \int\limits_{p}^{p+1} (x-p)^{n-1} \ln(x-p) f(2p-x) dx \\
& + \int\limits_{p}^{p+1} (x-p)^{n-1} \ln(x-p) f(x) dx \\
= & \int\limits_{p}^{p+1} (x-p)^{n-1} \ln(x-p) \left[f(x) - f(2p-x) \right] dx \geq 0. \qquad (49)
\end{aligned}
$$

For $d < 1$ the last inequality of the calculations for the second term is instead

$$
\begin{aligned}
& \int\limits_{p}^{p+1} (x-p)^{n-1} \ln(x-p) \left[f(x) - f(2p-x) \right] dx \\
\geq & \int\limits_{p}^{p+d} (x-p)^{n-1} \left[f(x) - f(2p-x) \right] dx \ln(d) \geq 0. \qquad (50)
\end{aligned}
$$

And for the last term we can write

$$
- \int\limits_{p+d}^{\overline{p}} (x-p)^{n-1} \ln(x-p) f(x) dx > - \int\limits_{p+d}^{\overline{p}} (x-p)^{n-1} \ln(d) f(x) dx. \qquad (51)
$$

Using (48), (49), and (51), (47) becomes

$$
(47) \quad > \left[- \int\limits_{p-d}^{p-1} (p-x)^{n-1} f(x) dx + \int\limits_{p+1}^{\overline{p}} (x-p)^{n-1} f(x) dx \right] \ln(d)
$$

$$
\begin{aligned}
&> \left[-\int_{p-d}^{p-1} (p-x)^{n-1} f(x)dx - \int_{p-1}^{p} (p-x)^{n-1} f(x)dx \right. \\
&+ \left. \int_{p}^{p+1} (x-p)^{n-1} f(x)dx + \int_{p+1}^{\overline{p}} (x-p)^{n-1} f(x)dx \right] \ln(d) \\
&= \left[-\int_{p-d}^{p} (p-x)^{n-1} f(x)dx + \int_{p}^{\overline{p}} (x-p)^{n-1} f(x)dx \right] \ln(d) \\
&= 0,
\end{aligned}
$$

where the inequality follows from the monotonicity of $f(\cdot)$, and the last equality follows from the first order condition (44).

Likewise, for $d < 1$, using (50) we have

$$
\begin{aligned}
(47) &> \left[-\int_{p-d}^{p} (p-x)^{n-1} f(x)dx + \int_{p}^{\underline{p}} (x-p)^{n-1} f(x)dx \right] \ln(d) \\
&= 0.
\end{aligned}
$$

Therefore the numerator is positive and

$$
\frac{dp(n)}{dn} > 0
$$

for right-skewed distributions. Combining this with the fact that $p(1) = \text{Median}(p^*)$ and $p(\infty) = \text{Midsupport}(p^*)$ we conclude that $p(n)$ monotonically increases from the median to the midpoint of the support of the efficient price distribution $f(\cdot)$, if $f(\cdot)$ is right-skewed. Analogously, for left-skewed $f(\cdot)$, $p(n)$ monotonically *de*creases from the median to the midpoint of the support.

A.6 Proof of Proposition 6

Proposition (Cross-correlation under market maker information)

If the distribution of the expected latent price with ex-ante support $[\underline{p}_t^*, \overline{p}_t^*]$ satisfies

$$
\left[\frac{\underline{p}_t^* + \overline{p}_t^*}{2} - p_{t-1}^* \right] \text{sgn}(\varepsilon_t) > s + \frac{\sigma}{E(|\varepsilon_t|)},
$$

then $\exists n_0 > 1$ such that $\forall n > n_0$ it holds that $Corr(\Delta p_t^*, \Delta u_t) > 0$.

Proof:

The new information each period now consists of two parts: First, as before, information about p_{t-1}^*, and second extra information about Δp_t^*. To be specific, we assume that this extra information is the direction of the latent price change $\{\text{sgn}(\varepsilon_t)\}$. If the distribution of expected latent price changes at the beginning of each period is the same, we can write the market maker response to this extra information as $R(\text{sgn}(\varepsilon_t)) = R\,\text{sgn}(\varepsilon_t)$. From (20)

$$\begin{aligned} E\left(\Delta p_t \varepsilon_t\right) &= E\left[\left(p_{t-1}^* + sq_t + R(\cdot) - p_{t-1}\right)\varepsilon_t\right] \\ &= \frac{1}{2}E\left[(sq_t + R(\cdot))\varepsilon_t \,|\varepsilon_t > 0\right] + \frac{1}{2}E\left[(sq_t - R(\cdot))\varepsilon_t \,|\varepsilon_t < 0\right] \\ &= RE\left(|\varepsilon_t|\right) + sE\left(q_t\varepsilon_t\right). \end{aligned} \tag{52}$$

Plugging (52) with $E\left(\Delta p_t \Delta p_t^*\right) = \sigma E\left(\Delta p_t \varepsilon_t\right)$ into (10) implies that the contemporaneous cross-covariance is positive if and only if

$$R > \frac{\sigma - sE\left(q_t\varepsilon_t\right)}{E\left(|\varepsilon_t|\right)}. \tag{53}$$

Because $E(q_t\varepsilon_t) > -E(|\varepsilon_t|)$ we have as sufficient condition

$$R > s + \frac{\sigma}{E\left(|\varepsilon_t|\right)}. \tag{54}$$

To satisfy (54) we need for $p_t^e = p_{t-1}^* + R(\cdot) = p_{t-1}^* + R\,\text{sgn}(\varepsilon_t)$ that

$$p_t^e \begin{cases} > p_{t-1}^* + s + \frac{\sigma}{E(|\varepsilon_t|)} & \text{for } \varepsilon > 0 \\ < p_{t-1}^* - s - \frac{\sigma}{E(|\varepsilon_t|)} & \text{for } \varepsilon < 0. \end{cases}$$

From Proposition 5 for any $p(n) \in \left[\text{Median}(p); \frac{p+\bar{p}}{2}\right]$ there is a risk aversion level n such that market makers will – after observing the signal $\{\text{sgn}(\varepsilon_t)\}$ – quote this price as midprice p_t^e. Therefore, for all distributions $f(p^*)$ which satisfy (22), a sufficiently large n leads to a market maker response which satisfies (53) and thus to a positive contemporaneous cross-covariance. *Q.E.D.*

A.7 Proof of Proposition 7

Proposition (Semi-strong form cross-correlation, one period model)

The contemporaneous cross-correlation is

$$Corr(\Delta\tilde{p}_t^e, \Delta\tilde{u}_t) = \frac{2\phi\lambda - \sigma E(q_t\varepsilon_t)}{\sqrt{\sigma^2 - 2\sigma\lambda E(q_t\varepsilon_t) + 2\phi\lambda^2}} \frac{\mathrm{sgn}(s-\lambda)}{\sqrt{2\phi}}.$$

The cross-correlation at displacement one equals

$$Corr(\Delta\tilde{p}_{t-1}^e, \Delta\tilde{u}_t) = \frac{-\phi\lambda}{\sqrt{\sigma^2 - 2\sigma\lambda E(q_t\varepsilon_t) + 2\lambda^2}} \frac{\mathrm{sgn}(s-\lambda)}{\sqrt{2\phi}}.$$

All cross-correlations at higher displacements are zero.

Proof:

The expressions for the cross-correlations follow directly from their multiperiod counterparts. In the setup of Section 4.1 the semi-strong form efficient price has the unconditional variance (33) and the corresponding noise has an unconditional variance of

$$\begin{aligned} Var(\Delta\tilde{u}_t) &= \frac{1}{T}\sum_{t=0}^{T-1} Var(\Delta\tilde{u}_t) \\ &= \frac{1}{T}\left\{\sum_{i=0}^{T-1}\left[\phi_i(\lambda_i - s_i)^2 + \phi_{i-1}(\lambda_{i-1} - s_{i-1})^2\right]\right. \\ &- \left. 2\sum_{i=1}^{T-1} E(q_t q_{t-1})(\lambda_i - s_i)(\lambda_{i-1} - s_{i-1})\right\}. \end{aligned} \tag{55}$$

The contemporaneous cross-correlation is

$$Corr(\Delta\tilde{p}_t^e, \Delta\tilde{u}_t) = \frac{Cov(\Delta\tilde{p}_t, \Delta\tilde{u}_t)}{\sqrt{Var(\Delta\tilde{p}_t^e)Var(\Delta\tilde{u}_t)}}.$$

where $Cov(\Delta\tilde{p}_t^e, \Delta\tilde{u}_t)$ is given by (23). All other cross-correlation can be obtained analogously.

For $T = 1$, spread and adverse selection parameter are constants, i.e. $s_t = s$ and $\lambda_t = \lambda$ $\forall t$, and the variance terms (33) and (55) simplify radically to

$$Var(\Delta\tilde{p}_t^e) = \sigma^2 - 2\sigma\lambda E(q_t\varepsilon_t) + 2\phi\lambda^2,$$

$$Var(\Delta\tilde{u}_t) = 2\phi(s-\lambda)^2,$$

where we have used that q_t is serially uncorrelated. Finally, from (23) and (25), we get for

$T = 1$ the covariances

$$Cov(\Delta\tilde{p}_t^e, \Delta\tilde{u}_t) = (s - \lambda)\left[2\phi\lambda - \sigma E(q_t\varepsilon_t)\right]$$

and

$$Cov(\Delta\tilde{p}_{t-1}^e, \Delta\tilde{u}_t) = \phi\lambda(\lambda - s).$$

Combining these with the variances immediately gives the cross-correlations stated in proposition 7. *Q.E.D.*

B Derivation of Noise-Robust Estimators for Realized Volatility

The estimators are based on the observable moments $\gamma_i \equiv E(\Delta p_t \Delta p_{t-i})$ and $\gamma_i^* \equiv E(\Delta p_t q_{t-i})$. They are based on the assumption that the latent price changes every period ($T = 1$), and remains unobserved for one or more periods, depending on the noise specification. The probability of no latent price change has measure zero. The two i.i.d. innovations to noise are $v_t \overset{iid}{\approx} (0, \eta_v^2)$ and $w_t \overset{iid}{\approx} (0, \eta_w^2)$, where $E(v_s w_t) = 0\ \forall s, t$. All noise specification are given in the second and third column of Table 3.

B.1 General Result

Consider the strong form efficient returns given by (7) and strong form noise given by (3). Then we have

$$\gamma_0 = \sigma^2 + 2\eta_v^2 - 2E(u_t u_{t-1}) + 2E(\Delta p_t^* u_t) - 2E(\Delta p_t^* u_{t-1}),$$

$$\gamma_1 = -\eta_v^2 + 2E(u_t u_{t-1}) - E(u_t u_{t-2}) + E(\Delta p_t^* u_{t+1}) - E(\Delta p_t^* u_t) + E(\Delta p_t^* u_{t-1}) - E(\Delta p_t^* u_{t-2})$$

and $\forall i \geq 2$

$$\begin{aligned} \gamma_i &= 2E(u_t u_{t-i}) - E(u_t u_{t-1-i}) \\ &+ E(\Delta p_t^* u_{t+i}) - E(\Delta p_t^* u_{t-1+i}) + E(\Delta p_t^* u_{t-i}) - E(\Delta p_t^* u_{t-1-i}). \end{aligned}$$

Recursively plugging in and solving for σ^2, we get

$$\sigma^2 = \gamma_0 + 2\sum_{i=1}^{k}\gamma_i - 2E(u_t u_{t-k}) + 2E(u_t u_{t-k-1}) - 2E(\Delta p_t^* u_{t+k}) + 2E(\Delta p_t^* u_{t-k-1}).$$

Absent insider information, $E(\Delta p_t^* u_{t-i}) = 0 \ \forall i \geq 1$, and therefore

$$\sigma^2 = \gamma_0 + 2\sum_{i=1}^{k}\gamma_i - 2E(u_t \Delta u_{t-k}) - 2E(\Delta p_t^* u_{t+k}). \tag{56}$$

If – as in the upper four rows of Table 3 – the market microstructure noise is asymptotically uncorrelated, i.e. $\lim_{k\to\infty} E(u_t \Delta u_{t-k}) = 0$ and $\lim_{k\to\infty} E(\Delta p_t^* u_{t+k}) = 0$, Equation (56) simplifies to

$$\sigma^2 = \gamma_0 + 2\sum_{i=1}^{\infty}\gamma_i. \tag{57}$$

Equation (57) does not hold in market maker inventory models, as the ones in the bottom row of Table 3. There, the noise follows a unit root process with $\Delta u_t = \alpha q_t$ so that $E(u_t \Delta u_{t-i}) = \alpha^2 \ \forall i$. In this case autocovariances alone are not sufficient, and σ^2 is given by Equation (56).

B.2 I.i.d. Measurement Error

$$\gamma_0 = \sigma^2 + 2\eta_v^2$$

$$\gamma_1 = -\eta_v^2$$

Therefore

$$\sigma^2 = \gamma_0 + 2\gamma_1.$$

B.3 Autoregressive Noise

$$\gamma_0 = \sigma^2 + \frac{2}{1+\alpha}\eta_v^2$$

$$\gamma_i = \alpha^{i-1}\frac{\alpha-1}{\alpha+1}\eta_v^2 \quad \forall i \geq 1$$

Therefore

$$\sigma^2 = \gamma_0 + 2\frac{\gamma_1^2}{\gamma_1 - \gamma_2}.$$

B.4 Autoregressive Noise with i.i.d. Measurement Error

$$\gamma_0 = \sigma^2 + \frac{2}{1+\alpha}\eta_v^2 + 2\eta_w^2$$

$$\gamma_1 = \frac{\alpha-1}{\alpha+1}\eta_v^2 - \eta_w^2$$

$$\gamma_i = \alpha^i \frac{\alpha-1}{\alpha+1}\eta_v^2 \quad \forall i \geq 2$$

Therefore

$$\sigma^2 = \gamma_0 + 2\gamma_1 + 2\frac{\gamma_2^2}{\gamma_2 - \gamma_3}.$$

B.5 Linear Noise Decay

$$\gamma_0 = \sigma^2 + \frac{S+1}{S}\alpha^2\eta_v^2$$

$$\gamma_i = -\frac{i}{S^2}\alpha^2\eta_v^2 \quad \text{for } i \in \{1 \dots S\}$$

$$\gamma_i = 0 \quad \forall i > S$$

$$\sum_{i=1}^{\infty}\gamma_i = \sum_{i=1}^{S}\gamma_i = -\frac{S+1}{2S}\alpha^2\eta_v^2 = \frac{S(S+1)}{2}\gamma_1$$

Therefore

$$\sigma^2 = \gamma_0 + 2\sum_{i=1}^{\infty}\gamma_i,$$

or

$$\sigma^2 = \gamma_0 + 2\sum_{i=1}^{S}\gamma_i = \gamma_0 + S(S+1)\gamma_1,$$

where S is either known or estimated by $\hat{S} = \sqrt{\frac{2}{\gamma_1}\sum_{i=1}^{\infty}\gamma_i + \frac{1}{4}} - \frac{1}{2}$. If S is estimated, this estimator is identical to the constant realized kernel estimator (29).

B.6 Market Maker Inventory from Noise Trading

$$\gamma_0 = \sigma^2 + \alpha^2$$

$$\gamma_1 = 0$$

$$\gamma_0^* = \alpha$$

$$\gamma_1^* = 0$$

Therefore

$$\sigma^2 = \gamma_0 - {\gamma_0^*}^2.$$

B.7 Dependent Measurement Error

$$\gamma_0 = \sigma^2 + 2\alpha(1+\alpha)\sigma^2 + 2\eta_v^2$$

$$\gamma_1 = -\alpha(1+\alpha)\sigma^2 - \eta_v^2$$

$$\gamma_2 = 0$$

Therefore

$$\sigma^2 = \gamma_0 + 2\gamma_1.$$

Likewise, for the bid-ask bounce from informed traders, we have

$$\gamma_0 = \sigma^2 + 2\alpha^2 + 2\alpha\sigma E\left(|\varepsilon_t|\right)$$

$$\gamma_1 = -\alpha^2 - \alpha\sigma E\left(|\varepsilon_t|\right)$$

$$\gamma_2 = 0$$

B.8 Nonstrategic Incompletely Informed Traders

$$\gamma_0 = \sigma^2 + \frac{2\beta}{1+\alpha}\left(\sigma^2(1+\alpha+\beta) + \beta\eta_v^2\right)$$

$$\gamma_i = -\frac{\beta^i(1-\alpha)}{1+\alpha}\left(\sigma^2(1+\alpha+\beta) + \beta\eta_v^2\right) \quad \forall i \geq 1$$

Therefore

$$\sigma^2 = \gamma_0 + 2\frac{\gamma_1^2}{\gamma_1 - \gamma_2}.$$

B.9 Autoregressive Noise with One-Period Private Information

$$\gamma_0 = \sigma^2 + 2\beta(1+\beta)\sigma^2 + \frac{2}{1+\alpha}\eta_v^2 + 2\eta_w^2$$

$$\gamma_1 = -\beta(1+\beta)\sigma^2 - \frac{1-\alpha}{1+\alpha}\eta_v^2 - \eta_w^2$$

$$\gamma_i = -\alpha^{i-1}\frac{1-\alpha}{1+\alpha}\eta_v^2 \quad \forall i \geq 2$$

Therefore

$$\sigma^2 = \gamma_0 + 2\gamma_1 + 2\frac{\gamma_2^2}{\gamma_2 - \gamma_3}.$$

Likewise, for nonstrategic informed traders, we have

$$\gamma_0 = \sigma^2 + \frac{2\alpha(1+\alpha+\beta)}{1+\beta}\sigma^2 + 2\eta_v^2$$

$$\gamma_1 = -(1-\beta)\frac{\alpha(1+\alpha+\beta)}{1+\beta}\sigma^2 - \eta_v^2$$

$$\gamma_i = -\beta^{i-1}(1-\beta)\frac{\alpha(1+\alpha+\beta)}{1+\beta}\sigma^2 \quad \forall i \geq 2$$

The Easley and O'Hara (1992)-type learning reflected in Equation (32) is a special case of this. Set $\alpha = -\sigma$ and $\beta = e^{-r}$, and drop the exogenous noise process v_t. Then $\forall i \geq 0$

$$\gamma_i = \frac{e^r - 1}{e^r + 1} e^{-ir}\sigma^2,$$

therefore $e^r = \frac{\gamma_0}{\gamma_1}$ and

$$\sigma^2 = \gamma_0 + \frac{2\gamma_1\gamma_0}{\gamma_0 - \gamma_1}.$$

Using $\frac{\gamma_1}{\gamma_2} = \frac{\gamma_0}{\gamma_1}$ reveals that this expression is equivalent to $\gamma_0 + \frac{2\gamma_1^2}{\gamma_1 - \gamma_2}$.

B.10 Strategic Informed Traders

$$\gamma_0 = \sigma^2 + \frac{(2+\alpha)S + \alpha}{S}\alpha\sigma^2 + \frac{S+1}{S}\alpha^2\eta_v^2$$

$$\gamma_i = -\frac{S + i\alpha}{S^2}\alpha\sigma^2 - \frac{i}{S^2}\alpha^2\eta_v^2 \quad \text{for } i \in \{1 \ldots S\}$$

$$\gamma_i = 0 \quad \forall i > S$$

$$\sum_{i=1}^{\infty}\gamma_i = \sum_{i=1}^{S}\gamma_i = -\frac{(2+\alpha)S+\alpha}{2S}\alpha\sigma^2 - \frac{S+1}{2S}\alpha^2\eta_v^2 = \frac{S-1}{2}\alpha\sigma^2 + \frac{S(S+1)}{2}\gamma_1$$

Therefore

$$\sigma^2 = \gamma_0 + 2\sum_{i=1}^{\infty}\gamma_i,$$

or

$$\sigma^2 = \gamma_0 + 2\sum_{i=1}^{S}\gamma_i = \gamma_0 + 2\left[\frac{S(S-1)}{2}\alpha\sigma^2 + \frac{S(S+1)}{2}\gamma_1\right],$$

which implies

$$\sigma^2 = \gamma_0 + S(3-S)\gamma_1 + S(S-1)\gamma_2,$$

where S is either known or estimated by $\hat{S} = \sqrt{\left(\frac{3\gamma_1-\gamma_2}{2(\gamma_2-\gamma_1)}\right)^2 + \frac{2}{\gamma_2-\gamma_1}\sum_{i=1}^{\infty}\gamma_i} - \frac{3\gamma_1-\gamma_2}{2(\gamma_2-\gamma_1)}$. If S is estimated, this estimator is identical to the constant realized kernel estimator (29).

B.11 Market Maker Inventory from Informed Trading

$$\gamma_0 = \sigma^2 + \alpha^2 + 2\alpha\sigma E\left(|\varepsilon_t|\right)$$

$$\gamma_1 = 0$$

$$\gamma_0^* = \alpha + \sigma E\left(|\varepsilon_t|\right)$$

$$\gamma_1^* = 0$$

If $\varepsilon_t \stackrel{iid}{\sim} N(0,1)$ we have

$$\gamma_0^* = \alpha + \sigma\sqrt{\frac{2}{\pi}}$$

and therefore

$$\sigma^2 = \frac{\pi}{\pi-2}\left(\gamma_0 - {\gamma_0^*}^2\right).$$