Local Mispricing and Microstructural Noise: A Parametric Perspective*

Torben G. Andersen† Ilya Archakov‡
Gökhan Cebiroglu§ Nikolaus Hautsch¶

This Version: February 2019

Abstract

We propose a generalization of the classical "martingale-plus-noise" model for high-frequency prices that accommodates both price endogeneity and an error correction mechanism originating from prevailing mispricing. The strength of price reversion relative to the noise-to-signal ratio determines two fundamental regimes of how mispricing is incorporated yielding a bias in standard realized variance estimates and triggering return serial correlations of different signs. We derive the model's properties, discuss parameter identifiability and locally estimate the model based on mid-quote returns of the NASDAQ 100 constituents. We show the existence of local regimes of positive and negative return autocorrelation arising from lagged feedback effects and sluggish price adjustment. We moreover document intraday periodicities in the speed of price reversion and noise-to-signal ratios. The model links notions from high- frequency statistics and market microstructure theory and thus opens up new paths for efficient (local) volatility estimation.

Keywords: volatility estimation; market microstructure noise; price reversal; momentum

---

*We thank participants of the "Jan Mossin Memorial Symposium on Financial Markets", Norwegian School of Economics, Bergen, 2016, the Guangzhou 2016 Symposium on Financial Engineering and Risk Management, the 9th Annual Meeting of the Society for Financial Econometrics, University of Aarhus, 2016, the Conference on "New Developments in Measuring and Forecasting Financial Volatility", Durham, USA, 2016, and seminars at Humboldt University, Berlin, and Universitat Pompeu Fabra, Barcelona for comments and discussions . Andersen acknowledges support from CREATES (DNRF78), funded by the Danish National Research Foundation. Hautsch acknowledges financial support from the Wiener Wissenschafts-, Forschungs- und Technologiefonds (WWTF).

†Kellogg School of Management, Northwestern University, NBER and CREATES. Email: t-andersen@kellogg.northwestern.edu.

‡Department of Statistics and Operations Research, University of Vienna. Email: ilya.archakov@univie.ac.at.

§Credit Suisse, London. Email: goekhan.cebi@gmail.com. The views, thoughts, and opinions expressed in this article belong solely to the author, and are not necessarily shared by the authors employer.

¶Department of Statistics and Operations Research as well as Research Platform "Data Science @ Uni Vienna", University of Vienna and CFS. Email: nikolaus.hautsch@univie.ac.at.

trading, contrarian trading

JEL classification: C58, C32, G14

1 Introduction

A major theme in the recent financial econometrics literature is the efficient extraction of daily or intra-daily volatility measures from high-frequency asset price data. A common starting point is to assume that micro-level prices, or quotes, are governed by an underlying latent semi-martingale process which, however, is polluted by noise. This noise component is typically attributed to market microstructure frictions, causing observed micro-level prices to deviate from the efficient (semi-)martingale process.

It is well known that this so-called "market microstructure noise" (henceforth MMN) induces autocorrelation in high-frequency returns and may cause substantial biases in volatility estimates if one utilizes returns on very high frequencies. A classical result is that discretely sampled high-frequency returns, generated by a (semi-)martingale process polluted by i.i.d. noise, will display negative serial correlation and follow an MA(1) process. This relationship is already described by Roll (1984) in the context of transaction returns subject to bid-ask bounce effects. In a fully parametric setting, this property can be exploited to construct estimators of both the "efficient," or fundamental variance, and the MMN variance, see, e.g., Zhang et al. (2005).

Hansen and Lunde (2006) show, however, that the dependence patterns in high-frequency returns typically are not consistent with an MA(1) representation, but tend to follow more general processes. This is particularly true if the high-frequency observations originate from the midpoint of bid and ask quotes (mid-quotes), rather than from transaction prices. Moreover, the former are less impacted by the bid-ask bounce and, almost by construction, provide a closer approximation to the underlying "efficient" price. Given that the negative serial correlation stemming from bid-ask bounces is mitigated in mid-quote returns, one may wonder if the standard discrete-time representation of high-frequency asset returns, as following a random walk plus i.i.d. noise process, remains a sensible first-order approximation to the price dynamics. Figure 1 displays so-called volatility signature plots for four different stocks on certain days during our sample. The upward sloping plots (for movement towards higher frequencies, i.e., closer to the origin) in the left panels correspond to the type of biases observed in the presence of traditional i.i.d. noise. The negative serial correlation creates excessive return variability, and the realized volatility (RV) measures diverge upward for the highest sampling frequencies. However, our sample of NASDAQ 100 stocks contains many examples of the exact opposite behavior. This is exemplified in the right panels, where the signature plots drop sharply as we approach the highest sampling frequencies, suggesting an underlying significant *positive*

return serial correlation. In fact, across the full cross-section of stocks, we find positive return dependence to be roughly as common as negative dependence, and it is represents one of the key empirical facts we scrutinize in this paper. It suggests that, in general, MMN cannot be i.i.d., but must have more general properties that accommodate *both* positive and negative serial correlation with these properties changing across local periods. Moreover, these findings suggest that we may also need to allow for the possibility that the noise is endogenous, i.e., correlated with fundamental price innovations.

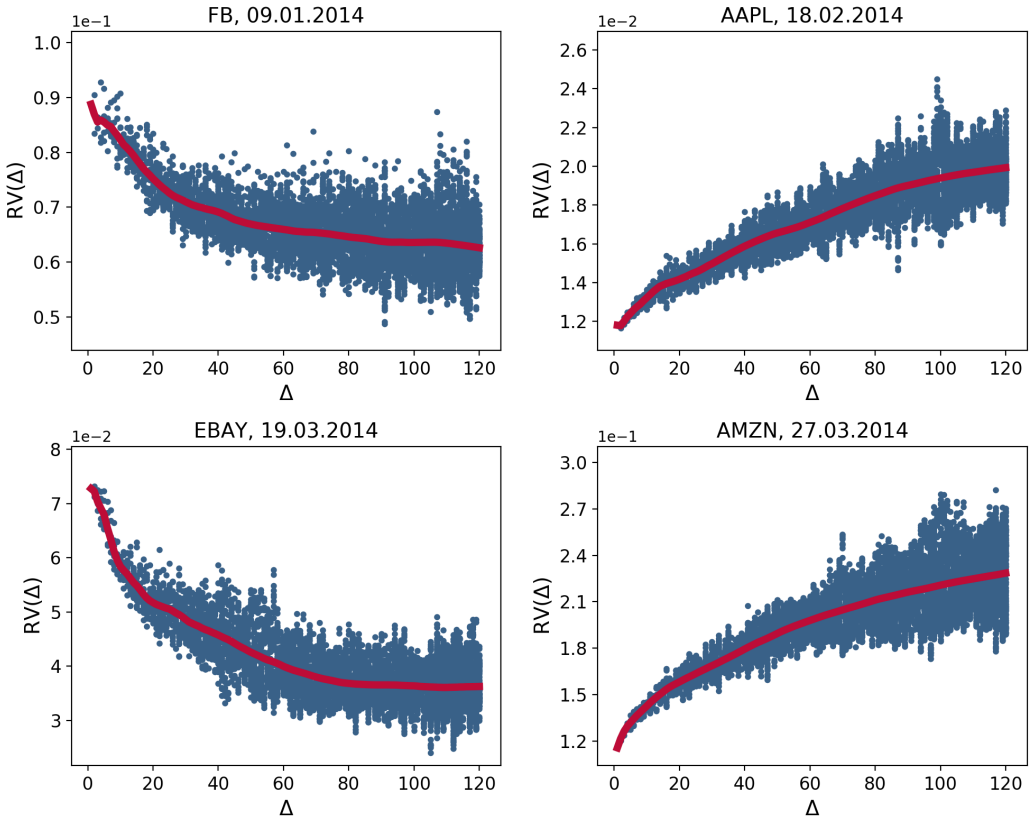

Figure 1: The figure depicts signature plots for the Realized Variances (RV) computed for 4 stocks on selected trading days. The left panels illustrate the cases with downward sloping signature plots and the examples of upward sloping signature plots are presented on the right side. For each considered sampling frequency Δ (from 1sec to 120sec with a second step), we compute RV on multiple price "grids" resulted by shifting the initial price observation by one second up to Δ seconds (thus, for Δ=1sec we have one "grid" and for Δ=120sec we receive 120 "grids"). Blue dots correspond to daily RV estimates computed for different sampling frequencies and for each price "grid" considered at a given frequency. Red solid lines represent the average values of RV obtained across all "grids" (sub-sampled RV).

The literature on high-frequency-based volatility estimation has responded by constructing modified procedures that allow for serial correlation of unknown form. The identification of the efficient price variation is obtained through the assumption that the systematic serial correlation stems from a stationary noise process. It follows that consistent estimators may concentrate on the long-range variation, exploiting techniques along the lines of the long-run variance

estimators in econometrics, as developed by Newey and West (1987) or Andrews (1991).[1] This nonparametric reduced-form approach is agnostic about the properties of the noise and focuses exclusively on estimating the quadratic variation of the random walk component.

Consequently, the implicit notion in the high-frequency volatility literature remains that of a latent semi-martingale process whose value is known at all times by the informed market participants. This requirement is obviously stringent for returns sampled at very high frequencies. The investor information set may not be perfect, and it may not be updated instantaneously, but rather be subject to information acquisition and processing delay. That is, investors may operate with slightly stale or incomplete information. For example, they may not be able to determine whether an increase in the asset price signals is due to the arrival of positive news or rather stems from the random arrival of net buying orders from liquidity traders (i.e., noise). In the former case, the price likely continues on an upward trajectory as more investors learn of the news or infer its presence from the trading activity, while, in the second case, the price will tend to revert back to the prior level. While an astute trader may avoid *systematic* mispricing of the asset, this only implies that the valuation is unbiased *on average*. Hence, during an episode with sustained imbalances in net liquidity demand or an unusual intensity of privately informed trading, even fully rational agents do not possess perfect information and they will err in their instantaneous asset valuation. On the other hand, rationality does imply that simple momentum and contrarian strategies are unprofitable. The fact that we find positive and negative return autocorrelation patterns about equally common is consistent with the latter prediction.

The primary objective of this paper is to explore whether the high-frequency price dynamics displays features consistent with delays in the acquisition, processing or reaction to innovations in the underlying efficient price. Hence, we extend the classical martingale-plus-noise model by incorporating a notion of sequential learning or lagged response, which allows for a more structural treatment of the observed temporal dependence patterns and serves to build a bridge to the related market microstructure literature originating from Kyle (1985) and Vives (1995). Motivated by empirical observations along the lines of Figure 1, we develop a model for mid-quote prices, which deviates from a pure martingale-plus-noise assumption by allowing for temporal feedback effects due to local discrepancies between the observed and efficient price. Our starting point is that those deviations, arising from frictions such as market microstructure noise or liquidity effects, in turn generate forces that tend to reverse these inefficiencies. As a consequence, the price discovery process is subject to a constant interplay between counteracting forces, and corrective feedback plays a substantial role in explaining the temporal dependencies in mid-quote returns.

[1] For a review of nonparametric noise-robust return variation estimators, see, e.g., the survey of Andersen et al. (2010) or Ait-Sahalia and Jacod (2014).

One striking empirical finding is that the dependence structure in the mid-quote returns fluctuates substantially over time. To retain a simple and transparent framework that facilitates interpretation, we stipulate that the model only applies locally over a short time interval, and then allow the parameters to vary across these (non-overlapping) windows. Estimating the model separately across short intervals allows the changing information and trading environment in the market to guide the time-varying parameter estimates.[2] Furthermore, the use of a local approximating model with constant volatility enables us to sidestep the difficult task of modeling the intraday patterns and persistent dynamics of the market activity, which almost inevitably would result in severe misspecification. Instead, we can focus the model on the critical features that distinguish it from existing stylized microstructure models and compare their performance across a large number of stocks and time intervals.

The main implication of our new representation is that the temporal dependence structure of observed returns arises from two separate channels. First, in line with the classical literature, market prices are disturbed by random noise arising due to market microstructure frictions or noise trading unrelated to information. The second channel is due to past "mispricing", i.e., deviations between past observed and efficient prices. The speed by which observed prices react to inherent mispricing is governed by a parameter that reflects the degree or "efficiency" with which investors are able to decipher the market signals and swiftly identify innovations in the efficient price. Mispricing generates trading opportunities for informed traders in the spirit of Kyle (1985). Hence, informed arbitrage activity ensures that deviations between the observed and efficient price remain bounded and that the effects of non-informational shocks get mitigated, as informed agents impound their knowledge into observed prices through trading, and other agents infer the presence of information.

To allow for this type of features, the model incorporates an error correction component, reminiscent of the partial adjustment mechanism proposed by Amihud and Mendelson (1987) for modeling daily returns. The motivation for thinking along the lines of sequential learning in this context is inspired by Kyle (1985) and Vives (1995). Thus, we combine notions from high-frequency statistics and market microstructure theory. As a result, we establish the mechanism of temporal feedback as an essential component in models for high-frequency prices. We study the properties of the model in a simple discrete-time framework of (locally) constant volatility. In such a setting, the model depends on three key parameters, namely the volatility of the random walk process (the so-called "fundamental volatility"), the "noise level" corresponding to the variance of the noise process, and the informational efficiency parameter governing the speed of reversals due to mispricing. We show that such a three-parameter setting is quite flexible and

[2]This mimics the asymptotic devise often used in high-frequency volatility estimation where the inference theory is established under the assumption of constant volatility over local (and asymptotically shrinking) time intervals.

can capture a wide range of stochastic behavior in mid-quote returns. As an important special case, our framework contains the classical "martingale-plus-noise" model, where the noise component, however, is negatively correlated with the efficient price. Our temporal feedback mechanism, arising from local mispricing, therefore provides a natural explanation for the occurrence of so-called endogenous noise as analyzed, e.g., by Kalnina and Linton (2008).[3] Moreover, the proposed structural approach for capturing the dynamics between the observed and efficient prices is related to a few prior studies that seek to link the properties of noise more explicitly to the process of trading and the underlying microstructure, e.g., Diebold and Strasser (2013), Chaker (2013) and Li et al. (2016). However, none of these studies explicitly consider temporal feedback effects which we deem crucial in explaining the return dynamics observed in Figure 1.

We show that our framework allows for a natural interpretation of the underlying model components. In particular, the speed of price reversion interacts with the ratio of the noise versus fundamental innovation volatility – the noise-to-signal ratio – to determine the sign of the return autocorrelation. This relation let us classify the (time-varying) market environment into two fundamental regimes. In one regime, the impact of "mispricing" is overcome by the strength of the price reversal, induced by idiosyncratic noise shocks, so that the returns display negative serial correlation, or "contrarian" price behavior. In the other regime, the feedback from "mispricing" generates "momentum," or positive autocorrelation, in the observed returns. Moreover, the regimes have important consequences for standard volatility measures based on high frequency returns, e.g., Andersen et al. (2003). In the contrarian regime, the realized volatility will tend to exaggerate the underlying quadratic variation of the efficient price, inducing the well-known upward-shaped (for increasing sampling frequencies) volatility signature plots, e.g., Hansen and Lunde (2006). In the other regime, however, high-frequency sampling causes downward-shaped signature plots which are not readily explained within the traditional "martingale-plus-noise" setting, but arise naturally in our framework. As noted previously, Figure 1 provides empirical illustrations of both types of signature plots.

In its basic form, the model can be conveniently represented in a linear state-space form, allowing for standard pseudo-maximum likelihood estimation via the Kalman filter. Thus, we may readily construct a consistent local volatility estimator. Moreover, we develop a generalized representation in which specific parameter configurations allow us to encompass alternative models of interest, including the Amihud and Mendelson (1987) version of the lagged price adjustment hypothesis and the standard martingale-plus-noise specification, among others. Within such a framework we moreover provide a systematic discussion on the (parametric) identifiability of effects arising from mispricing and noise endogeneity. In particular, we

[3] See also Sheppard (2013) for a related concept to measure the speed of a market.

show that the signal-to-noise ratio and the covariance between market microstructure noise and efficient price increments are only partially identified. As such, our approach provides the first step towards the development of new high-frequency based volatility estimators that simultaneously retain a link to the existing market microstructure literature and thus allow for a more structural treatment of effects resulting from market microstructure noise. These insights open up new paths for the construction of more efficient (local) volatility estimates.

Utilizing high-frequency data from NASDAQ trading, processed via the data platform LOBSTER[4], we estimate the model over local intraday intervals based on high-frequency mid-quote returns sampled at different frequencies. We confirm that the model parameters vary over time. In line with our model, we identify local regimes of contrarian and momentum behavior which are triggered by the (estimated) speed of price reversion in relation to the noise-to-signal ratio. These regimes are persistent, but generally quite short-lived, with slightly higher persistence for the more liquid assets. The local estimates for the speed of price reversion point towards a relatively high degree of informational efficiency, although the price updating clearly also displays an element of sluggishness. We also document the presence of strong intraday periodicities in the speed of price discovery and the relative proportion of informational variance to total variance. Overall, our evidence strongly favors our newly developed specification with lagged price adjustment relative to other popular stylized microstructure models.

The remainder of the paper is structured as follows. In Section 2, we provide comprehensive evidence for locally significant return autocorrelation of varying sign. We explore the persistence of these distinct regimes in serial correlation and undertake a variety of additional tests to verify the results. Section 3 introduces the underlying model. We first provide the general representation, and then explore the different features it captures, including the information feedback, the potential for delayed recognition of the news impact, and the presence of idiosyncratic noise. We then explore the implications for the (local) second order return moments. We also show that a number of existing models arise as special cases within our framework, so we inherently allow features discussed within specific models to impact the return dynamics at different points in time. We conlcude by considering the implications for realized volatility measurement. Section 4 focuses on a fundamental identification in our most general model specification. This problem arises naturally in models with latent fundamental information shocks combined with endogenous and serially correlated noise. We then explain our estimation strategy, which entails standard inference for most key model parameters but only achieves set identification for a set of coefficients governing the presence of endogeneity in the noise and the extent of idiosyncratic noise in the system. Our empirical results are presented in Section 5. We illustrate how the results appear for some representative trading days, but then focus on

[4]See `https://lobsterdata.com/`.

the aggregate results regarding return variation and autocorrelation patterns. In addition, we explore the pronounced systematic variation observed across the trading day. Finally, Section 6 concludes. Auxiliary results and proofs are collected in the Appendix.

2 Local Regimes of Return Autocorrelation

This section provides empirical evidence for the existence of short-lived regimes with persistent return autocorrelations of changing signs. We quantify return autocorrelations over short consecutive intervals using high-frequency mid-quote prices of the NASDAQ 100 equity index constituents. The underlying data stem from the LOBSTER database (https://lobsterdata.com), building on NASDAQ's historical TotalView-ITCH data. It provides information on all trade and quote activity on NASDAQ at nano-second time stamp precision. To keep the analysis manageable, we restrict our sample to the first three months of 2014 yielding a total of 61 trading days. For each day, we consider the full 6.5 hour trading period.

Given the objective to identify the properties of the mid-quote return series across short intervals (spanning 10-20 minutes), we primarily focus on the more liquid stocks with frequent quote revisions. Hence, we divide our sample into quintiles based on the average daily number of mid-quote revisions, with Group 1 representing the 20% of stocks with most frequent quote revisions and Group 5 standing for the 20% of stocks with the least frequent ones. While we concentrate our discussion on the most liquid groups, we convey a sense of robustness by reporting results for the full spectrum of stocks.

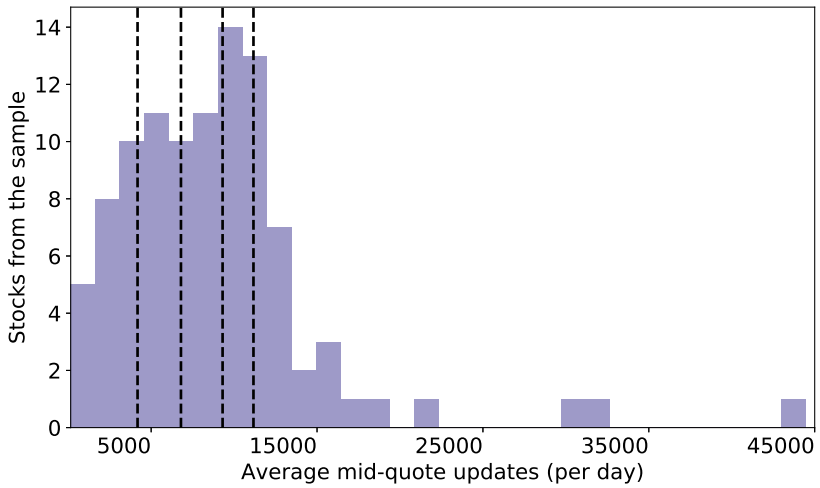

Figure 2: Empirical distribution of per-stock averages of daily mid-quote revisions for 100 considered stocks, based on the first 61 trading days in 2014. The vertical dashed lines corresponds to the 20'th, 40'th, 60'th and 80'th percentiles of the empirical distribution, thus partitioning the sample into five corresponding liquidity groups.

Figure 2 displays the distribution of daily mid-quote revisions, with the vertical lines provid-

ing the partitioning into the five liquidity groupings. The corresponding quintiles correspond to 17,485, 10,365, 8,166, 5,531 and 2,497 mid-quote revisions for the five groups, respectively. If we sample every 5 seconds (over 6.5 hours of trading), we have 4,680 return observations per day. Consequently, for a typical stock in Group 5, we usually have fewer quote revisions than intra-day returns, whereas this ratio is much more favorable for stocks in Groups 1-4.

2.1 Significance of Return Autocorrelations in Local Intervals

We face the challenge of identifying return autocorrelations over local intervals of unknown (and possibly) time-varying lengths. If individual stocks truly do undergo alternating periods with different degrees of return autocorrelation – ranging from significantly negative to significantly positive – but such episodes are fairly short, of varying duration and not directly observed, we need to adopt a methodology designed to detect this type of deviation from the null hypothesis of constant return autocorrelation.

On the one hand, autocorrelations computed over long intervals provide limited power against local fluctuations in the sign and size of serial correlations, as the values are averaged across regimes.Consequently, the autocorrelation statistics are biased towards zero, and regular test procedures lack power. On the other hand, we require a non-trivial number of return observations to obtain statistics with reasonable precision.

To address this bias-variance trade-off, we employ different test procedures and robustify our results by analyzing local intervals of different lengths (10, 15, and 20 minutes) and sampling at different, but relatively high frequencies, $\Delta = 2$, 3, and 5 seconds, guaranteeing a reasonably large number of observations within each interval. For the most illiquid stocks, this is an aggressive choice, and we may expect the findings for Group 5 to reflect the impact of microstructure effects and, especially, irregular trading. Nonetheless, we include findings across all stocks, time intervals, and sampling frequencies to convey a sense of the robustness.

Invoking the null hypothesis of zero return autocorrelation, it is natural to initially test whether an abnormal number of intervals generates excessively high absolute serial correlations. To formalize the procedure, denote the intra-day mid-quote log-return for a given stock over interval i by r_i, with i ranging across an equidistant grid $i = 1, \ldots, M$. Throughout the paper, we assume that the drift (unconditional mean) over short (intraday) intervals are negligible, imposing $\mathbb{E}[r_i] = 0$. We then split the intra-daily trading period into K_d local intervals (ranging from 10 to 20 minutes), with $N = \lfloor M/K_d \rfloor$ observations in each interval. For a sample with D trading days, we have a total of $K = DK_d$ local intervals in our sample. For $k = 1, \ldots, K$, the empirical (first-order) autocorrelation for interval k is $\hat{\rho}_k = \left(\sum_{i=1}^{N-1} r_{i,k} r_{i-1,k} \right) / \sum_{i=1}^{N} r_{i,k}^2$, where $r_{i,k}$ denotes the i-th midquote return in interval k.

To perform inference robust to heteroskedasticity and volatility clustering in the high-

frequency returns, stemming from stochastic volatility and intraday periodicity effects, we rely on the asymptotic results derived by Kokoszka and Politis (2011). Under standard (mild) assumptions on the moments of the return process and the dependence in the squared return process, they demonstrate that,

$$\sqrt{N}(\hat{\rho}_k - \rho_k) \xrightarrow{d} \mathcal{N}(0, V_k)\,,$$

where the asymptotic variance V_k is given by,

$$V_k = \frac{(N-1)^{-1}\sum_{i=1}^{N-1} r_{i,k}^2 r_{i-1,k}^2}{(N^{-1}\sum_{i=1}^{N} r_{i,k}^2)^2}\,.$$

Under the null hypothesis $H_0 : \rho_k = 0$, for all k, the local estimates $\hat{\rho}_k$ are noisy estimates of zero, stemming from asymptotically independent, non-identical Gaussian distributions. To gain power, we do not aim at testing the significance of local autocorrelations on each individual interval, but focus on *averages* of absolute (local) autocorrelations. Under the aymptotics given above, the absolute autocorrelations $|\hat{\rho}_k|$ are asymptotically distributed as a half-normal random variate with $\mathbb{E}[|\hat{\rho}_k|] = \sqrt{2V_k/(N\pi)}$ and $\mathbb{V}[|\hat{\rho}_k|] = (1-2/\pi)V_k/N$ for all k. Then, by a standard central limit theorem, under H_0, the *average* absolute autocorrelation $\bar{\rho} := K^{-1}\sum_k |\hat{\rho}_k|$, asymptotically follows the distribution,

$$\bar{\rho} \overset{a}{\sim} \mathcal{N}\left(\frac{1}{K}\sum_k \sqrt{\frac{2V_k}{N\pi}},\quad \frac{1}{K^2 N}\left(1-\frac{2}{\pi}\right)\sum_k V_k\right). \tag{1}$$

Hence, by constructing a test for $\bar{\rho}$ using the quantiles associated with the distribution in equation (1), we obtain a test with power against the alternative of local autocorrelations of varying signs. To test separately for positive and negative autocorrelations, we split the sample of K intervals into the subsamples I_- and I_+, consisting of intervals where $\hat{\rho}_k$ is positive or negative, respectively. Accordingly, *one-sided* average absolute autocorrelations are given by $\bar{\rho}_+ = K_+^{-1}\sum_{k\in I_+}|\hat{\rho}_k|$ and $\bar{\rho}_- = K_-^{-1}\sum_{k\in I_-}|\hat{\rho}_k|$, respectively, where $K_+ = |I_+|$ and $K_- = |I_-|$. The asymptotic distributions of $\bar{\rho}_+$ and $\bar{\rho}_-$ are obtained similarly as for $\bar{\rho}$. Hence, this test has power to separate average empirical autocorrelation estimates of given sign from noisy realizations of the distribution under the null.

Table 1 reports the corresponding test results. Choosing T and Δ trades off the estimation efficiency (increasing for low Δ and large T), biases due to the possible impact of market microstructure noise (worsens for small Δ) and the superposition of local autocorrelation regimes (problem for large T) versus the resulting power of the test. We report inference based

	T = 10 min							T = 15 min							T = 20 min						
	(I)		(I_+)		(I_-)			(I)		(I_+)		(I_-)			(I)		(I_+)		(I_-)		
Δ	$\#_{5\%}$	$\#_{1\%}$	$\#_{5\%}$	$\#_{1\%}$	$\#_{5\%}$	$\#_{1\%}$	$\bar{\pi}_+$	$\#_{5\%}$	$\#_{1\%}$	$\#_{5\%}$	$\#_{1\%}$	$\#_{5\%}$	$\#_{1\%}$	$\bar{\pi}_+$	$\#_{5\%}$	$\#_{1\%}$	$\#_{5\%}$	$\#_{1\%}$	$\#_{5\%}$	$\#_{1\%}$	$\bar{\pi}_+$
Liquidity group 1 (20 most liquid assets from NASDAQ 100)																					
2 sec	20	20	20	20	14	13	0.57	20	20	20	20	14	12	0.58	20	20	20	20	14	12	0.59
3 sec	20	20	20	20	11	10	0.58	20	20	20	20	12	10	0.59	20	20	20	20	12	9	0.60
5 sec	20	20	20	20	11	9	0.59	20	20	20	20	11	10	0.60	20	20	19	19	10	9	0.62
Liquidity group 2																					
2 sec	20	20	20	20	17	17	0.53	20	20	20	19	17	17	0.54	20	20	20	20	17	16	0.54
3 sec	20	20	20	20	16	15	0.55	20	20	20	20	15	14	0.55	20	20	20	20	15	15	0.56
5 sec	20	20	20	20	15	12	0.57	20	20	20	20	14	11	0.58	20	20	20	20	11	10	0.59
Liquidity group 3																					
2 sec	20	20	18	18	20	19	0.49	20	20	18	18	18	18	0.49	20	20	18	18	18	16	0.50
3 sec	20	20	18	18	20	15	0.52	20	20	18	18	18	16	0.52	20	20	18	18	14	11	0.53
5 sec	20	20	19	19	16	12	0.55	20	20	18	18	14	10	0.56	20	20	18	18	9	7	0.57
Liquidity group 4																					
2 sec	20	20	18	18	20	19	0.44	20	20	18	17	19	19	0.45	20	20	18	15	19	19	0.45
3 sec	20	20	18	17	20	19	0.48	20	20	18	17	18	18	0.48	20	20	17	17	18	17	0.48
5 sec	20	20	19	19	17	17	0.51	20	20	19	18	16	16	0.52	20	20	20	19	17	14	0.53
Liquidity group 5 (20 least liquid assets from NASDAQ 100)																					
2 sec	20	20	19	14	20	19	0.27	20	20	14	10	20	19	0.30	20	20	11	8	20	19	0.31
3 sec	20	20	18	15	20	20	0.29	20	20	14	11	20	20	0.32	20	20	11	10	20	19	0.33
5 sec	20	20	18	14	20	19	0.32	20	20	17	11	20	19	0.35	20	20	14	7	20	18	0.36
	K = 2379							K = 1586							K = 1159						

Table 1: Significance of first order return autocorrelations calculated for local intra-daily intervals of length $T =$ 10, 15, and 20 minutes. Columns $\#_{5\%}$ and $\#_{1\%}$ report the number of stocks (out of 20 in each liquidity group) for which the null hypothesis is rejected at 5% and 1% significance level, respectively. Columns marked by (I) refer to the results of the test which is based on all intervals in the sample, while columns under (I_+) and (I_-) refer to the results based on intervals with only positive and only non-positive autocorrelations, respectively. Columns $\bar{\pi}_+$ provide average (across stocks) fractions of intervals with positive return autocorrelations in the total sample of K intra-daily intervals (K are reported in the bottom of the table). The tests are conducted for local intervals observed during 61 consecutive trading days spanning the first three month of 2014. The results are reported separately for five groups of stocks sorted by the average daily mid-quote revisions.

on local intervals of $T = 10$, 15, and 20 minutes, computed at sampling frequencies $\Delta = 2$, 3, and 5 seconds. The results are based on the first 61 trading days of 2014 and are reported separately for stocks from the five "liquidity" groups defined above.

We report on the test for absolute autocorrelation $\bar{\rho}$ in the columns marked by (I), while those marked (I_+) and (I_-) concern $\hat{\rho}_k$ for $k \in I_+$ and $k \in I_-$, respectively. We indicate the number of stocks (out of 20 in each liquidity class) for which we reject the null at 1% and 5% confidence levels, together with the relative proportions of intervals belonging to I_+ (or I_-).

First, for every stock, we reject the null hypothesis of zero average first-order autocorrelations. For a vast majority of stocks, we also reject the null for most configurations of T and Δ, when positive and negative autocorrelations are considered separately. The fraction

of intervals with positive serial correlation ranges between 27% and 62%, depending on the liquidity group, sampling frequency Δ and interval length T. Hence, the evidence clearly supports the hypothesis of significant local return serial correlation of varying signs.

Secondly, for most of the liquid stocks (group 1 and 2), we reject the null more often for regimes of positive autocorrelations than for negative autocorrelations. This pattern reverses for the low-liquidity groups 4 and 5, revealing asymmetries with respect to the intensity of the quote revisions. For longer intervals, T, rejections of the null are slightly less frequent, likely due to a loss of power caused by variations in the return dynamics which generate offsetting effects. Conversely, the proportion of rejections is largely independent of the sampling frequency, Δ.

Finally, we generally observe more intervals with positive than negative autocorrelations ($K_+ > K_-$) for the stocks from the first three liquidity groups. This suggests that a form of "momentum" is a major driver of local return dependence structure.

2.2 Persistence of Return Autocorrelations

We now explore the persistence of the return autocorrelation regimes. Predictability of the sign of the serial correlation across local windows is inconsistent with the hypothesis of zero return dependencies. Thus, rejections of this null hypothesis points to the existence of local regimes with changing return autocorrelation structures.

Define the sequence of indicator variables θ_k, $k = 1, \ldots, K$, taking on the value one in case of $\hat{\rho}_k \geq 0$, and zero otherwise,

$$\theta_k = 1_{\left\{\hat{\rho}_k \geq 0\right\}} .$$

We apply a standard Wald–Wolfowitz runs test to assess whether θ_k is purely random. A "run" is defined as a string of identical elements in the indicator sequence. Under the assumption of independent elements, the number of runs, R_θ, follows the asymptotic distribution,

$$R_\theta \sim \mathcal{N}\left(2\pi_+\pi_- K + 1,\ 4\pi_+^2\pi_-^2 K^2 - 2\pi_+\pi_- \frac{K}{K-1}\right),$$

where $\pi_+ = \frac{1}{K}\sum_k \theta_k = K_+/K$ and $\pi_- = 1 - \pi_+ = K_-/K$ denote the empirical fractions of positive and negative autocorrelation regimes, respectively. Table 2 reports results for different choices of T and Δ. The null is rejected for a clear majority of stocks even at the 99% confidence level. Moreover, the rejections fall consistently in the left tail, indicating *positive* persistence.

To explore the dependence structure in more detail, we implement a corresponding test based on conditional transition probabilities. Specifically, we quantify how likely it is to observe a local regime of positive (negative) autocorrelation, conditional on having observed positive (negative) autocorrelations in the preceding interval. Suppose that $\{\theta_k\}$ follows a two-state homogeneous

Δ	T = 10 min $\#_{2.5}$	$\#_{97.5}$	$\#_{0.5}$	$\#_{99.5}$	$\overline{\pi}_+$	T = 15 min $\#_{2.5}$	$\#_{97.5}$	$\#_{0.5}$	$\#_{99.5}$	$\overline{\pi}_+$	T = 20 min $\#_{2.5}$	$\#_{97.5}$	$\#_{0.5}$	$\#_{99.5}$	$\overline{\pi}_+$
Liquidity group 1 (20 most liquid assets from NASDAQ 100)															
2 sec	19	0	19	0	0.57	19	0	17	0	0.58	19	0	17	0	0.59
3 sec	19	0	18	0	0.58	20	0	18	0	0.59	19	0	17	0	0.60
5 sec	20	0	17	0	0.59	19	0	18	0	0.60	19	0	15	0	0.62
Liquidity group 2															
2 sec	20	0	20	0	0.53	18	0	16	0	0.54	14	0	14	0	0.54
3 sec	18	0	18	0	0.55	17	0	14	0	0.55	19	0	16	0	0.56
5 sec	19	0	17	0	0.57	19	0	18	0	0.58	17	0	17	0	0.59
Liquidity group 3															
2 sec	17	0	14	0	0.49	17	0	14	0	0.49	15	0	11	0	0.50
3 sec	15	0	15	0	0.52	18	0	16	0	0.52	18	0	15	0	0.53
5 sec	16	0	15	0	0.55	18	0	16	0	0.56	16	0	14	0	0.57
Liquidity group 4															
2 sec	20	0	14	0	0.44	16	0	14	0	0.45	13	0	12	0	0.45
3 sec	19	0	16	0	0.48	15	0	12	0	0.48	17	0	12	0	0.48
5 sec	18	0	17	0	0.51	16	0	13	0	0.52	16	0	11	0	0.53
Liquidity group 5 (20 least liquid assets from NASDAQ 100)															
2 sec	18	0	15	0	0.27	15	0	13	0	0.30	16	0	12	0	0.31
3 sec	19	0	17	0	0.29	17	0	13	0	0.32	15	0	10	0	0.33
5 sec	18	0	15	0	0.32	17	0	10	0	0.35	14	0	11	0	0.36
	K = 2379					K = 1586					K = 1159				

Table 2: The table reports the results of the Wald–Wolfowitz runs test applied for local intra-daily intervals of length $T = 10$, 15, and 20 minutes. The null hypothesis claims that the occurrence of positive vs. negative return autocorrelation regimes follows an independent random sequence. We test the null hypothesis at 5% and 1% significance levels. The number of stocks for which we reject the null hypothesis are given in columns marked by $\%_{2.5}$, $\%_{0.5}$ (rejections in the left tail) and $\%_{97.5}$, $\%_{99.5}$ (rejections in the right tail). Columns $\overline{\pi}_+$ provide average (across stocks) fractions of intervals with positive return autocorrelations in the total sample of K intra-daily intervals (K are reported in the bottom of the table). The tests are conducted for local intervals observed during 61 consecutive trading days spanning the first three month of 2014. The results are reported separately for five groups of stocks sorted by the average daily mid-quote revisions.

Markov chain of order one and define the transition probability as $p_{ij} = \text{Prob}(\theta_{k+1} = j \big| \theta_k = i)$, where i and j take values in $\{0, 1\}$. The null hypothesis stipulates that θ_k is an i.i.d. Bernoulli process, i.e., $p_{00} = p_{10} = p$ and $p_{01} = p_{11} = 1 - p$, so θ_k should be a Bernoulli distributed random variable with $\mathbb{E}[\theta_k] = p$ and $\mathbb{V}[\theta_k] = p(1 - p)$, $\forall k$. A natural estimator for p_{ij}, $\hat{p}_{ij}$, is given by the proportion of observed transitions from state i to state j relative to the total number of observations for state i. A standard central limit theorem implies, under the null and for $K \to \infty$, that $\hat{p}_{ij}$ is asymptotically distributed as,

$$\hat{p}_{ij} \sim \mathcal{N}\left(1_{\{i=1\}}(2\pi_+ - 1) + \pi_-, : \frac{\pi_+ \pi_-}{K} \right),$$

where π_+ denotes the sample estimate of p.

Table 3 reports results for different choices of T and Δ. The null hypothesis of purely random variation in the sign of the autocorrelation is rejected for the vast majority of stocks. Moreover, the failure of the null hypothesis is associated exclusively with extreme right tail realizations of the test statistic, indicating that the dependence structure is symmetric. That is, both the positive and negative autocorrelation regimes display a significant degree of persistence in their transitional dynamics across the local intervals.

	Persistence of positive regimes (H_0: $p_{11} = p$)									Persistence of negative regimes (H_0: $p_{00} = 1 - p$)								
	10 min			15 min			20 min			10 min			15 min			20 min		
Δ	$\#_{2.5}$	$\#_{97.5}$	$\overline{\pi}_+$	$\#_{2.5}$	$\#_{97.5}$	$\overline{\pi}_+$	$\#_{2.5}$	$\#_{97.5}$	$\overline{\pi}_+$	$\#_{2.5}$	$\#_{97.5}$	$\overline{\pi}_-$	$\#_{2.5}$	$\#_{97.5}$	$\overline{\pi}_-$	$\#_{2.5}$	$\#_{97.5}$	$\overline{\pi}_-$
Liquidity group 1 (20 most liquid assets from NASDAQ 100)																		
2 sec	0	20	0.57	0	19	0.58	0	16	0.59	0	19	0.43	0	19	0.42	0	19	0.41
3 sec	0	18	0.58	0	18	0.59	0	17	0.60	0	19	0.42	0	20	0.41	0	18	0.40
5 sec	0	18	0.59	0	18	0.60	0	15	0.62	0	20	0.41	0	20	0.40	0	19	0.38
Liquidity group 2																		
2 sec	0	20	0.53	0	17	0.54	0	13	0.54	0	20	0.47	0	18	0.46	0	16	0.46
3 sec	0	18	0.55	0	14	0.55	0	16	0.56	0	18	0.45	0	19	0.45	0	20	0.44
5 sec	0	18	0.57	0	18	0.58	0	15	0.59	0	20	0.43	0	19	0.42	0	17	0.41
Liquidity group 3																		
2 sec	0	16	0.49	0	17	0.49	0	14	0.50	0	15	0.51	0	17	0.51	0	14	0.50
3 sec	0	16	0.52	0	17	0.52	0	18	0.53	0	15	0.48	0	18	0.48	0	16	0.47
5 sec	0	16	0.55	0	17	0.56	0	15	0.57	0	16	0.45	0	17	0.44	0	17	0.43
Liquidity group 4																		
2 sec	0	20	0.44	0	16	0.45	0	15	0.45	0	15	0.56	0	15	0.55	0	13	0.55
3 sec	0	17	0.48	0	15	0.48	0	16	0.48	0	18	0.52	0	14	0.52	0	16	0.52
5 sec	0	17	0.51	0	15	0.52	0	14	0.53	0	19	0.49	0	17	0.48	0	16	0.47
Liquidity group 5 (20 least liquid assets from NASDAQ 100)																		
2 sec	0	19	0.27	1	18	0.30	0	19	0.31	0	12	0.73	0	9	0.70	0	9	0.69
3 sec	0	20	0.29	1	18	0.32	0	16	0.33	0	16	0.71	0	10	0.68	0	8	0.67
5 sec	0	19	0.32	0	18	0.35	0	16	0.36	0	14	0.68	0	12	0.65	0	10	0.64
	K = 2379			K = 1586			K = 1159			K = 2379			K = 1586			K = 1159		

Table 3: The table reports the number of stocks for which the null hypothesis of independent transitions between local autocorrelation regimes is rejected. On the left side, we report the results of the test for the persistence of regimes with positive return autocorrelation under H_0: $p_{11} = p$. On the right side, we test the presence of persistence of regimes with non-positive return autocorrelation under H_0: $p_{00} = 1 - p$. We reject the null hypothesis for the left tail at the 2.5% and for the right tail at the 97.5% level. The corresponding number of stocks for which we reject the null hypothesis is given in columns marked by $\#_{2.5}$ and $\#_{97.5}$. Columns $\overline{\pi}_+$ and $\overline{\pi}_-$ provide average (across stocks) fractions of intervals with positive and non-positive return autocorrelations in the total sample of K intra-daily intervals, respectively (K are reported in the bottom of the table). The tests are conducted for local intervals observed during 61 consecutive trading days spanning the first three month of 2014. The results are reported separately for five groups of stocks sorted by the average daily mid-quote revisions.

2.3 Monte Carlo design

We now investigate whether the strong evidence in favor of significant and persistent autocorrelation regimes in the individual stock returns documented above are likely to be an artifact of microstructure noise or price jumps in the high-frequency return series. We therefore perform simulations based on a standard one-factor stochastic volatility setting that has been used in a number of existing studies, for instance, Huang and Tauchen (2005), Barndorff-Nielsen et al. (2008), and Goncalves and Meddahi (2009).

Normalizing the trading day to $t \in [0, 1]$, the efficient price process $y(t)$ is assumed to evolve according to,

$$\begin{aligned} dy(t) &= \mu dt + \sigma(t) dW_y(t) + dJ(t)\,, \\ \sigma(t) &= \sigma_d(t)\sigma_s(t)\,, \\ \sigma_s(t) &= \exp\big(\beta_0 + \beta_1 \tau(t)\big)\,, \\ d\tau(t) &= \alpha\tau(t)dt + dW_\tau\,, \end{aligned}$$

where $W_y(t)$ and $W_\tau(t)$ are correlated Brownian motions with $\text{corr}(dW_y, dW_\tau) = \rho$.

We rely on the following parametrization: $\mu = 0.03$, $\beta_1 = 0.125$, $\alpha = -0.025$, and $\rho = -0.3$. Following Barndorff-Nielsen et al. (2008), we set $\beta_0 = \beta_1^2/(2\alpha)$ to normalize the integrated variance, ensuring that $\mathbb{E}[\sigma_s^2(t)] = 1$.

The efficient price volatility, $\sigma(t)$, has a multiplicative structure and consists of two components. Apart from the stochastic component $\sigma_s(t)$, we introduce a deterministic component $\sigma_d(t)$ to capture an empirically documented diurnal J-shape volatility pattern. We model it as an exponential polynomial, along the lines of Hasbrouck (1999) and Andersen et al. (2012),

$$\sigma_d(t) = C + Ae^{-at} + Be^{-b(1-t)}\,,$$

where the parameters are set to $A = 0.75$, $B = 0.25$, $C = 0.88929198$, and $a = b = 10$. The specification ensures a sensible normalization of the diurnal effect, implying $\int_0^1 \sigma_d(t)dt = 1$.

Jumps are introduced through a sum of two compound Poisson processes $J(t) = J_s(t) + J_l(t)$, independent of $W_y(t)$ and $W_\tau(t)$. The first component refers to small frequent jumps and the second component attributes to large rare jumps. In case a small jump occurs at time t, we assume that its size depends on the spot volatility and is drawn from $N\big(0, 25\sigma^2(t)\big)$. The intensity of small jump arrivals is calibrated to have on average one jump every three minutes. We assume that each large jump contributes 25% of the expected integrated variance (0.25) and occurs on average once per week.

We let the efficient price $y(t)$ be contaminated by a market microstructure noise component,

$u(t)$, so the observed price (mid-quote) reads,

$$p(t) \;=\; y(t) \;+\; u(t)\,.$$

The noise component is assumed to be independently and identically distributed as $N\big(0,\, 0.1\sigma^2(t)\big)$ at time t, whenever the mid-quote $p(t)$ is observed. This standard specification of the noise process generates an MA-structure for the observed returns, with a negative autocorrelation and a noise-to-signal ratio that varies with the sampling frequency.

Using an Euler discretization scheme, we simulate the observed price series across each trading day on a second-to-second basis, generating T=23,400 recorded mid-quotes. To align the setting with our empirical analysis, we simulate 61 trading days and break them into a single ordered sequence of intraday intervals. We then compute the first-order autocorrelations within each local interval and repeat the statistical tests applied to the individual stock price series in Section 2.1. We repeat this procedure 1,000 times and report the percentage of local intervals for which the null hypotheses are rejected.

2.4 Simulation Results

Table 4 reports the same statistics as Table 1, but based on *simulated* data. In the scenario without jumps and microstructure noise (top panel), the tests using the average absolute or signed autocorrelations tend to over-reject the null hypothesis of zero autocorrelation. The size distortion is most prevalent for relatively short local intervals and low sampling frequencies. Conversely, for longer intervals and higher sampling frequencies, the test size approaches its nominal value. This suggests that the size distortion arises from a finite-sample bias, since the tests eventually behave as prescribed by the Kokoszka and Politis (2011) asymptotic approximation.

The second panel reports on a simulation scenario *with* jumps. Jumps do not induce serial dependence in the returns, but they can seriously affect the distribution of $\hat{\rho}_k$ and, consequently, the statistical properties of the test. In fact, we observe a slight increase in the rejection rates compared to the scenario with no jumps. Again, this size distortion tends to vanish when considering combinations of T and Δ implying a higher number of observations. Finally, we note that, in the absence of microstructure noise, the simulated observed price series generates equal proportions of local intervals with positive and negative estimated return autocorrelations.

The two bottom panels provide results based on scenarios *with* microstructure noise. The presence of (independent) noise causes negative autocorrelations in high-frequency returns. Accordingly, for sampling frequencies of 2 seconds and 3 seconds, nearly all tests reject the null hypothesis across the various simulation designs. When the test is applied to the subsets of

intervals with exclusively positive and exclusively negative estimated autocorrelations, the (one-sided) tests correctly reject the null only for the negative serial correlation regimes (columns (I_-)), but never for positive serial dependence regimes (columns (I_+)). This is as expected, because a pronounced negative mean value is implied by the data generating process, while the positive values for $\hat{\rho}_k$ only sporadically in finite samples due to random shocks, and they are unlikely to generate strong outliers relative to a null hypothesis of zero return autocorrelation.

	T = 10 min							T = 15 min							T = 20 min						
	(I)		(I_+)		(I_-)			(I)		(I_+)		(I_-)			(I)		(I_+)		(I_-)		
Δ	$\%_{5\%}$	$\%_{1\%}$	$\%_{5\%}$	$\%_{1\%}$	$\%_{5\%}$	$\%_{1\%}$	$\bar{\pi}_+$	$\%_{5\%}$	$\%_{1\%}$	$\%_{5\%}$	$\%_{1\%}$	$\%_{5\%}$	$\%_{1\%}$	$\bar{\pi}_+$	$\%_{5\%}$	$\%_{1\%}$	$\%_{5\%}$	$\%_{1\%}$	$\%_{5\%}$	$\%_{1\%}$	$\bar{\pi}_+$
Simulation design without jumps and noise																					
2 sec	7.2	1.9	7.1	2.1	5.5	1.6	0.50	5.5	1.5	5.1	0.6	5.4	1.2	0.50	4.7	0.7	4.7	1.0	4.5	0.9	0.50
3 sec	10.6	2.6	8.6	1.8	8.1	2.3	0.50	7.0	1.7	6.3	1.8	6.8	1.5	0.50	6.6	1.3	5.5	1.6	5.3	1.3	0.50
5 sec	15.8	4.7	11.8	3.3	10.7	2.9	0.50	9.7	2.2	8.5	1.9	8.2	1.8	0.50	7.2	1.4	6.4	1.4	7.0	1.0	0.50
Simulation design with jumps, but without noise																					
2 sec	12.9	3.3	10.0	2.1	11.0	3.1	0.50	10.2	2.3	7.7	2.0	10.5	2.3	0.50	7.8	1.2	6.2	1.2	7.2	1.9	0.50
3 sec	17.3	5.6	12.9	2.9	14.6	4.2	0.50	12.2	2.9	8.9	2.4	10.6	3.1	0.50	10.1	2.4	7.7	1.9	7.7	1.9	0.50
5 sec	23.6	7.8	16.2	4.3	15.9	3.8	0.50	12.6	3.9	10.3	3.0	9.5	2.4	0.50	9.8	2.4	7.3	1.6	9.3	2.4	0.50
Simulation design without jumps, but with i.i.d. noise																					
2 sec	100	100	0.0	0.0	100	100	0.22	100	100	0.0	0.0	100	100	0.17	100	100	0.0	0.0	100	100	0.13
3 sec	100	100	0.0	0.0	100	100	0.33	100	100	0.0	0.0	100	100	0.29	100	100	0.0	0.0	100	100	0.27
5 sec	63.2	38.2	0.0	0.0	100	99.8	0.42	64.3	38.3	0.0	0.0	99.9	99.9	0.40	67.3	44.9	0.0	0.0	99.9	99.5	0.38
Simulation design with jumps and i.i.d. noise																					
2 sec	100	100	0.0	0.0	100	100	0.25	100	100	0.0	0.0	100	100	0.20	100	100	0.0	0.0	100	100	0.17
3 sec	100	99.8	0.0	0.0	100	100	0.35	100	100	0.0	0.0	100	100	0.32	100	100	0.0	0.0	100	100	0.30
5 sec	63.5	34.7	0.0	0.0	99.9	99.1	0.43	57.3	29.0	0.0	0.0	99.9	99.0	0.41	57.7	30.5	0.0	0.0	99.8	98.0	0.40
	K = 2379							K = 1586							K = 1159						

Table 4: Significance of first order return autocorrelations calculated for local intervals with simulated returns of length $T = 10$, 15, and 20 minutes. Columns $\%_{5\%}$ and $\%_{1\%}$ report the percentage of instances, for which the null hypothesis is rejected at 5% and 1% significance level, respectively. Columns marked by (I) refer to the results of the test which is based on all simulated intervals, while columns under (I_+) and (I_-) refer to the results based on intervals with only positive and only non-positive autocorrelations, respectively. Columns $\bar{\pi}_+$ provide average (across instances) fractions of intervals with positive return autocorrelations in the total sample of K simulated intervals (K are reported in the bottom of the table). The tests are conducted for 1000 simulated instances, each consisted of 61 simulated trading days to mirror the empirical test design from Section 2.1.

Table 5 is the counterpart to Table 2 and reports results for the Wald–Wolfowitz runs test based on simulated data. Since none of the simulation scenarios introduces any dependence structure in the return autocorrelations across consecutive intra-daily intervals, we expect the test size to be satisfactory. In fact – even in the presence of noise and jumps – the rejection rates are well aligned with their nominal values for all simulation designs, interval lengths and observation frequencies. It is interesting to note, however, that rejections from the left tail occur slightly more often than from the right tail. This suggests that, under the null, the test statistic is

mildly skewed in small samples.

	T = 10 min					T = 15 min					T = 20 min				
Δ	$\%_{2.5}$	$\%_{97.5}$	$\%_{0.5}$	$\%_{99.5}$	$\bar{\pi}_+$	$\%_{2.5}$	$\%_{97.5}$	$\%_{0.5}$	$\%_{99.5}$	$\bar{\pi}_+$	$\%_{2.5}$	$\%_{97.5}$	$\%_{0.5}$	$\%_{99.5}$	$\bar{\pi}_+$
	Simulation design without jumps and noise														
2 sec	3.2	1.9	0.9	0.3	0.50	1.1	2.5	0.5	0.3	0.50	2.8	1.5	0.6	0.2	0.50
3 sec	2.7	3.0	0.4	0.3	0.50	2.3	2.2	0.4	0.3	0.50	2.7	2.3	0.1	0.4	0.50
5 sec	3.5	2.6	0.6	0.3	0.50	2.7	2.4	0.5	0.6	0.50	1.8	3.9	0.2	1.2	0.50
	Simulation design with jumps, but without noise														
2 sec	2.6	2.6	0.3	0.7	0.50	2.5	1.2	0.3	0.1	0.50	3.3	2.5	0.6	0.2	0.50
3 sec	3.0	1.2	0.2	0.2	0.50	2.6	1.7	0.7	0.2	0.50	2.1	3.1	0.4	0.8	0.50
5 sec	2.5	2.2	0.6	0.1	0.50	2.5	2.5	0.5	0.3	0.50	2.5	1.7	0.3	0.1	0.50
	Simulation design without jumps, but with i.i.d. noise														
2 sec	2.2	2.4	0.1	0.4	0.22	3.4	1.2	1.2	0.3	0.17	4.3	1.0	1.0	0.2	0.13
3 sec	2.8	2.7	0.4	0.6	0.33	2.0	2.4	0.2	0.3	0.29	3.4	1.8	0.4	0.7	0.27
5 sec	3.7	2.2	0.6	0.7	0.42	2.6	3.3	0.4	0.5	0.40	1.7	1.3	0.3	0.1	0.38
	Simulation design with jumps and i.i.d. noise														
2 sec	3.0	2.2	0.3	0.5	0.25	3.1	1.8	0.5	0.4	0.20	3.8	1.8	0.9	0.3	0.17
3 sec	1.8	1.6	0.3	0.8	0.35	2.6	1.6	0.5	0.6	0.32	2.5	1.8	0.8	0.1	0.30
5 sec	2.6	2.6	0.8	0.5	0.43	2.6	2.6	0.7	0.3	0.41	3.1	1.9	0.7	0.1	0.40
	K = 2379					K = 1586					K = 1159				

Table 5: The table reports the results of the Wald–Wolfowitz runs test applied for local intervals with simulated returns of length $T = 10$, 15, and 20 minutes. The null hypothesis claims that the occurrence of positive vs. negative return autocorrelation regimes follows an independent random sequence. We test the null hypothesis at 5% and 1% significance levels. The fractions of instances for which we reject the null hypothesis are given in columns marked by $\%_{2.5}$, $\%_{0.5}$ (rejections in the left tail) and $\%_{97.5}$, $\%_{99.5}$ (rejections in the right tail). Columns $\bar{\pi}_+$ provide average (across instances) fractions of intervals with positive return autocorrelations in the total sample of K simulated intervals (K are reported in the bottom of the table). The tests are conducted for 1000 simulated instances, each consisted of 61 simulated trading days to mirror the empirical test design from Section 2.2.

Finally, Table 6 mimics Table 3, providing evidence on the persistence in the sign of the autocorrelation regimes. The presence of noise has a small, yet notable, impact on the rejection patterns. We observe moderate over-rejections in the transition rate from intervals with positive autocorrelation (left panel) while, conversely, the test tends to under-reject for transitions from intervals with negative autocorrelations (right panel). Since the noise process induces *negative* return autocorrelation, the simulations generate only a few instances of (estimated) positive autocorrelations and, as a result, the estimates of $\hat{p}_{00}$ and $\hat{p}_{11}$ are imprecise, rendering size distortions more likely in finite samples.[5] We note, however, that the rejections occur equally often in the left and right tail. Thus, the test (correctly) does not indicate the presence of persistence or anti-persistence in the autocorrelation regimes.

In summary, we find that our tests possess reasonable statistical properties. Even if there are

[5]This conjecture is corroborated by the fact that size distortions grow more pronounced for higher frequencies and longer intraday intervals. In these settings, we obtain even more negative estimates for $\hat{\rho}_k$, further reducing the number of transitions between different regimes.

size distortions, they differ by an order of magnitude from the test outcomes in Sections 2.1 and 2.2. Hence, the empirical evidence on the time variability, significance and persistence of local autocorrelation regimes documented above is likely not spurious, but rather driven by structural mechanisms of the price formation process in the market.

	Persistence of positive regimes (H_0: $p_{11} = p$)									Persistence of negative regimes (H_0: $p_{00} = 1 - p$)								
	10 min			15 min			20 min			10 min			15 min			20 min		
Δ	$\%_{2.5}$	$\%_{97.5}$	$\overline{\pi}_+$	$\%_{2.5}$	$\%_{97.5}$	$\overline{\pi}_+$	$\%_{2.5}$	$\%_{97.5}$	$\overline{\pi}_+$	$\%_{2.5}$	$\%_{97.5}$	$\overline{\pi}_-$	$\%_{2.5}$	$\%_{97.5}$	$\overline{\pi}_-$	$\%_{2.5}$	$\%_{97.5}$	$\overline{\pi}_-$
Simulation design without jumps and noise																		
2 sec	2.2	3.1	0.50	2.7	1.2	0.50	1.8	2.6	0.50	2.3	2.9	0.50	3.0	1.0	0.50	1.9	2.1	0.50
3 sec	3.1	2.4	0.50	2.5	2.3	0.50	3.0	2.2	0.50	3.2	2.3	0.50	2.7	2.0	0.50	3.4	2.3	0.50
5 sec	3.0	3.4	0.50	2.6	2.6	0.50	4.4	1.4	0.50	3.6	3.1	0.50	3.3	2.4	0.50	4.6	1.2	0.50
Simulation design with jumps, but without noise																		
2 sec	3.1	2.0	0.50	1.8	2.6	0.50	3.0	2.9	0.50	3.2	2.1	0.50	2.1	1.8	0.50	3.6	2.4	0.50
3 sec	1.5	2.3	0.50	2.2	2.7	0.50	3.3	1.9	0.50	1.8	1.9	0.50	2.2	2.2	0.50	3.7	1.4	0.50
5 sec	2.7	2.5	0.50	2.9	2.3	0.50	1.9	2.6	0.50	3.2	2.1	0.50	3.0	2.2	0.50	2.7	1.8	0.50
Simulation design without jumps, but with i.i.d. noise																		
2 sec	14.7	14.9	0.22	17.2	18.9	0.17	20.3	23.8	0.13	0.0	0.0	0.78	0.0	0.0	0.83	0.0	0.0	0.87
3 sec	8.7	9.1	0.33	9.9	9.2	0.29	13.4	11.5	0.27	0.5	0.1	0.67	0.0	0.0	0.71	0.2	0.1	0.73
5 sec	4.7	5.7	0.42	6.9	5.2	0.40	5.2	5.3	0.38	1.4	0.8	0.58	1.6	0.4	0.60	0.3	0.1	0.62
Simulation design with jumps and i.i.d. noise																		
2 sec	13.9	13.6	0.25	14.7	14.2	0.20	16.8	17.3	0.17	0.1	0.0	0.75	0.0	0.0	0.80	0.0	0.1	0.83
3 sec	7.7	5.1	0.35	9.4	8.8	0.32	11.3	9.9	0.30	0.9	0.1	0.65	0.3	0.0	0.68	0.1	0.3	0.70
5 sec	5.2	3.8	0.43	5.7	4.7	0.41	4.4	4.9	0.40	1.4	1.0	0.57	1.2	0.9	0.59	0.8	0.8	0.60
	K = 2379			K = 1586			K = 1159			K = 2379			K = 1586			K = 1159		

Table 6: The table reports the number of instances for which the null hypothesis of independent transitions between autocorrelation regimes in simulated local intervals is rejected. On the left side, we test the presence of persistence of regimes with non-positive return autocorrelation under H_0: $p_{11} = p$. On the right side, we test the presence of persistence of regimes with non-positive return autocorrelation under H_0: $p_{00} = 1 - p$. We reject the null hypothesis for the left tail at the 2.5% and for the right tail at the 97.5% level. The corresponding fractions of instances for which we reject the null hypothesis are given in columns marked by $\%_{2.5}$ and $\%_{97.5}$. Columns $\overline{\pi}_+$ and $\overline{\pi}_-$ provide average (across instances) fractions of intervals with positive and non-positive return autocorrelations in the total sample of K simulated intervals (K are reported in the bottom of the table). The tests are conducted for 1000 simulated instances, each consisted of 61 simulated trading days to mirror the empirical test design from Section 2.2.

3 Models for the High-Frequency Asset Price Dynamics

This section introduces a parsimonious parametric model that can accommodate the salient features of the short-term price dynamics documented in Section 2. The objective is to assist in the identification and interpretation of the drivers behind the time-variation in the shape of the volatility signature plots and the return serial correlation patterns. The aspiration is to provide a starting point for building more realistic models of real-time price discovery and market fluctuations in which we can gauge the relative impact of learning, information heterogeneity

and elements of market microstructure frictions.

3.1 The Discrete-Time Local Regime Model

We consider a model in discrete time, $i \in \{0, 1, 2, \ldots, T\}$, defined over an equidistant time grid with (minimal) interval length $\Delta = 1$. At each discrete point in time, we observe the logarithm of the quote midpoint for a single financial asset, p_i. Hence, for the time period $[0, T]$, we have a total of $T + 1$ separate log-price observations, and T consecutive continuously compounded returns, $r_i = p_i - p_{i-1}$, $i = 1, \ldots, T$. For the interval length $\Delta = 1$, the number of return observations equals $n = T/\Delta = T$, while for coarser sampling, governed by the integer $\Delta > 1$, the number of non-overlapping returns is $n = \lfloor T/\Delta \rfloor$. While the notation is general, we focus on the salient features of the intraday return dynamics, so we typically think of Δ as representing a short trading interval, on the order of multiple seconds, but not fractions of a second. Even for mid-quote returns, the tick size is likely a major factor in the computation of the realized return variation at very high sampling frequencies.

An important consideration for the design of our model is the identification of a channel for price discovery, distinct from standard specifications of market microstructure noise, allowing for significant return autocorrelation over non-trivial time intervals. Since we cannot disentangle such effects from microstructure frictions by nonparametric techniques, we impose identifying structure through specific parametric assumptions that may be subject to empirical scrutiny.

The "efficient" (full information and frictionless) log price at time i, p_i^*, follows a random walk, consistent with a no-arbitrage representation for the log-return over short intraday intervals,

$$p_i^* \;=\; p_{i-1}^* + \varepsilon_i^* \qquad \text{and} \qquad \varepsilon_i^* \;=\; r_i^* \;=\; p_i^* \,-\, p_{i-1}^*\,, \tag{2}$$

where $r_i^* = \varepsilon_i^*$ is the "efficient" return with $\mathbb{E}[\varepsilon_i^*] = 0$ and $\mathbb{V}[\varepsilon_i^*] = \sigma_*^2$.

The evidence from Section 2 suggests that the dynamics of the mid-quote price deviates from that of the efficient full-information price in significant ways. Hence, we propose a simple representation that accommodates both price endogeneity and correlation in the pricing errors via readily identifiable components. Specifically, we assume that the return dynamics, locally, evolves according to the scheme,

$$r_i \;=\; p_i \,-\, p_{i-1} \;=\; -\,\alpha\,(p_{i-1} \,-\, p_{i-1}^*) \,+\, (\gamma\,\varepsilon_i^* \,+\, \varepsilon_i)\,, \quad 0 < \alpha < 2, \tag{3}$$

where ε_i is an i.i.d. return component, uncorrelated with ε_i^*, $\mathbb{E}[\varepsilon_i\,] \,=\, \mathbb{E}[\varepsilon_i\,\varepsilon_i^*] \,=\, 0$ and $\mathbb{V}[\varepsilon_i] = \sigma_\varepsilon^2$. Finally, we denote the noise-to-signal ratio by $\lambda = \sigma_\varepsilon^2/\sigma_*^2$. This ratio is known to be critical in classical settings. For example, the optimal sampling frequency in the "classical"

martingale plus noise models hinges on the noise-to-signal ratio, see, e.g., Ait-Sahalia et al. (2005a) and Bandi and Russell (2006).

For $\alpha = \gamma = 1$, we obtain the special case of the "classical" martingale plus (i.i.d.) noise model $p_i = p_i^* + \varepsilon_i^*$. If $\gamma \neq 1$ and $\alpha \neq 1$, however, the terms involving γ and α allow for endogeneity and an error correction mechanism in the price dynamics. These components can induce inversions in the volatility signature plot and persistent return autocorrelation patterns, respectively. We illustrate each feature in turn.

Uncorrelated Endogenous Pricing Errors

For $\alpha = 1$, we eliminate persistent return serial dependence. Specifically,

$$p_i \;=\; p_i^* \;+\; [\,(\gamma \;-\; 1)\,\varepsilon_i^* \;+\; \varepsilon_i\,]\,, \tag{4}$$

so the system does feature i.i.d. noise and, in general, there is endogeneity, as the error term is correlated with the fundamental price innovation ε_i^*, whenever $\gamma \neq 1$.

Allowing γ to differ across local intervals is one way to mimic intertemporal heterogeneity in the information environment. Traders with incomplete information constantly seek to infer the fundamental value from the concurrent market dynamics, including signals derived from observable price innovations, evolving order book imbalances, consummated trades, and incoming news items. At a given point in time, rational agents with incomplete information tend to draw similar conclusions from available public signals. Hence, cumulative random shocks in either direction generate correlated, albeit perhaps short-lived, pricing errors. Thus, even if investors price assets correctly on average, they will, for short intervals, induce $\gamma > 1$ in some scenarios and $\gamma < 1$ for others. Such temporary periods of over- or under-reaction to underlying latent news affects the properties of the noise and, consequently, the price dynamics. To formalize the discussion, we explicitly characterize the return variation and first-order return dependence in the system.

Direct computations show that the return variance equals

$$\mathbb{V}[r_i] \;=\; \mathbb{V}[p_i - p_{i-1}] \;=\; \sigma_*^2 \;+\; 2\;[\,\gamma(\gamma-1) \;+\; \lambda\,]\,\sigma_*^2\,, \tag{5}$$

while the first-order return auto-covariance takes the form

$$\mathbb{C}ov[r_i, r_{i-1}] \;=\; [\,\gamma(1-\gamma) \;-\; \lambda\,]\,\sigma_*^2\,, \tag{6}$$

and all higher order return autocorrelation coefficients ($h > 1$) are zero.

Hence, the return variation is a function of the quadratic term $\gamma(\gamma - 1)$. If $\gamma = 1$, we have

the usual noise inflation of the return variance. For $\gamma < 1$, the concurrent return innovation does not fully incorporate the efficient price innovation. This induces a smoothing of the price path which counteracts the impact of the idiosyncratic noise component. As a result, the return variation declines, and if this effect dominates the idiosyncratic noise component, the observed return variation may fall below the level associated with the fundamental return, i.e., $\mathbb{V}[r_i] < \sigma_*^2$. This smoothing effect becomes stronger, as γ declines, attaining a minimum for $\gamma = 1/2$. As γ drops below $1/2$, too little of the initial efficient price innovation is incorporated so that, combined with the future price innovations, the system overall generates larger deviations from the efficient price path. That is, $\mathbb{V}[r_i]$ is decreasing in γ for $\gamma \in (0, 1/2)$ and increasing for $\gamma \in (1/2, 1)$.

The impact of γ on the return autocorrelation is reserved, as it depends on $\gamma(1-\gamma)$. Accordingly, the serial dependence is maximized for $\gamma = 1/2$, and it is positive if $\gamma(\gamma-1) > \lambda$, i.e., in scenarios where the noise-to-signal ratio is not too big. Hence, in these cases, endogeneity – correlation between the fundamental returns and the pricing errors – causes momentum which overwhelms the impact of the idiosyncratic noise and can generate positive serial dependence in the return series. Of course, the opposite is also possible. For $\gamma < 0$ or $\gamma > 1$, the price response is excessive and exacerbates the excess return variation and negative return correlation induced by the idiosyncratic noise component.

Note that cases for values of γ with the same absolute distance to $1/2$, i.e. for same values of $|\gamma - 1/2|$ for $\gamma \in [0, 1]$ imply identical return variances and autocovariances. Hence, positive and negative deviations from the variance-minimizing level $\gamma = 1/2$ imply an extent of smoothing which ultimately counteracts the idiosyncratic noise component in a way implying identical second moments. In this context, the case $\gamma = 1$ is equivalent to $\gamma = 0$. Hence, as long as $\alpha = 1$, the case of $\gamma = 0$ coincides with the classical martingale plus noise model ruling out any endogeneity. The coincidence of these cases results into a statistical identification problem which will be discussed in more detail below.

Error Correction Dynamics

Focusing instead on the error correction mechanism, the scenario $\gamma = 1$ and $\alpha \neq 1$ yields

$$p_i \;=\; p_i^* \;-\; \alpha\,(p_{i-1} \;-\; p_{i-1}^*) \;+\; \varepsilon_i\,, \tag{7}$$

Now, the i.i.d. error term is independent of the fundamental return innovation, but for $0 < \alpha < 1$, the observed price responds sluggishly to past pricing errors, inducing longer-run return autocorrelation. The latter is, empirically, a realistic scenario, as documented in Section 2.

This error correction mechanism is motivated by the intuition that risk-averse agents with

incomplete information form (unconditionally) unbiased, yet not error-free, expectations about the efficient price. In this setting, the agents are aware that temporary mispricing is likely. This creates opportunities for (risky) speculative trading among partially informed agents, generating the pull back towards efficiency, i.e., when prices appear overvalued, informed agents short the asset. As such, the model captures temporal feedback effects through a minimalistic reduced-form approach. Earlier contributions, including Kyle (1985) and Vives (1995), discuss various settings and economic factors, that render prolonged price reaction patterns consistent with sequential learning models.

More formally, it is straightforward in this case to establish that

$$\mathbb{V}[r_i] \;=\; \sigma_*^2 \;+\; \frac{2\,\lambda}{2\;-\;\alpha}\,\sigma_*^2\,, \tag{8}$$

and the h^{th}-order return auto-covariance is

$$\mathbb{C}ov[r_i, r_{i-h}] \;=\; -\;(1\;-\;\alpha)^{h-1}\,\frac{\alpha\,\lambda}{2\;-\;\alpha}\,\sigma_*^2\,. \tag{9}$$

Absent endogeneity, the error term invariably inflates the return variation in equation (8) and the return autocorrelations in equation (9) are all negative. When α is smaller, the information transmission is slower, i.e., a given pricing error induces a more protracted correction implying more return smoothing and yielding a reduction of the price variation. Likewise, if $\alpha > 1$, price reactions imply an overshooting increasing the return variation.

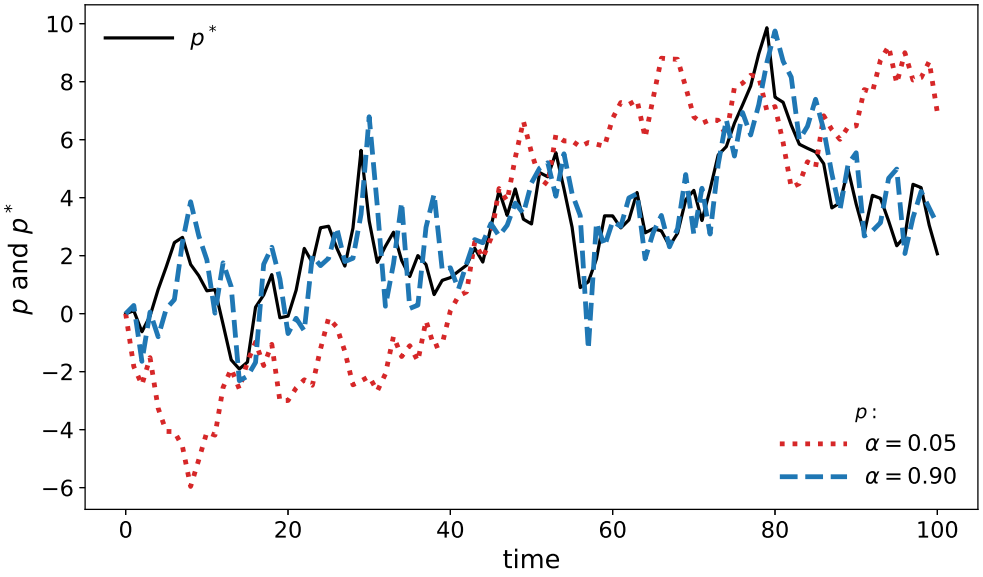

Figure 3: Simulated trajectories for p_i^* and p_i from the model (2) - (3) given two different values of α. The remaining parameters take the values $\sigma_*^2 = 1$, $\gamma = 0$ and $\lambda = 0.1$.

The Random-Walk-plus-Noise Representation

To facilitate a more direct comparison to the extant literature, it is useful to couch our model (2)-(3) in a standard random-walk-plus-noise type format. Towards this end, we introduce explicit notation for the mispricing component, its innovation, and its variance,

$$\mu_i \;=\; p_i \,-\, p_i^*, \qquad \varepsilon_i^{\mu} \,=(\gamma-1)\,\varepsilon_i^* \,+\, \varepsilon_i \quad \text{and} \quad \mathbb{V}[\varepsilon_i^{\mu}] \;=\; \sigma_{\mu}^2 \;=\; (\gamma-1)^2\sigma_*^2 \,+\, \sigma_{\epsilon}^2\,. \tag{10}$$

The system (2)–(3) now implies,

$$p_i \;=\; p_i^* \,+\, \mu_i, \tag{11}$$

$$\mu_i \;=\; (1-\alpha)\,\mu_{i-1} \,+\, \varepsilon_i^{\mu}\,. \tag{12}$$

Equation (11) highlights the role of μ_i as a "standard" microstructure noise component, creating a wedge between p_i and p_i^*. However, the temporal feedback and the correlation between the innovations in equations (2) and (3) imply that the pricing error is affected by shocks to both the efficient price and noise process. Econometrically, $\alpha\,\mu_{i-1}$ serves as an error correction term, pushing the mid-quote back towards the "equilibrium" level, $p_i = p_i^*$.

The role of ε_i^* in the error ε_i^{μ} mirrors the discussion abvoe in the case of uncorrelated endogenous errors. If $\gamma = 1$, the efficient price innovation is fully incorporated in p_i and therefore does not contribute to ε_i^{μ}. In contrast, if $\gamma = 0$, the efficient price innovation is not taken into account. Hence, agents do not respond instantaneously to new information, which implies a lagged price response. Consequently, the noise is endogenous with ε_i^* fully embedded within ε_i^{μ}. This case is discussed in depth in Section 3.4.5.

Figure 3 displays a simulated efficient price path along with a few corresponding price series generated by the model, given specific parametric configurations. The choice of $\gamma = 0$ reflects a short (one-period) delay associated with information retrieval and processing, while the different values for α control the degree of adherence to the efficient price.[6] The feedback generated by the error correction mechanism induces short-term serial correlation, as it is evident from the $\alpha = 0.05$ and $\alpha = 0.70$ scenarios. Both of these series constitute smoothed and lagged variants of the more jagged efficient price path, but it is particularly striking for the low $\alpha = 0.05$ value. In the extreme case, $\alpha = 0$, the efficient and observed prices are entirely decoupled. In contrast, for $\alpha = 1$, all current mispricing is eliminated in the subsequent period.

The temporal feedback in equations (11)-(12) takes the form of a mean-zero AR(1) process.

[6] As we establish later, as long as $\alpha \neq 1$, the case $\gamma = 0$ is empirically relevant, and a canonical case of particular interest.

It is stationary for $0 < \alpha < 2$ with unconditional variance,

$$\mathbb{V}[\mu_i] \;=\; \frac{\sigma_\mu^2}{\alpha\,(2-\alpha)} \;=\; \frac{(\gamma-1)^2 \,+\, \lambda}{\alpha\,(2-\alpha)}\,\sigma_*^2 . \tag{13}$$

Expression (13) may be interpreted as a measure of mispricing. The mid-quote price equals the efficient price in expectation, i.e., $\mathbb{E}[\mu_i] = 0$ and $\mathbb{E}[p_i] = \mathbb{E}[p_i^*]$, so the average squared pricing error is determined by the variance, $\mathbb{V}[\mu_i]$. It is evident that $\mathbb{V}[\mu_i]$ is minimized for $\alpha = 1$ implying, rather intuitively, that rapid error correction enhances price efficiency.

The second factor governing the extent of inefficiency is the numerator in equation (13), namely the size of the pricing error innovation. This component is small if, all else equal, there is no (local) over- or under-reaction to underlying latent news ($\gamma = 1$), and the idiosyncratic error is minimal (low λ).

In summary, the price dynamics reflects the covariance structure of the innovations to the mispricing component and to the fundamental price, and how these shocks get amplified through the temporal adjustment process. This covariance structure takes the form

$$\Sigma \;=\; \begin{bmatrix} \mathbb{E}[(\varepsilon_i^\mu)^2] & \mathbb{E}[\varepsilon_i^\mu \varepsilon_i^*] \\ \mathbb{E}[\varepsilon_i^\mu \varepsilon_i^*] & \mathbb{E}[(\varepsilon_i^*)^2] \end{bmatrix} \;=\; \begin{bmatrix} (\gamma-1)^2+\lambda & \gamma-1 \\ \gamma-1 & 1 \end{bmatrix} \sigma_*^2 \,. \tag{14}$$

Equation (14) demonstrates that the off-diagonal entries are non-zero, if and only if $\gamma \neq 1$. The ramifications of the endogeneity of the pricing error is further discussed below in our review of familiar models from the extant literature.

We reiterate that our model, deliberately, is stylized and focuses on specific channels behind salient features of the return dynamics. It is intended as an approximation for short intraday periods, like 10-30 minutes, for which the factors governing the return dynamics remain stable. Over longer horizons, the size of the volatility innovation is clearly not constant. Allowing model parameters to shift across adjacent short intervals, indicating a transition from one informational regime to another, is in line with the findings of significant time-varying return autocorrelations in Section 2. We may move from a period of under-reaction to one of over-reaction due to the shifting nature of the information flow and the broader market environment.[7]

3.2 Second Order Return Moments

This section characterizes the conditions for excess volatility and return serial correlation implied by the model. These features are interrelated and closely associated with the impact of

[7]The spirit of this approach mimics the assumptions often invoked for developing inference techniques with high-frequency data, see, e.g., Mykland and Zhang (2009) and Bibinger et al. (2014), who explicitly rely on local windows in which the quantity of interest, in their case return volatility, may be assumed to remain fixed.

idiosyncratic noise versus smoothing in the return variation process. The latter also affects our ability to identify the underlying sources behind the price dynamics from the observed return moments without imposing auxiliary distributional assumptions.

3.2.1 Excess Return Variation

We recall that the continuously compounded mid-quote returns in equation (3) may be stated as $r_i = -\alpha\mu_{i-1} + (\varepsilon_i^* + \varepsilon_i^\mu)$. The following lemma provides an explicit expression for the return variation in this general case.

Lemma 1 (The Return Variance)**.** *Assume* $\sigma_\varepsilon^2 > 0,\ \ 0 < \alpha \leq 1,\ \ 0 \leq \gamma < 2.$ *Then,*

$$\mathbb{V}[r_i] \ = \ \alpha^2\,\mathbb{V}[\mu_i] \ + \ (\gamma^2 + \lambda) \tag{15}$$

$$= \ \sigma_*^2 \ + \ \frac{2}{2\,-\,\alpha}\,F_\alpha(\gamma,\lambda)\ \sigma_*^2\ , \tag{16}$$

where, for later convenience, for given α, we define $F_\alpha(\gamma,\lambda)$ as a function of γ and λ,

$$F_\alpha(\gamma,\lambda) \ = \ \gamma^2 \ + \ (\,1\,-\,\gamma\,)\ \alpha \ + \ \lambda \ - \ 1.$$

Proof. Follows straightforwardly from equation (13). □

From equation (15), one readily observes the intuitive relationship that the return variance is increasing in σ_*^2 and λ. Likewise, reflecting the effect of a "smoothing" of the price path for low values of α, the return variance is increasing in α (for $\gamma < 1$). Finally, we again observe that γ affects volatility in a non-monotonic way. The return variance is minimized for $\gamma = \alpha/2$, it is increasing in γ for $\gamma > \alpha/2$, and decreasing for $\gamma < \alpha/2$. This generalizes the findings for the uncorrelated endogenous noise setting above in the case of $\alpha = 1$.

Hence, in the general setting of equation (15), featuring both error correction and endogeneity, the overall impact of γ on the return variation is ambiguous and depends on α. We therefore directly compare the variance expression (15) to the fundamental variance, σ_*^2, as expressed in the following corollary.

Corollary 1. *Assume* $\sigma_\varepsilon^2 > 0,$ *and* $0 < \alpha \leq 1.$

$$\mathbb{V}[r_i] \ \leq \ \mathbb{V}[r_i^*] \qquad \textit{if} \qquad F_\alpha(\gamma,\lambda) \ \leq \ 0\,, \tag{17}$$

$$\mathbb{V}[r_i] \ > \ \mathbb{V}[r_i^*] \qquad \textit{otherwise}. \tag{18}$$

Hence, the observed return variation undercuts the efficient return variation if price updating and the extent of noise are slow (i.e., α and λ are small) and γ does not deviate too much from $\alpha/2$,

implying variance-minimizing smoothing. The condition $F_\alpha(\gamma, \lambda) < 0$ can be alternatively stated as $\gamma^2 + \lambda < 1 - \alpha(1 - \gamma)$. Hence, the two regimes are governed by the relationship between the extent of *effective return smoothing* $1 - \alpha(1 - \gamma)$ and the size of the innovations to the price variation $\gamma^2 + \lambda$.

Consequently, we have two separate excess volatility regimes, governed by the instantaneous response to new fundamental information, γ, the noise-to-signal ratio, λ, and the smoothing effect of partial price adjustment, captured by low values of α. In the classical model, $\gamma = 1$, the presence of noise, $\lambda > 0$, trivially leads to excess volatility. If noise is absent, however, the partial price adjustment can drive the return variation below the level of the fundamental innovation process.

3.3 Return Auto-Covariances

The determinants of whether returns display excess volatility are also critical for the sign of the return auto-covariances. This is a consequence of the following lemma and corollary.

Lemma 2 (Return Auto-Covariances)**.** *Assume* $\sigma_\varepsilon^2 > 0$, $0 < \alpha \leq 1$, , $h \geq 1$. *Then,*

$$\mathbb{C}ov[r_i, r_{i-h}] = \psi(h-1)\ \sigma_*^2\ \left[1 - \alpha\,(1-\gamma) - (\gamma^2 + \lambda)\right] \tag{19}$$

$$= -\ F_\alpha(\gamma, \lambda)\ \psi(h-1)\ \sigma_*^2\,, \tag{20}$$

with $\psi(h-1) = \frac{\alpha}{2-\alpha}\,(1-\alpha)^{h-1}$, *and* $\psi(0) = 1$, *if* $\alpha = 1$.

Proof. See Appendix. □

Hence, the relation between $1 - \alpha\,(1 - \gamma)$ and $\gamma^2 + \lambda$, also governs the sign of the auto-covariance for the observed returns. We summarize this result in the form of a corollary.

Corollary 2. *Assume* $\sigma_\varepsilon^2 > 0$, $0 < \alpha \leq 1$, *and* $h \geq 1$.

$$\mathbb{C}ov[r_i, r_{i-h}] \geq 0 \qquad \textit{if} \qquad F_\alpha(\gamma, \lambda) \leq 0\,, \tag{21}$$

$$\mathbb{C}ov[r_i, r_{i-h}] \leq 0 \qquad \textit{otherwise}. \tag{22}$$

Furthermore, we note that, for $0 < \alpha < 1$, the sign of the auto-correlations are identical for all $h \geq 1$, indicating that both the "contrarian" and "momentum" regimes imply a certain persistence in the return dependencies. To the extent these regimes provide a good approximation to the local return dynamics over short, yet non-trivial, trading intervals, it may represent us with an opportunity to capture such episodes empirically. In contrast, for $\alpha > 1$, the first-order return serial correlation is negative, and the signs alternate for $h > 1$. While such behavior may

be plausible over longer horizons, this type of high-frequency shifts in the return dependencies for mid-quotes across short intraday trading intervals is counterfactual.

For models without endogeneity, $\gamma = 1$, the condition for positive return autocovariance, $\lambda < 0$, is infeasible – such models never generate momentum regimes with positive correlation in returns. We discuss this point in more detail below in Section 3.4.1. In contrast, for $\gamma < 1$, positive return autocorrelation is attainable. In contrast to the pure price endogeneity setting discussed in Section 3.1, this is even possible in case of $\gamma = 0$ (if $\lambda < 1 - \alpha$).

In fact, we find mid-quote returns, on average, to display only limited serial dependence. For low values of γ, this is possible in our general setting, only if $1 - \alpha \approx \lambda$. This scenario is plausibly, if the noise component is small, while prices impound information swiftly, so that $1 - \alpha$ is small as well. More generally, zero return correlation is consistent with some noise, as long as the impact is mitigated by a certain sluggishness in market prices. Nonetheless, a relatively large noise component, $\lambda \geq 1$, always induces negative serial correlation and excess return volatility due to the strong reversals required to correct the idiosyncratic mispricing.

A natural application of these results is to realized volatility estimation. They imply that the traditional realized variance measure, obtained from cumulative high-frequency squared returns, is biased and inconsistent unless the market environment generates a return dynamic that satisfies the condition for zero excess volatility or, equivalently, for no (significant) return serial correlation, namely, $\gamma^2 + \lambda \approx 1 - \alpha\,(1-\gamma)$ or $F_\alpha(\gamma, \lambda) \approx 0$.

Finally, note that for given α and λ, scenarios for which $|\gamma - \alpha/2|$ are identical, e.g., the cases $\gamma = 0$ and $\gamma = \alpha$, yield the same return variances and autocovariances. We will discuss these aspects in more detail in the context of statistical identification in Chapter 4.1.

3.4 Standard Microstructure Model Specifications

Our general model seeks to capture the basic structural mechanisms that are present in market microstructure models within a uniform setting. This provides a starting point for pinpointing strengths as well as limitations of existing models. Towards this end, it is useful that our model nests important special cases arising from distinct parametric restrictions. Since we subsequently estimate a number of these representations over short intraday intervals, we obtain guidance in terms of which stylized models best approximate critical salient features of the high-frequency return process. A second objective is to highlight the identification issues that arise in microstructure models with a latent endogenous noise component. This becomes evident in our review of the empirical implications of some existing models below.

3.4.1 The Classical Model with Idiosyncratic Noise Autocorrelation

A critical feature of our model (2)–(3) is the allowance for noise endogeneity by having $\gamma \neq 1$. In fact, if $\gamma = 1$ then, independently of the noise dynamics, we obtain,

$$\Sigma \;=\; \begin{bmatrix} \mathbb{E}[(\varepsilon_i^{\mu})^2] & \mathbb{E}[\varepsilon_i^{\mu}\varepsilon_i^*] \\ \mathbb{E}[\varepsilon_i^{\mu}\varepsilon_i^*] & \mathbb{E}[(\varepsilon_i^*)^2] \end{bmatrix} \;=\; \begin{bmatrix} \lambda & 0 \\ 0 & 1 \end{bmatrix} \sigma_*^2 \,.$$

The absence of correlation among the innovations has important implications for the return variation. Given this assumption, we have, for *any* stationary noise dynamics,

$$r_i \;=\; p_i \,-\, p_{i-1} \;=\; \varepsilon_i^* \,+\, \Delta\mu_i \,, \tag{23}$$

where $\Delta\mu_i = \mu_i - \mu_{i-1}$.

Since in this case $\mathbb{E}[\,\varepsilon_i^*\,\Delta\mu_j\,] \,=\, 0$ for all integers i, j, it follows that the actual return variation must exceed the efficient return variation,

$$\mathbb{E}[\,r_i^2\,] \;=\; \sigma_*^2 \,+\, \mathbb{E}[\,(\Delta\mu_i)^2\,] \;\geq\; \sigma_*^2$$

implying that the return variance and autocovariance are of the form in equations (8) and (9). Hence, uniformly, the volatility signature plot should be above the efficient return variance. This is at odds with the empirical evidence in Section 2, where the signature plot often drops significantly as we sample at increasingly high frequencies.

We reiterate that this feature arises, even with dependence in the return series induced by autocorrelation in the noise. Our model specification (11) and (12) now implies,

$$p_i \;=\; p_i^* \,+\, \mu_i \,, \tag{24}$$

$$\mu_i \;=\; (1-\alpha)\,\mu_{i-1} \,+\, \varepsilon_i \,, \tag{25}$$

with $\mathbb{E}[\varepsilon_i\,\varepsilon_i^*] \,=\, 0$ and $\mathbb{E}[\,\varepsilon_i\;\mu_j\,] \,=\, 0$ for all integers i, j. Hence, in our parametric setting, the mispricing component, μ_i, follows an AR(1) process, and is serially correlated for $\alpha \neq 1$, inducing return dependence. This feature is empirically relevant for observations at high frequencies, as the bid-ask spread, the discrete price grid and order splitting tend to induce serial dependence in both transaction and mid-quote returns.Recognizing this property, the original realized volatility estimation procedures recommended sampling relatively sparsely, e.g., Andersen and Bollerslev (1998). Subsequently, a variety of robust approaches were developed, like realized kernels, pre-averaging, or multi-scale procedures, e.g., Barndorff-Nielsen et al. (2008), Jacod et al. (2009), and Zhang et al. (2005).

3.4.2 The Classical Model

The i.i.d. noise model is obtained from the model in the preceding section by, in addition, imposing $\alpha = 1$, so

$$p_i \;=\; p_i^* \;+\; \varepsilon_i \,, \tag{26}$$

with $\mathbb{E}[\varepsilon_i \, \varepsilon_i^*] = 0$.

It implies that fundamental news is embodied in the price instantaneously and past pricing errors are corrected without delay, so noise shocks are fully dissipated by the next observation yielding $\mathbb{V}[r_i] = \sigma_*^2(1+2\lambda)$ and $\mathbb{C}ov[r_i \, , r_{i-1}] = -\lambda\sigma_*^2$. Hence, the immediate error correction induces a negative first-order serial correlation, but also implies that the return series is bereft of correlation beyond lag one.

This scenario corresponds to the "classical" random-walk-plus-iid-noise model, see, e.g., Zhou (1996), Ait-Sahalia et al. (2005b) and Bandi and Russell (2008). This formulation still provides the basic reference for empirical microstructure effects, yet it rules out both feedback effects and noise endogeneity. Here, we simply note that this model is invalidated, if we find significant evidence in favor of $0 < \alpha < 1,$ and, secondarily, $\gamma \neq 1$.

3.4.3 The Uncorrelated Endogenous Noise Model

By imposing $\alpha = 1$, we obtain a scenario involving endogenous but serially uncorrelated noise as discussed in Section 3.4.2 with $p_i \;=\; p_i^* \;+\; [\,(\gamma \;-\; 1)\,\varepsilon_i^* \;+\; \varepsilon_i\,]$ and the autocovariance given by $\mathbb{C}ov(r_i \, , r_{i-1}) \;=\; [\,(1 \;-\; \gamma)\,\gamma \;-\; \lambda\,]\;\sigma_*^2$. This specification can cause (local) positive first-order return auto-correlation if $\lambda < 1$, i.e., the concurrent return innovation does not fully incorporate the efficient price innovation, inducing subsequent momentum. This specification is explored at some length in Hansen and Lunde (2006). They note that the return variation may be estimated consistently within this setting through a simple first-order lag adjustment to the realized volatility estimator, as originally suggested by Zhou (1996).

The general point is that noise endogeneity and information feedback are distinct features and have different implications. The endogeneity can alter the sign of the first-order return autocorrelation, while the noise correlation will generate longer lasting correlation effects. One objective of this paper is to determine conditions under which we separately can identify the underlying source of noise from the observed high-frequency return dependencies.

3.4.4 The Amihud-Mendelson Model

Our model (2)–(3) modifies the standard representation to allow for alternative economic mechanisms that generate more persistent return autocorrelation effects. Meanwhile, the finance literature contains several models designed to accommodate serial correlation at low frequencies. One prominent such model, Amihud and Mendelson (1987), emerges as a special case.

If $\gamma = \alpha$, our model implies that,

$$p_i \;=\; (1 \,-\, \alpha)\, p_{i-1} \,+\, \alpha\, p_i^* \,+\, \varepsilon_i\,. \tag{27}$$

In this scenario, the mid-quote price is a weighted average of the past price and the *current* efficient price, with an innovation term, that is uncorrelated with the efficient price innovation.

Consequently, the price dynamics in equation (27) mimics the one in equation (3), except that the price adjustment in the latter case is based on the discrepancy between the lagged observed price and *past* efficient price. Equation (27) is motivated by the idea that traders update their assessment of the fundamental price within each observation unit and modify their quotes (and anticipated transaction prices) accordingly. Amihud and Mendelson (1987) apply their specification to *daily* returns. In that setting, updating based on the *current* efficient price (established over the full trading day) is sensible relative to one-day-old information. In contrast, within a high-frequency environment featuring incessant order book revisions, multiple trades per second, and near continuous news feeds, it is less plausible that traders have the capability (and all relevant information) to gauge the efficient price every instant. Instead, we assume that efficient price updates occur with a slight, potentially time-varying, delay, representing a few seconds. In other words, we hypothesize that our model is more suitable at high frequencies (small Δ), while model (27) may be appropriate for large Δ. A major advantage of our local regime model is that alternative price formation mechanisms can be subjected to empirical scrutiny within a unified general framework.

3.4.5 The Information Delay Model

Importantly, a representation that yields first and second return moments identical to those in Amihud and Mendelson (1987) may be obtained through a different mechanism. As discussed in Section 3.2, we may impose the restriction $\gamma = 0$ to reflect the fact that even well-informed agents cannot respond instantaneously to newly arriving private information, but only with a (small) delay because of difficulties in interpreting the ongoing changes to the order book, the actual transactions consummated in the market place, and the presence of asymmetric information. Instead, this may lead to a lagged price response or feedback, which is stronger or weaker depending on the information and market environment at the time, motivating our

approach of keeping the model parameters fixed only over short intraday intervals.

In this "information processing delay" representation ($\gamma = 0$), the only contemporaneous shock impacting the price is idiosyncratic noise. Simultaneously, the price will display adjustments reflecting past pricing errors, enabling the efficient price innovations to govern the dynamics over time, as the market disentangles the noise and fundamental shocks to prices.

Letting $\gamma = 0$, equations (2)–(3) generate the following price dynamic, due to information processing delay,

$$p_i \;=\; (1 \,-\, \alpha)\, p_{i-1} \,+\, \alpha\, p^*_{i-1} \,+\, \varepsilon_i \,. \tag{28}$$

A few comments are in order. First, since ε^*_i is latent and uncorrelated with ε_i, we obtain an equivalent representation by relabeling p^*_{i-1} as p^*_i. In turn, this renders equations (27) and (28) identical. Hence, we cannot separately identify the two models – they generate the identical process for the observed returns. This type of issue was noted already in Section 3. It complicates identification and inference in general, as we discuss further in Section 4.1 below.

Second, the high-frequency covariance structure of the mispricing and efficient price innovations now take the simple form,

$$\Sigma \;=\; \begin{bmatrix} \mathbb{E}[(\varepsilon^\mu_i)^2] & \mathbb{E}[\varepsilon^\mu_i \varepsilon^*_i] \\ \mathbb{E}[\varepsilon^\mu_i \varepsilon^*_i] & \mathbb{E}[(\varepsilon^*_i)^2] \end{bmatrix} \;=\; \sigma^2_* \begin{bmatrix} 1+\lambda & -1 \\ -1 & 1 \end{bmatrix}$$

This result reflects that, with $\gamma = 0$, the efficient innovation can only be incorporated into market prices with a delay, so the mispricing component absorbs the full impact of any ε^*_i shock, generating strong endogeneity of the noise process. In this scenario, the model dynamics is the only source of coherence between the market and efficient price. Hence, it features a market in a perpetual state of transition, driven by an ongoing process of price discovery. Conceptually, this separates our specification from standard microstructure models. In the latter, prices are in equilibrium, except for short-lived distortions induced by exogenous mean-zero noise shocks.

Third, the above represents one aspect in which the "information delay" model differs in terms of economic interpretation from the high-frequency Amihud and Mendelson (1987) model, as the latter has a covariance structure identical to the one for the general representation (14), except for $\gamma = \alpha$. For either model, once the information processing is initiated, it continues at a rate governed by α. The latency of the efficient innovation implies that we cannot determine if this occurs with a temporal delay or not. Irrespective, the prolonged adjustment renders the noise process endogenous. Typically, this type of endogeneity is imposed through explicit statistical assumptions, see, e.g., Kalnina and Linton (2008), but it arises more structurally in our specification.

In summary, our information delay model or, equivalently, our high-frequency reinterpretation of Amihud and Mendelson (1987), featuring just four parameters (α, γ, λ, σ_*), endows the asset pricing process with distinct dynamic properties that reflect underlying economic forces. In particular, the inclusion of γ allows us to embed many existing structural microstructure models and test various specifications against one another. The subsequent sections explore the theoretical ramifications for the high-frequency return process, while the empirical evidence in Section 5 illustrates how this minimalistic framework captures salient features of the data generating process over short intraday trading periods, providing a superior framework for assessing the economic forces at play relative to traditional structural microstructure representations.

3.5 The Realized Variance

This section reviews the implications of the preceding results for realized variance measures computed from return data generated by model (2)–(3). The realized variance for interval $[0, T]$, using equidistant log-price observations, $p_{i\Delta}\,, i = 0, \ldots, n = n(\Delta) = \lfloor T/\Delta \rfloor$, is obtained as

$$RV_T(\Delta) \;=\; \sum_{i=1}^{n(\Delta)} \left(p_{i\Delta} - p_{(i-1)\Delta} \right)^2 \;=\; \sum_{i=1}^{n(\Delta)} \left(r_{i\Delta} \right)^2 , \tag{29}$$

where we have chosen to make the dependence of n on Δ explicit.

In the frictionless and informationally efficient version of our model, $\alpha = \gamma = 1$ and $\lambda = 0$, we have,

$$\mathbb{E}\left[\, RV_T(\Delta) \,\right] \;=\; T\,\sigma_*^2 . \tag{30}$$

In the case of general exogenous noise, $r_i \;=\; \varepsilon_i^* \;+\; \Delta\mu_i$, as discussed in Section 3.4.1, we find, in analogy to the result established there, that,

$$\mathbb{E}\left[\, RV_T(\Delta) \,\right] \;=\; T\,\left[\, \sigma_*^2 \,+\, \mathbb{E}[\,(\Delta\mu_i)^2\,] \,\right] \;\geq\; T\,\sigma_*^2 . \tag{31}$$

The point is that the realized volatility measure cannot be downward biased in the absence of endogenous noise. Moreover, the extent of any upward bias is determined by the properties of the noise process. The usual assumption is that every price observation contains a noise component unrelated to the sampling interval, so $\Delta\mu_i$ constitutes an $O_p(1)$ term which, for $\Delta \to 0$, dominates the $O_p(\sqrt{\Delta})$ term stemming from the efficient price innovation. Consequently, standard models generate a growing upward bias as the sampling frequency increases. Since the volatility signature plot displays the $RV_T(\Delta)$ measure over $[0, T]$, as a function of the underlying return interval, Δ, this implies a positively sloped volatility signature plot for $\Delta \to 0$.

Our empirical evidence indicates that the above is a poor characterization of the realized

volatility process in general. The illustrations in Figure 1 show that we encounter sample periods in which the traditional volatility signature plot is sharply upward sloping and others in which it is sharply downward sloping as a function of the sampling frequency. The discussion in Section 3.4.1 reminds us that the latter only can happen if the noise process is endogenous. Specifically, for our parametric model (2)–(3), but imposing endogenous noise, $\lambda = 1$, we obtain,

$$\mathbb{E}\,[\,RV_T(\Delta)\,] \;=\; T\,[\,1 \,+\, 2\,\lambda\,]\,\sigma_*^2\,. \tag{32}$$

In contrast, for the full parametric model with $\gamma \neq 1$,

$$\mathbb{E}\,[\,RV_T(\Delta)\,] \;=\; T\,\left[\,1 \,+\, \frac{2}{2\,-\,\alpha}\,F_\alpha(\gamma,\lambda)\right]\,\sigma_*^2\,. \tag{33}$$

As before, the sign of $F_\alpha(\gamma,\lambda)$ determines whether the return volatility measure is upward or downward biased . Both cases are empirically important, as they each occur with regularity, showing that a (time-varying) degree of exogenous and endogenous noise is required to capture the observed features of the volatility process for individual large capitalization equity returns.

4 Statistical Inference

4.1 Identification

As illustrated already in Section 3, our general model (2)-(3) is not uniquely identified from the second order return moments. This type of identification issue is well recognized in the literature. There are two distinct sets of assumptions regarding the noise, leading to starkly different conclusions. One approach casts the model in continuous time and conducts inference via an in-fill asymptotic scheme, where the number of observations (ticks) diverges within a fixed interval. In this case, the noise dependence (in tick time) can generally be identified and estimated. The reason is that the size of the tick-by-tick innovations in the fundamental price shrinks in line with the observation interval, while the noise correlation is fixed in tick time. As a consequence, the noise is asymptotically "big" and it will constitute the dominant term in the limit, allowing for nonparametric identification and inference. A thorough development of this approach is provided by Jacod et al. (2017). Likewise, Da and Xiu (2017) obtain consistent inference from a tick-time MA(q) process for the noise in a continuous-time setting with independence between the noise and fundamental price, but otherwise very general assumptions.

In contrast, if the noise dependence is given in calendar time, then the relative size of the noise and of the fundamental price increments does not shrink asymptotically. In this case, separate identification of the noise and fundamental return process hinges on the stipulation

that the fundamental price process is a martingale, while the noise is stationary. The latter implies some form for mean reversion, but this can only be ascertained over time intervals of increasing (asymptotic) size. Moreover, dependence among the noise and fundamental price process complicates identification even further, as noted in Ait-Sahalia et al. (2006). Typically, one can only accomplish this task via careful parametric modeling.

Our general (discrete-time) specification illustrates the latter point. Without endogeneity ($\gamma = 1$), the system is identified through the parametric specification, but once such dependence is introduced, we have only partial identification. Nonetheless, we can estimate the key parameter determining the fundamental price volatility, namely σ_*^2. It is the finer details of the noise dynamics that eludes us, as the second order moments cannot fully disentangle the relative size (variance) of the noise component from the delay with which the observed price incorporates the fundamental price innovation.

Specifically, in our setting, the challenge revolves around separate identification of the parameters λ and γ. They enter the expressions for the return variation and all return autocovariances in equations (15) and (20) only through the function $F_\alpha(\gamma, \lambda) = \gamma^2 - \alpha\gamma + \alpha + \lambda - 1$. In particular, any combination of γ and λ yielding the same value for $F_\alpha(\gamma, \lambda)$ imply identical second-order return moments (for given α and σ_*^2). In contrast, it is readily shown that α and σ_*^2 can be identified. Henceforth, in this section, we occasionally simplify notation and refer to $F_\alpha(\gamma, \lambda)$ simply by F_α.

There are natural constraints on γ and λ, such as non-negativity. Yet, whenever they are not at a boundary value, any small shift in γ may be offset by a shift in λ, so as to keep F_α unchanged. Thus, while we can identify F_α from the second order return moments, we cannot, in general, obtain separate identification for γ and λ. Since F_α is linear in λ, but quadratic in γ, fixing γ allows us to uniquely identify λ (and all other parameters). Conversely, as noted in Section 3, fixing λ does *not* generally allow us to uniquely identify γ, as scenarios, in which $|\gamma - \alpha/2|$ are identical, are observationally equivalent. When estimating the general model, we therefore focus on the lower-dimensional parameter vector $(\alpha, \sigma_*^2, F_\alpha)$, while being cognizant of the fact that the value for F_α is consistent with a range of admissible (γ, λ) combinations for the given return dynamics.

We illustrate the geometry for regions of non-identification for the pair (γ, λ) by considering two scenarios associated with positive and negative return autocorrelations, respectively. In Regime I, we choose the parameter values $\alpha = 0.835$, $\sigma_*^2 = 0.986 \cdot 10^{-8}$ and $F_\alpha = -0.100$ producing a *positive* first-order autocorrelation of 0.087. In Regime II, we set $\alpha = 0.976$, $\sigma_*^2 = 0.524 \cdot 10^{-8}$ and $F_\alpha = 0.091$, resulting in a *negative* first-order autocorrelation of -0.074. The values chosen for α, σ_*^2 and F_α are representative for our empirical results in Section 5, as they reflect the median estimates $\hat{\alpha}$, $\hat{\sigma}_*^2$ and $\hat{F}_\alpha$ obtained across local 10-minute intervals with

positive and negative return autocorrelations, respectively.

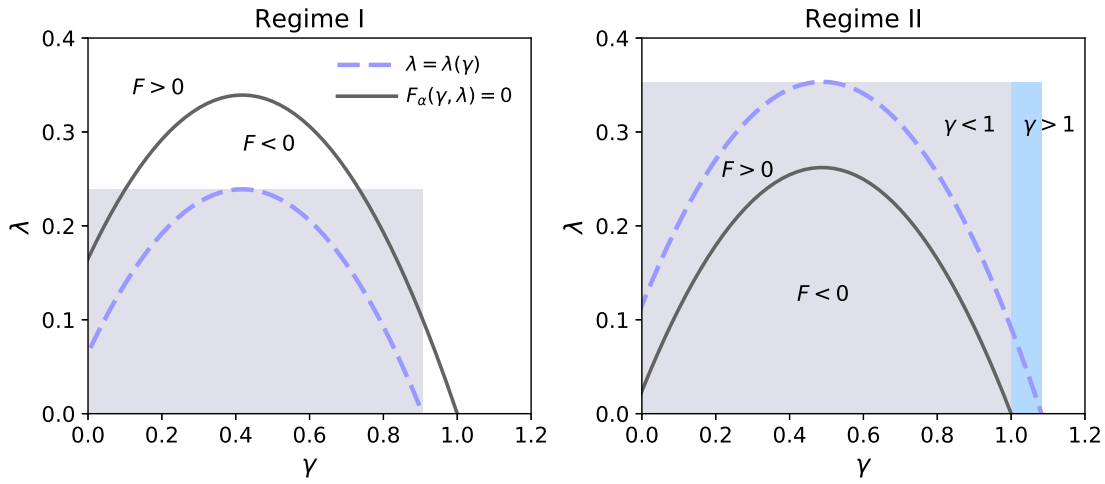

Figure 4: Dashed blue lines describe admissible values of (γ, λ) which correspond to the model parameters of Regime I (left plot) and Regime II (right plot). Regime I is parametrized by $\alpha = 0.835$, $\sigma_*^2 = 0.986 \cdot 10^{-8}$ and $F_\alpha = -0.100$ and generates a positive return autocorrelation. Regime II is parametrized by $\alpha = 0.976$, $\sigma_*^2 = 0.524 \cdot 10^{-8}$ and $F_\alpha = 0.091$ and generates a negative return autocorrelation. Solid black lines represent values of (γ, λ) which solve equation $F_\alpha(\gamma, \lambda) = 0$ for a given value of of α.

Figure 4 depicts the loci of (γ, λ) admissible for the parametrization of Regime I (left) and Regime II (right) under the constraint of non-negativity for γ and λ. The shaded rectangular area indicates the range of values which γ (horizontal axis) or λ (vertical axis) can take given the remaining model parameters. The loci of admissible (γ, λ) are indicated by the dashed blue line, implying that λ is a function of γ: $\lambda(\gamma) = -\gamma^2 + \alpha\gamma - \alpha + F_\alpha + 1$, for F_α and α determined by the given regime.

The solid black line depicts the loci of (γ, λ) solving $F_\alpha(\gamma, \lambda) = 0$ for a given value of α. All points below this line produce $F_\alpha(\gamma, \lambda) < 0$, whereas all points above it imply $F_\alpha(\gamma, \lambda) > 0$. Accordingly, for Regime I (positive autocorrelation) and negative F_α, all admissible pairs (γ, λ) lie below the parabolic curve $F_\alpha(\gamma, \lambda) = 0$. In contrast, for Regime II (negative autocorrelation), all admissible values of (γ, λ) are located above $F_\alpha(\gamma, \lambda) = 0$.

This analysis enables us to determine the maximum admissible values of γ and λ within a given regime. For λ, the largest admissible value, λ_{max}, is attained at the level of γ for which the return variance is minimized, i.e., $\gamma = \alpha/2$,

$$\lambda_{max} = \frac{\alpha^2}{4} - \alpha + F_\alpha + 1. \tag{34}$$

Correspondingly, it is readily seen that γ_{max} is attained for $\lambda = 0$, so that,

$$\gamma_{max} = \frac{\alpha}{2} + \sqrt{\lambda_{max}}\,. \tag{35}$$

Hence, although the parameters γ and λ are not point identified, we can compute algebraic upper bounds from estimates for the identified parameters $(\alpha, \sigma_*^2, F_\alpha)$.

We further note that the maximum γ satisfying the boundary condition $F_\alpha(\gamma, \lambda) = 0$ is attained for $\lambda = 0$, and equals one. Therefore, $F_\alpha < 0$ always implies that $\gamma_{max} < 1$ and $F_\alpha > 0$ implies $\gamma_{max} > 1$. This property is nicely illustrated in Figure 4. When $F_\alpha < 0$ (Regime I), the dashed curve $\lambda(\gamma)$ is always below the black curve $F_\alpha(\gamma, \lambda) = 0$, thus crossing the horizontal axis below 1, so $\gamma_{max} < 1$. When $F_\alpha > 0$ (Regime II), the situation is reversed. Here, market overreactions to concurrent return innovations ($\gamma > 1$) can only occur if $F_\alpha > 0$, thus implying a negative first-order autocorrelation. Conversely, whenever the serial correlation is positive, concurrent return innovations are never fully incorporated.

These identification relationships have several important implications for the special cases of our model discussed in Section 3.4. Firstly, the equivalence between the Amihud-Mendelson (Section 3.4.4) and Information Delay model (Section 3.4.5) follows from the fact that the constraints $\gamma = \alpha$ and $\gamma = 0$ generate the same $F_\alpha(\gamma, \lambda) = \alpha + \lambda - 1$. Secondly, the classical model with idiosyncratic noise (Section 3.4.1) is characterized by $\gamma = 1$ and $F_\alpha(1, \lambda) = \lambda > 0$, thus excluding return dynamics with positive serial correlation. As demonstrated in Section 2, this feature is not supported by empirical evidence. Hence, the classical model with i.i.d. noise (Section 3.4.2), often viewed as a benchmark in the literature, cannot accommodate the most salient features of the high-frequency return dynamics. As a consequence, relying on the classical model for volatility measurement is tantamount to the imposition of restrictions on α and F_α, that are incompatible with the basic features of the return dynamics, and estimates of the efficient return variance, σ_*^2, will be biased.

The model with endogenous noise (Section 3.4.3) and $\alpha = 1$ implies $F_1(\gamma, \lambda) = \gamma^2 - \gamma + \lambda$. Contrary to the classical idiosyncratic noise model, This specification allows for a positive serial correlation for the observed returns. Due to the constraints discussed above, we have $F_1(\gamma, \lambda) \geq -\frac{1}{4}$, thus the first-order autocorrelation is bounded by $(-\frac{1}{2}, \frac{1}{2})$.

Finally, we note that the information delay, or equivalently the Amihud-Mendelson, model ($\gamma = 0$) provides a representation facilitating identification without severely restricting the empirical implications. Fixing γ lets us uniquely identify the remaining parameters, and thus $F_\alpha = \alpha + \lambda - 1$, without imposing additional tight constraints on the return dynamics. The key is that the specification allows for instant partial incorporation of the fundamental price innovations along with the feedback mechanism associated with dynamic price discovery.

4.2 Estimation

For estimation purposes, it is convenient to restate the model as,

$$\begin{aligned} r_i &= -\alpha\,\mu_{i-1} + \tilde{\varepsilon}_i, \\ \mu_i &= (1-\alpha)\,\mu_{i-1} + \varepsilon_i^{\mu}, \end{aligned}$$

where $\tilde{\varepsilon}_i = \gamma\varepsilon_i^* + \varepsilon_i$ and $\varepsilon_i^{\mu} = \tilde{\varepsilon}_i - \varepsilon_i^* = (\gamma-1)\varepsilon_i^* + \varepsilon_i$. Then, returns r_i are driven by a latent variable μ_i following an AR(1) process with covariance structure given by,

$$\begin{aligned} \mathbb{E}[\tilde{\varepsilon}_i\,\mu_{i+h}] &= (1-\alpha)^h\,\left(\gamma^2+\lambda-\gamma\right)\,\sigma_*^2\,, \qquad \forall h \geq 0, \\ \mathbb{E}[\tilde{\varepsilon}_i\,\mu_{i-h}] &= 0, \qquad \forall h > 0, \\ \mathbb{E}[\tilde{\varepsilon}_i\,\varepsilon_i^{\mu}] &= [\,\gamma\,(\gamma-1)+\lambda\,]\,\sigma_*^2, \\ \mathbb{E}[\tilde{\varepsilon}_i\,\varepsilon_{i-h}^{\mu}] &= 0, \qquad \forall h \neq 0. \end{aligned}$$

Define Y as the vector of observed returns, $Y = (r_1, \ldots, r_T)$ and $\theta = (\sigma_*^2, \sigma_\varepsilon^2, \alpha, \gamma)'$ as the parameter vector to be estimated. Assuming ε_i and ε_i^* are i.i.d. Gaussian with variances σ_ε^2 and σ_*^2, the log-likelihood function is (up to the constant $-T\ln(2\pi)/2$) given by the log density of the corresponding T-dimensional zero-mean normal distribution with covariance matrix Σ,

$$\ell(Y,\theta) = -\ln|(\Sigma)|\,/\,2 - Y'\,\Sigma^{-1}\,Y/2, \tag{36}$$

where the diagonal and off-diagonal elements of Σ are given, respectively, by,

$$\Sigma^{ii} = \sigma_*^2\,(1 + 2\,F_\alpha(\gamma,\lambda)\,/\,(2-\alpha)\,), \tag{37}$$

$$\Sigma^{ij} = -\sigma_*^2\,F_\alpha(\gamma,\lambda)\,\psi(\,|i-j|-1\,)\,, \tag{38}$$

where $\psi(\,\cdot\,)$ is defined below equation (20) in Lemma 2.

Under standard regularity conditions for maximum likelihood (ML) estimation, we obtain the following asymptotic limiting distribution for the ML estimates $\hat{\theta}$,

$$T^{1/2}\,(\hat{\theta}-\theta) \xrightarrow{d} N(\,0, \mathcal{I}_1^{-1}\,),$$

with $\mathcal{I}_1 = T^{-1}\mathcal{I} = -T^{-1}\,\mathbb{E}[\,\partial^2\ell(Y,\theta)/\partial\theta\partial\theta'\,]$. Consequently, the quadratic variation over all T observations, $T\sigma_*^2$, is estimated at rate $T^{1/4}$. Define DC as the $mn \times q$ matrix $\partial vec(C)/\partial\theta'$, where C is a $m \times n$ matrix depending on a $q \times 1$ parameter vector θ. Then, $\partial\ell(Y,\theta)/\partial\theta'$ and

$\partial^2 \ell(Y,\theta)/\partial\theta\partial\theta'$ can be computed as,

$$\partial \ell(Y,\theta)/\partial\theta' = -\frac{1}{2}\left(vec'[\Sigma^{-1}(I_T - yy'\Sigma^{-1})]D\Sigma\right) \tag{39}$$

$$\begin{aligned}\partial^2 \ell(Y,\theta)/\partial\theta\partial\theta' = -\frac{1}{2}\{&-D\Sigma'((I_T - \Sigma^{-1}yy')\Sigma^{-1}\otimes\Sigma^{-1})D\Sigma \\ &+D\Sigma'(\Sigma^{-1}\otimes\Sigma^{-1}yy'\Sigma^{-1})D\Sigma\},\end{aligned} \tag{40}$$

where $\otimes$ denotes the Kronecker product. Then, $\mathbb{E}[\partial^2 \ell(Y,\theta)/\partial\theta\partial\theta']$ is given by,

$$\mathbb{E}[\partial^2 \ell(Y,\theta)/\partial\theta\partial\theta'] = -\frac{1}{2}\left\{D\Sigma'(\Sigma^{-1}\otimes\Sigma^{-1})D\Sigma\right\}. \tag{41}$$

$D\Sigma$ is readily obtained from partial derivatives of equations (37) and (38) with respect to θ.

Analytically tractable formulas for Σ^{-1}, however, are available only for the special case $\alpha = 1$, when Σ collapses to a tri-diagonal matrix, as $\Sigma^{ij} = 0$ for $|i-j| > 1$. In this scenario, the elements of Σ^{-1} may be computed using results by Usmani (1994). If $\alpha = 1$ and $\gamma = 1$, the autocovariance structure of r_i is governed by an MA(1) process, and Σ^{-1} takes a convenient form, see Hamilton (1994). Ait-Sahalia et al. (2005a) exploit this formulation to derive a tractable approximation to $\mathbb{E}[\partial^2 \ell(Y,\theta)/\partial\theta\partial\theta']$ for $T \to \infty$.

In the general case, however, Σ is of a special Toeplitz form for which the inverse of Σ is not available in closed form. It is then more convenient to formulate the log likelihood in terms of the prediction error decomposition, see, e.g., Harvey (1989),

$$\ell(Y,\theta) = -\frac{1}{2}\sum_{i=1}^{T}\ln(s_i) - \frac{1}{2}\sum_{i=1}^{T}\frac{\nu_i^2}{s_i^2}, \tag{42}$$

where ν_i denote (optimal) linear predictions of r_i given past returns $(r_{i-1},\ldots,r_1)$, and s_i^2 denotes the conditional variance $s_i^2 = \mathbb{E}[\nu_i^2 | r_{i-1},\ldots,r_1]$.

Given the normality for ε_i and ε_i^*, $\ell(Y,\theta)$ can be readily computed by the Kalman filter based on a linear state-space representation of the model. Denote X_i as a state vector at i with, $X_i = (\mu_i, \mu_{i-1}, \tilde{\varepsilon}_i)'$. Then, r_i may be written as,

$$\begin{aligned} r_i &= FX_i, \\ X_i &= GX_{i-1} + w_i, \end{aligned}$$

with $F = (0, -\alpha, 1)$ and

$$G := \begin{pmatrix} (1-\alpha) & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

The residual vector is given by $w_i = (\varepsilon_i^{\mu}, 0, \tilde{\varepsilon}_i)$ with covariance matrix,

$$\Sigma_w = \begin{pmatrix} (\gamma-1)^2\sigma_*^2 + \sigma_\varepsilon^2 & 0 & \sigma_*^2[\gamma(\gamma-1)+\lambda] \\ 0 & 0 & 0 \\ \sigma_*^2[\gamma(\gamma-1)+\lambda] & 0 & \gamma^2\sigma_*^2 + \sigma_\varepsilon^2 \end{pmatrix}.$$

To capture the covariance structure between ε_i^{μ} and ε_i , we treat ε_i and μ_{i-1} as separate latent variables and thus components of X_i. Accordingly, the system features three state equations.

The Kalman filter (Kalman (1960), Kalman (1963)) produces optimal forecasts for the latent variables X_i given observations up to $i-1$, and consequently optimal predictions $r_{i,i-1}$, minimizing the mean-squared error s_i, see Hamilton (1994). Given the linear Gaussian state space model, the Kalman filter therefore enables us to compute the exact log-likelihood function yielding (asymptotically) efficient inference. The construction of the log likelihood and the corresponding derivatives $\partial\ell(Y,\theta)/\partial\theta'$ and $\partial^2\ell(Y,\theta)/\partial\theta\partial\theta'$ is provided in the Appendix.

If the errors ε_i^* and ε_i (and state variables X_i) are *not* normally distributed, equation (42) may be interpreted as a *quasi* maximum likelihood function. Under non-normality, the linear predictions of the state variables X_i (implied by the Kalman filter algorithm) are not optimal among all prediction functions, but remain the best *linear* ones. This argument renders the (linear) Kalman filter applicable in a quasi maximum likelihood setting, as discussed in Gourieroux et al. (1984). A formal proof for the consistency of the Kalman-filter based (Q)ML estimator in a non-Gaussian linear state space system is given by Schlemm and Stelzer (2012).

Alternatively, the system could be estimated directly from the unconditional moment restrictions implied by the model following the standard GMM procedure of Hansen (1982).

5 Evidence from Model Estimation

In this section, we explore the high-frequency price dynamics within local intra-daily intervals by estimating the parametric model introduced in Section 3. In our analysis, we use intra-daily mid-quote returns from the 100 stocks described in Section 2 and conduct estimation for 10 minute segments for the first 3 months of 2014 (61 trading day). Thus, in our empirical analysis we inspect a total of 237,900 local intervals.

5.1 Parameter Restrictions for Model Identification

Since the full set of parameters of the general model are only partially jointly identified, we initially impose a convenient restriction, namely $\gamma = 0$. This choice is motivated by our prior analysis of model identification. In Section 4.1, we demonstrated that $\gamma = 0$ is admissible irrespective of the regime governing the return dynamics. Therefore, the model parameters α, σ_*^2 and $F_\alpha(\gamma, \lambda)$ remain unrestricted and are uniquely identified, with λ simply adjusting to accommodate the estimated value for $F_\alpha(0, \lambda) = \widehat{F}$. Next, we remove the $\gamma = 0$ constraint and instead determine the set of (γ, λ) values compatible with $F_\alpha(\gamma, \lambda) = \widehat{F}$. In the process, we readily obtain algebraic upper bounds for λ_{max} and γ_{max}.

5.2 Parameter Estimates

We estimate the model by (pseudo-)maximum likelihood using the Kalman filter as described in Section 4.2. We obtain estimates for each stock in the NASDAQ 100 over the first 61 trading days of 2014, for each T = 10 min interval, sampling at the Δ= 2 second frequency.

5.2.1 The Empirical Parameter Distribution

Figure 5 displays the empirical distributions of the parameter estimates obtained across all stocks in the first, fourth and fifth quintile, respectively, according to the intensity of quote-midpoint revisions, as described in Section 2. All the distributions appear unimodal and approximately symmetric, apart from the distinct right skew in the distribution for σ_*^2.

For the stocks in Group 1, the estimates of α have a mode of around 0.9, corroborating the hypothesis of a feedback effect and a mildly sluggish price adjustment. The distribution of γ_{max} is centered around 1. This is consistent with our prior finding in Table 1 that the empirical first-order return autocorrelation is centered on zero, but with a significant number of both positive and negative coefficients across the individual local intervals. These two regime scenarios imply, respectively, that $\gamma_{max} < 1$ and $\gamma_{max} > 1$, and these values are achieved only for $\lambda = 0$, as is also evident from the left and right panels of Figure 4. Hence, at a minimum, for a large set of intervals, we must have $\gamma < 1$. Moreover, returns measured over seconds invariably reflect some degree of noise, as price discreteness and other frictions are non-trivial at such frequencies. Hence, $\lambda > 0$, and the relevant parameter $\gamma < \gamma_{max}$, reflecting a leftward movement along the dashed curves in Figure 4, starting from $(\gamma_{max}, 0)$.

The above conclusion reflects the fact that there are two sources for violations of the standard martingale result for returns in our setting. Idiosyncratic noise will be "corrected" over time, inducing negative return serial dependence, while incomplete incorporation of fundamental price innovations will imply positive dependence, as the market subsequently adjusts to the

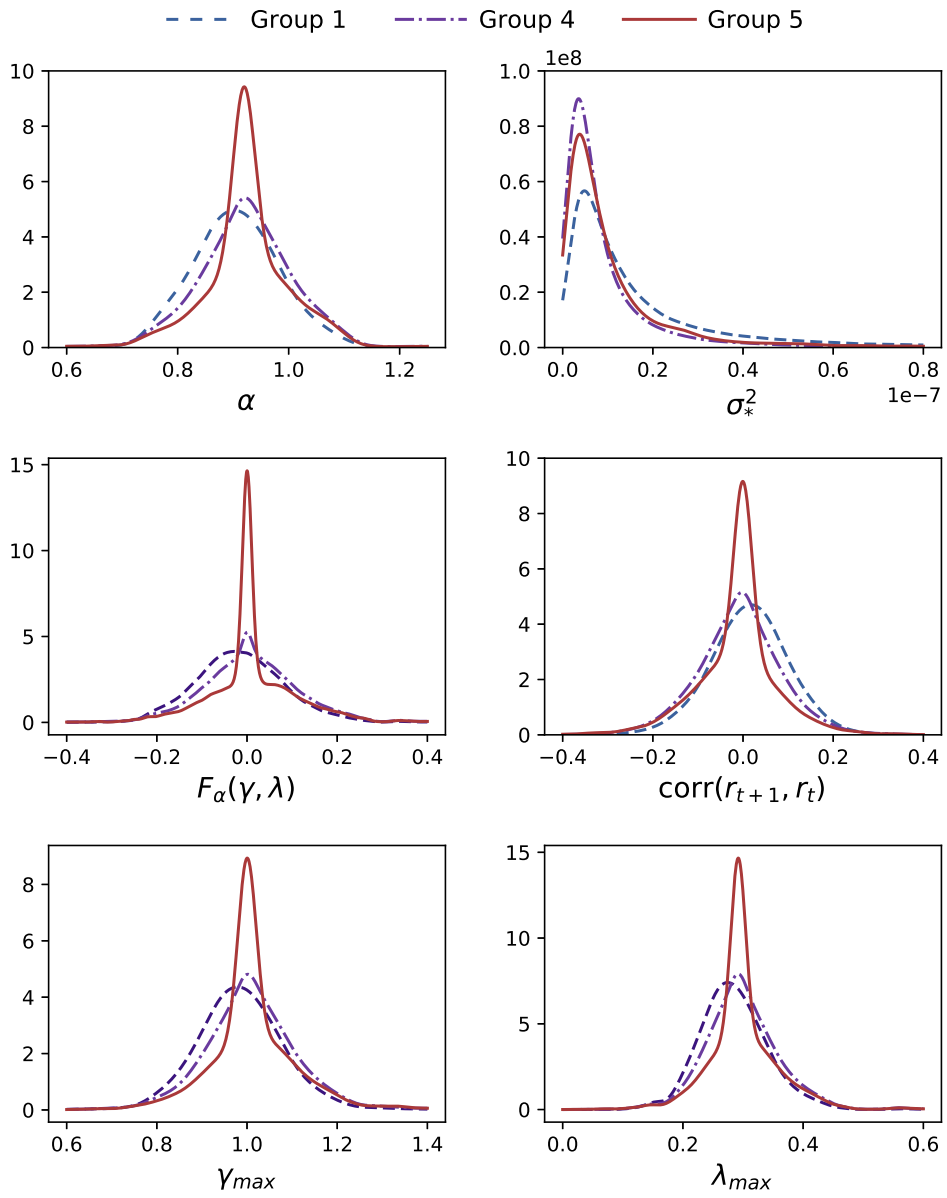

Figure 5: Empirical distributions of the model parameter estimates (estimated under restriction $\gamma = 0$) obtained for T = 10 min local intervals with returns sampled at frequency Δ= 2 sec. Resulted densities are shown for the stocks from 1st, 4th and 5th groups of stocks sorted by the average number of daily mid-quote revisions (with 20 stocks in each group) for the first 61 trading days in 2014. In total, each depicted density curve is constructed with parameter estimates from 47,580 local intervals.

shift in the underlying efficient price (for $1 > \gamma > \alpha/2$). Therefore, given the observed level of serial correlation, an increase in noise must be counterbalanced by a smaller instantaneous adjustment to fundamental price innovations.

From the preceding discussion we infer that, in the majority of intervals, the return dynamics is affected by endogenous noise. This typically induces a smoothing of the price path, and the overall effect may be a positive or negative return autocorrelation depending on the relative

size of the idiosyncratic noise component, λ. On the other hand, a non-trivial proportion of the local intervals featuring negative return autocorrelation may also have $\gamma > 1$, indicating an information environment where agents are concerned about the asymmetric information and tend to overreact to information signals, generating "overshooting." This induces subsequent price reversals that reinforce the impact of idiosyncratic noise.

We further note that our maximal estimates for the noise-to-signal ratio, λ_{max}, fall below 0.4 for the majority of the intervals. This confirms findings in the recent literature of a small idiosyncratic noise component for stocks in the major large cap indices. Consistent with the evidence provided in Section 2, the distribution of the estimated autocorrelations for stocks in Group 1 has a small positive mode and is right-skewed, indicating that, for the more liquid stocks, regimes with positive return autocorrelation are common and perhaps even in the majority.

In contrast, the stocks in Group 5 generate estimates of α and λ_{max} that are somewhat higher on average, have smaller dispersion around the mode and are distinctly right-skewed. That is, less liquid stocks with fewer quote revisions and larger relative tick sizes generate returns that are noisier and react more strongly to past pricing errors. Consequently, the model-implied return autocorrelatiosn are predominantly negative, as also indicated by the slight left skew in the associated empirical distribution. Table 7 in the Appendix provides detailed summary statistics for the parameter estimates in each of the five liquidity quintiles. The qualitative patterns discussed above are present in all quintiles and the relevant average values shift almost monotonically across the quintiles, thus corroborating our earlier findings.

5.2.2 Illustrative Daily Results

We now provide the price path, parameter estimates and empirical return autocorrelations for 10-minute intervals covering a single trading day for a couple of stocks. Figure 6 depicts the time series of estimated model parameters and the model-implied versus empirical first-order return autocorrelations using 2-second sampling on January 9, 2014, for Apple and Netflix. The figure is quite representative for other stocks and trading days.

From the top panel in Figure 6, we see that both stocks suffered a drop of about 1.5% over the course of the trading day. Moreover, their price paths are clearly correlated, with sharp losses realized between 10:00 and 11:30 and moderate recoveries taking place after 14:30. Nonetheless, they also display stark differences. The AAPL price path is generally more orderly, with only minor sudden sharp price swings and overall lower volatility. In contrast, the initial half hour of trading for NFLX is tumultuous, and the majority of the price decline then takes place between 10:00 to 10:30, followed by a largely flat but volatile price path for the remainder of the trading day. These discrepancies manifest themselves in the parameter estimates. Over the first half day of trading, the fundamental volatility (σ_*^2) is much higher for NFLX than

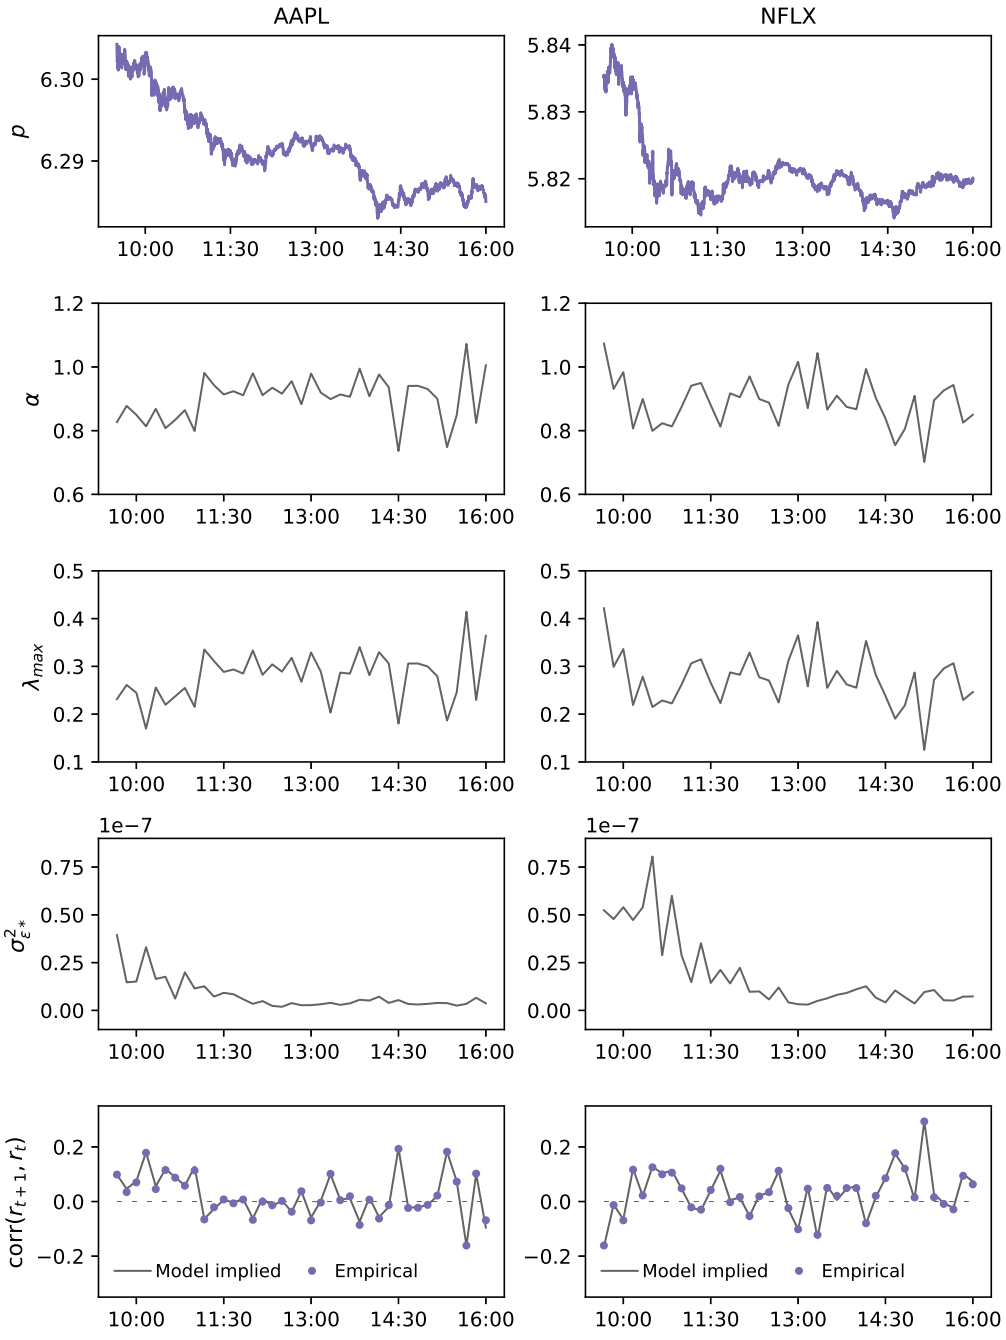

Figure 6: Model parameters estimated for AAPL (left side) and NFLX (right side) for T = 10 min local intervals on January 9th, 2014 (Δ= 2 sec).

AAPL, the speed with which pricing errors are corrected (α) is initially much more elevated for NFLX, the return dependence is of opposite sign at the early stages of trading – negative for NFLX and positive for AAPL – and the λ_{max} coefficient suggests that the idiosyncratic noise is particularly high for NFLX in the first trading hour. After a relatively stable period for the series across the middle of the trading day, it is arguably AAPL that displays the more erratic patterns

towards the close of trading, with the price volatility intensifying, the strength of the pricing error correction (and possibly idiosyncratic noise) picking up, and the return autocorrelations oscillating quite rapidly.

In summary, the market environment for these two stocks evolved quite differently on this trading day. For NFLX, one can sense a suspicion that new information needed to be incorporated into the price. After an initial phase of rapid moves up and down, characterized by a strong negative autocorrelation in the underlying high-frequency returns, the sentiment turned distinctly negative, and the price dropped quickly, while the fundamental volatility increased to an even higher level, and the return serial dependence turned positive. In contrast, AAPL underwent a relatively calm decline in prices from a lower level of fundamental price volatility, with positive return autocorrelation throughout this period, and a lower speed of pricing error correction. Later on in the trading day, however, some of these features become more apparent for AAPL, reflecting relatively more uncertainty about the state of the underlying information manifesting itself in the trading for AAPL than NFLX stocks towards the market close.

The purpose of our model is to render distinct and rapidly shifting features of the market environment, as those discussed above, quantifiable and interpretable. Our approach does this by estimating a flexible model to the locally observed second order return moments. The bottom panels of Figure 6 illustrate how well our model accommodates the varying first-order return autocorrelations. Importantly, the objective is not to simply fit these return moments, but rather to convert the observations into a simple structural framework that allow us to better disentangle the factors driving the market dynamics. Specifically, a period with observed negative return correlations may reflect either a high degree of idiosyncratic noise, a low degree of fundamental information flow, or a quicker pricing error correction, or a combination of such features. Each scenario will naturally translate into a different measure for the intensity of the information and frictional forces driving the significant intra-daily variation in the return autocorrelation regimes documented in Section 2.

5.3 Parameter Variation and Time-of-Day Effects

Figure 7 depicts time-of-day average estimates for α, λ_{max}, σ_*^2 and the first-order return autocorrelation across all trading days for the stocks in Groups 1 and 4. Qualitatively identical patterns are present for the other liquidity quintile groups.

Several of the panels reveal a pronounced intraday pattern. Most strikingly, the average $\hat{\sigma}_*^2$ values form a sharp L-shape, with high volatility in the morning followed by a quick decay and a largely flat pattern for the remainder of the trading day. The picture is similar for both set of stocks, except that the decay is more gradual for the Group 1 stocks. The similarities and differences between the group of stocks are even more evident for the return serial correlation.

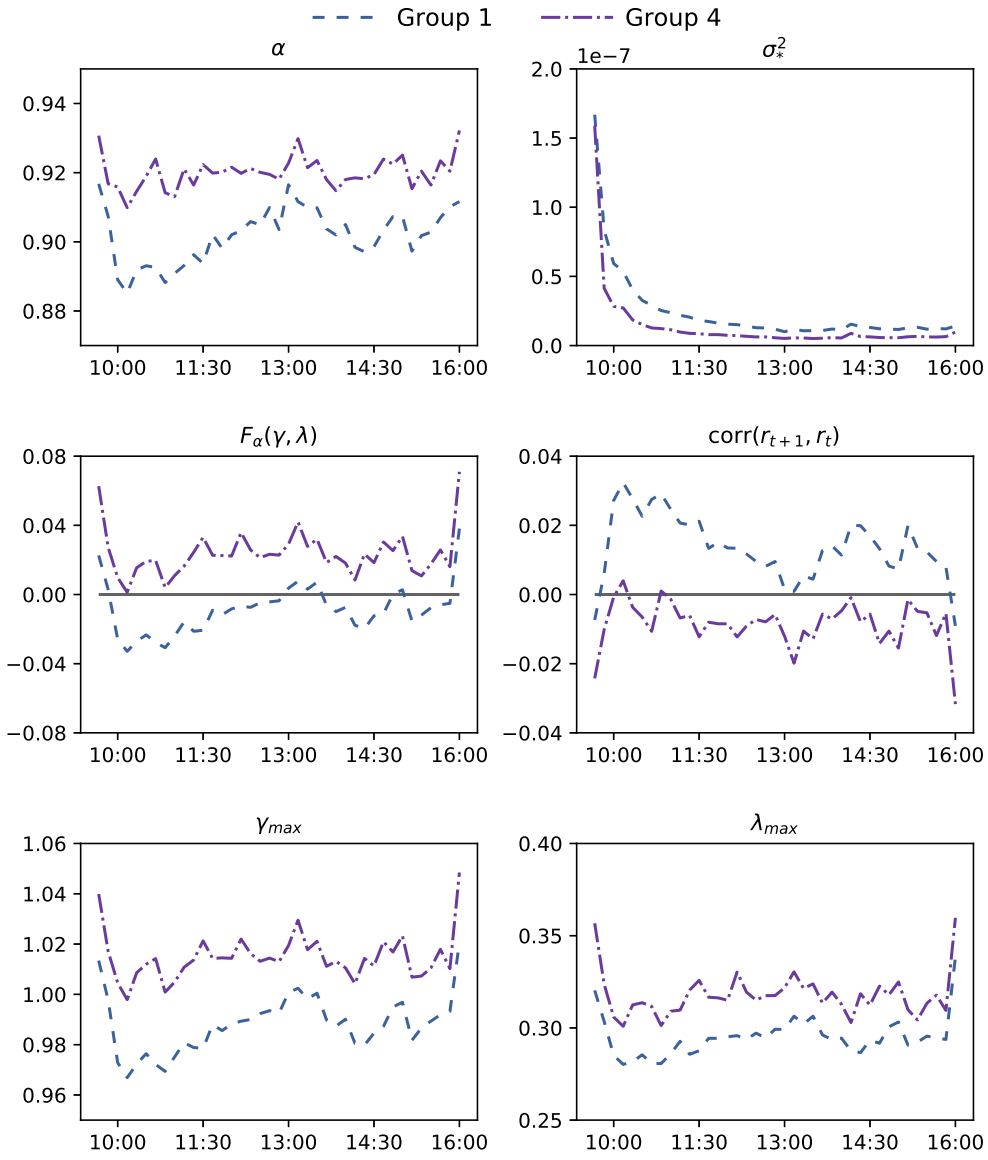

Figure 7: Intra-daily patterns of average parameter estimates (estimated under restriction $\gamma = 0$) obtained for T = 10 min local intervals with returns sampled at frequency Δ= 2 sec. Resulted patterns are calculated by averaging the estimates from all local intervals for the first 61 trading days in 2014 for the stocks from 1st and 4th groups sorted by the average number of daily mid-quote revisions (with 20 stocks in each group).

Both groups display sharply lower autocorrelation at the opening and closing of the trading day, while the level across the middle of the day resembles a mild u-shape. At the same time, the Group 1 values are consistently above those for Group 4, with the former being positive throughout except for the periods around the open and close, while the latter is almost uniformly negative.

The above features will naturally manifest themselves in our estimates for the remaining model parameters. For α, we observe an elevation at the open and close, while the average

estimates form a mild reverse u-shape across the main phase of the trading session. The bottom two panels display the (average) maximal γ and λ values consistent with the subset of fully identified parameters. By construction, they merely provide upper bounds, and they cannot be attained simultaneously. Specifically, γ_{max} applies for $\lambda = 0$, while λ_{max} is achieved for $\gamma = \alpha/2$. Since we, unquestionable, must have $\lambda > 0$ to capture idiosyncratic noise stemming from market frictions, it is clear that the average γ is below one almost uniformly for Group 1 stocks and, likewise, for most of the time in Group 4. Hence, the presence of endogenous noise is ubiquitous throughout the trading day. These considerations aside, the spike in the upper bound for both parameters at the boundary of the trading day are noteworthy, pointing again to an unusual market environment at these times.

How do we interpret this evidence? The most striking feature is the huge asymmetry between the open and close for the fundamental volatility and the largely symmetric behavior of all other series. Although return volatility is highly elevated both at the open and close, our model clearly associates only the opening with the arrival and assimilation of an unusual amount of fundamental information. The remaining parameters suggest that this raises the speed by which pricing errors are corrected (α), and likely also the instantaneous response of market prices to new information (γ) and the amount of idiosyncratic noise (λ). We conclude that the extraordinarily high fundamental volatility at the start of the trading day likely stems from price discovery associated with information processing of overnight news and orders. Given the (expected) elevation in the arrival of efficient price innovations following the market opening, dealers and market makers are unwilling to commit a large amount of capital to intermediation so we will see larger bid-ask spreads and less deep order books. Moreover, the sensitivity to every piece of market news will be enhanced. Both features tend to elevate the amount of idiosyncratic noise in the price discovery process. Finally, understanding that the efficient price is shifting quickly and idiosyncratic noise is present, resources are expended on adjusting rapidly to the incoming signals, so α rises. In all cases, we expect the return autocorrelation to drop, as we observe more immediate incorporation of fundamental news, a larger presence of idiosyncratic noise and less price smoothing.

The period prior to the close is different from the open in the sense that customer orders must be executed before the trading terminates to not be exposed to price uncertainty in the post exchange trading market and overnight hours. On the other hand, some popular incentive schemes for trade execution, like the VWAP criterion, motivate intermediaries to delay transactions towards the tail end of the day, as they acquire a better sense of the benchmark their trades will be assessed against. Executing trades early in the day carries an outright exposure to subsequent price movements that impact the "value-weighted average price." Hence, there is a natural non-informational rationale for an elevation in the trading volume and return volatility

towards the close. Absent an unusual degree of asymmetric information, the market is deeper. The end result is enhanced volatility that raises the responsiveness of market participants to the market signals, generating quicker price corrections, but also additional idiosyncratic noise given the prevalence of non-informational transactions.

Overall, we find evidence corroborating the idea that liquid market conditions early and late in the trading day foster efficient discovery and processing of fundamental news. At such times, the main component behind temporary misalignments in prices is market microstructure noise, consistent with negative return serial correlation. In contrast, when the information and order flow is slower, the resources devoted to price discovery shrink and it may be more difficult to identify less prevalent news items among ever-present noise. This effect tends to enhance price sluggishness, generating conditions that favor an element of momentum in the price formation.

Finally, the systematic differences across stocks with different liquidity characteristics are also noteworthy. The more rapid drop in the fundamental volatility level in the morning and the negative return autocorrelation suggest that the information environment is characterized by a higher degree of idiosyncratic noise for the less liquid stocks. If this renders the level of the efficient price less difficult to decipher–as most market signals are uninformative in such settings–this can also rationalize the higher level of the increase in the α parameter.

In the Appendix, we provide additional summary statistics of the parameter estimates for an individual analysis of the first trading interval from 9:30 to 9:40 (Figure 8 and Table 8) and the last trading interval from 15:50 to 16:00 (Figure 9 and Table 9). In these intervals, market conditions are unusual and the return dynamics differs substantially from the remainder of the trading day. In addition, we report summary statistics across all intraday intervals excluding the first and last 20 minutes of trading (Figure 10 and Table 10). This may be viewed as a robustness check for the statistics reported in Figure 5 and Table 7. In all cases, our qualitative findings remain unaltered.

6 Conclusions

This paper provides a novel perspective on "classical" martingale-plus-noise models for high-frequency mid-quote prices, as commonly used in the literature. The proposed modification of this model class by just one additional parameter is simple but powerful, as it yields a more structural perspective and opens up canonical links to existing approaches in market microstructure theory. We show that well-known properties of market microstructure noise, such as serial correlation as well as correlation with the efficient price ('endogeneity'), arise naturally from temporal feedback due to MMN-induced mispricing.

We identify error correction due to prevailing mispricing as a key source of high-frequency

dynamics in financial returns. The interplay of the speed of price adjustment and the noise-to-signal ratio governs the autocovariance structure of intraday returns and is shown to have a profound effect on realized volatilities sampled on high frequencies.

Serial correlation in high-frequency returns may therefore be rationalized by the extent of sluggishness in prices and the relative noisiness (or informativeness) of the price innovations. Our empirical analysis shows that the interplay between these two key parameters varies over time and triggers local intervals of 'momentum' behavior (positive autocorrelations) and 'contrarian' trading (negative autocorrelations). An important empirical finding is that particularly regimes of statistically significant positive serial correlations in high-frequency mid-quote returns (inducing downward sloped realized volatility signature plots) occur way more often than commonly expected.

Our results show price updating in line with the information delay model causing *endogenous* noise is empirically supported and is sufficiently flexible to capture high-frequency return dynamics over local intervals. An immediate consequence is that the "classical" martingale-plus-iid-noise model is clearly rejected. We conclude that temporal feedback, as captured by the model developed in this paper, may be a critical determinant for the dynamics of high-frequency returns.

From a statistical viewpoint, this paper makes important contributions. First, the proposed model gives a more structural treatment of market microstructure noise capturing the origins of endogeneity and temporal dependence. These insights can be exploited to construct more efficient estimators of volatility explicitly taking into account this noise structure. Second, we show that extensions of the classical martingale-plus-noise model are not necessarily identifiable in all generality. In fact, the signal-to-noise ratio is only partly identifiable if the covariance between market microstructure noise and efficient price increments is allowed to vary freely. This implies restrictions on the parameter space for which statistical inference can be conducted. In particular, it implies the observational equivalence between mis-pricing adjustments based on current (efficient) prices (according to Amihud and Mendelson (1987)) and mis-pricing corrections based on *lagged* (efficient) prices ("information delay model"). Third, our results provide strong evidence for locally changing return dynamics. These findings are important for local volatility estimation but also for the estimation of volatility over longer intervals, such as, e.g., a day or a week. Our results suggest that high-frequency-based volatility estimators should be constructed based on aggregates of local estimators that take the regime-switching nature of return processes explicitly into account. This gives interesting avenues for future work on the construction of more efficient volatility estimators.

References

AIT-SAHALIA, Y. AND J. JACOD (2014): *High-Frequency Financial Econometrics*, Princeton University Press.

AIT-SAHALIA, Y., P. MYKLAND, AND L. ZHANG (2005a): "How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise," *Review of Financial Studies*, 351–416.

——— (2006): "Comment," *Journal of Business & Economic Statistics*, 162–167.

AIT-SAHALIA, Y., P. A. MYKLAND, AND L. ZHANG (2005b): "How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise," *Review of Financial Studies*, 17, 351–416.

AMIHUD, Y. AND H. MENDELSON (1987): "Trading Mechanisms and Stock Returns: An Empirical Investigation," *Journal of Finance*, 42 (3), 533–553.

ANDERSEN, T., D. DOBREV, AND E. SCHAUMBURG (2012): "Jump-robust volatility estimation using nearest neighbor truncation," *Journal of Econometrics*, 169, 75–93.

ANDERSEN, T. G. AND T. BOLLERSLEV (1998): "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885–905.

ANDERSEN, T. G., T. BOLLERSLEV, F. DIEBOLD, AND P. LABYS (2003): "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 579–625.

ANDERSEN, T. G., T. BOLLERSLEV, AND F. X. DIEBOLD (2010): "Parametric and Non-Parametric Volatility Measurement," in *Handbook of Financial Econometrics*, ed. by Y. Ait-Sahalia and L. P. Hansen, Elsevier, 67–137.

ANDREWS, D. (1991): "Heteroscedasticity and Autocorrelation Consistent Covaraince Matrix Estimation," *Econometrica*, 59, 817–858.

BANDI, F. AND J. RUSSELL (2006): "Separating Microstructure Noise from Volatility," *Journal of Financial Economics*, 79, 655–692.

——— (2008): "Microstructure Noise, Realized Variance, and Optimal Sampling," *Review of Economic Studies*, 75, 339–369.

BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, AND N. SHEPHARD (2008): "Designing Realized Kernels to Measure the ex post Variation of Equity Prices in the Presence of Noise," *Econometrica*, 76, 1481–1536.

BIBINGER, M., N. HAUTSCH, P. MALEC, AND M. REISS (2014): "Estimating the quadratic covariation matrix from noisy observations: local method of moments and efficiency," *Annals of Statistics*, 42, 1312–1364.

CHAKER, S. (2013): "Volatility and Liquidity Costs," Tech. Rep. 2013-29, Bank of Canada.

DA, R. AND D. XIU (2017): "When Moving Average Models Meet High-Frequency Data: Uniform Inference on Volatility," Working Paper, University of Chicago.

DIEBOLD, F. AND G. STRASSER (2013): "On the Correlation Structure of Microstructure Noise: A Financial Economic Approach," *Review of Economic Studies*, 80, 13041337.

GONCALVES, S. AND N. MEDDAHI (2009): "Bootstrapping Realized Volatility," *Econometrica*, 77, 283–306.

GOURIEROUX, C., A. MONFORT, AND A. TROGNON (1984): "Pseudo maximum likelihood methods: theory," *Econometrica*, 52, 681–700.

HAMILTON, J. D. (1994): "State Space Models," in *Handbook of Econometrics*, ed. by R. Engle and D. L. McFadden, Elsevier, vol. 4, 3039–3080.

HANSEN, L. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1054.

HANSEN, P. R. AND A. LUNDE (2006): "Realized Variance and Market Microstructure Noise," *Journal of Business and Economic Statistics*, 24, 127–161.

HARVEY, A. C. (1989): *Forecasting. Structural Time Series Models and the Kalman Filter.*, Cambridge University Press.

HASBROUCK, J. (1999): "The Dynamics of Discrete Bid and Ask Quotes," *Journal of Finance*, 54, 2109–2142.

HUANG, X. AND G. TAUCHEN (2005): "The Relative Contribution of Jumps to Total Price Variance," *Journal of Financial Econometrics*, 3, 456–499.

JACOD, J., Y. LI, P. MYKLAND, M. PODOLSKIJ, AND M. VETTER (2009): "Mirostructure noise in the continuous case: the pre-averaging approach," *Stochastic Processes and Their Applications*, 26, 2803–2831.

JACOD, J., Y. LI, AND X. ZHENG (2017): "Statistical Properties of Microstructure Noise," *Econometrica*, 85, 1133–1174.

KALMAN, R. (1960): "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering, Transactions of the ASME, Series D*, 82, 35–45.

——— (1963): "New Methods in Wiener Filtering Theory," in *Proceedings of the First Symposium of Engineering Applications of Random Function Theory and Probability*, ed. by J. Bogdanoff and F. Kozin, Wiley, 270–388.

KALNINA, I. AND O. LINTON (2008): "Estimating Quadratic Variation Consistently in the Presence of Endogenous and Diurnal Measurement Error," *Journal of Econometrics*, 147, 47–59.

KOKOSZKA, P. AND D. POLITIS (2011): "Non-linearity of ARCH and stochastic volatility models and Bartlett's formula," *Probability and Mathematical Statistics*, 31, 47–59.

KYLE, A. S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53 (6), 1315–1335.

LI, Y., S. XIE, AND X. ZHENG (2016): "Efficient Estimation of Integrated Volatility Incorporating Trading Information," Tech. rep., Hong Kong University of Science and Technology.

MYKLAND, P. AND L. ZHANG (2009): "Inference for Continuous Semimartingales Observed at High Frequency," *Econometrica*, 77, 1403–1445.

NEWEY, W. AND K. WEST (1987): "A Simple, Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covaraince Matrix," *Econometrica*, 55, 703–708.

ROLL, R. (1984): "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *Journal of Finance*, 1127–1140.

SCHLEMM, E. AND R. STELZER (2012): "Quasi maximum likelihood estimation for strongly mixing state space models and multivariate Levy-driven CARMA processes," *Electronic Journal of Statistics*, 6, 2185–2234.

SHEPPARD, K. (2013): "Measuring Market Speed," Tech. rep., University of Oxford.

USMANI, R. (1994): "Inversion of a tridiagonal Jacobi matrix," *Linear Algebra and its Applications*, 212-213, 413–414.

VIVES, X. (1995): "The Speed of Information Revelation in a Financial Market Mechanism," *Journal of Economic Theory*, 67, 187–204.

ZHANG, L., P. MYKLAND, AND Y. AIT-SAHALIA (2005): "A Tale of Two Time Scales: Determining Integrated Volatility With Noisy High-Frequency Data," *Journal of the American Statistical Association*, 100, 13941411.

ZHOU, B. (1996): "High-Frequency Data and Volatility in Foreign-Exchange Rates," *Journal of Business and Economic Statistics*, 45–52.

Appendix

Proofs

For general reading, the next lemma can be skipped and only serves as a technical support for the results in the next sections.

Lemma A. Let μ_0 be a fixed given initial value for the process $\{\mu_i\}_{i\geq 0}$.
The cross-moments of ε_i and μ_i obey, for $i \geq 1$,

(i) $\mathbb{E}[\mu_{i-h}\,\varepsilon_i] \;=\; 0 \qquad \forall h > 0,$

(ii) $\mathbb{E}[\mu_{i+h}\,\varepsilon_i] \;=\; (1-\alpha)^h\,\left(\sigma_\varepsilon^2 \,-\, \tilde{\gamma}\right) \;=\; (1-\alpha)^h\,\left(\gamma^2+\lambda-\gamma\right)\,\sigma_*^2 \qquad \forall h \geq 0.$

Proof. From equation (12), we have $\mu_{i+1} = (1-\alpha)\mu_i + \varepsilon_{i+1}^{\mu}$. Hence, backward recursion generates a moving average representation for μ_{i+h} and μ_i,

$$\mu_{i+h} = \;(1-\alpha)^{h+1}\mu_{i-1} \,+\, \sum_{j=0}^{h}(1-\alpha)^j\,\varepsilon_{i+h-j}^{\mu}, \qquad \forall i \geq 1,\, h \geq 0, \tag{43}$$

$$\mu_{i-h} = \;(1-\alpha)^{i-h}\mu_0 \,+\, \sum_{j=0}^{i-h-1}(1-\alpha)^j\,\varepsilon_{i-h-j}^{\mu}, \qquad \forall i > h \geq 0. \tag{44}$$

The innovation to the mispricing term is defined as $\varepsilon_i^{\mu} = \varepsilon_i - \varepsilon_i^*$, where $(\varepsilon_i\,, \varepsilon_i^*)$ is a bivariate i.i.d. random noise process with contemporaneous covariance $\tilde{\gamma} = \gamma\,\sigma_*^2$. Therefore,

$$\mathbb{E}[\varepsilon_i\,\varepsilon_k^{\mu}] \;=\; \delta(i-k)\,(\sigma_\varepsilon^2 \,-\, \tilde{\gamma}) \;=\; \delta(i-k)\,\left(\gamma^2 \,+\, \lambda \,-\, \gamma\right)\,\sigma_*^2\,, \tag{45}$$

where $\delta(\cdot)$ denotes the Dirac-Delta function. The lemma now follows straightforwardly.

(i) Equation (44) reveals that μ_{i-h} is a function of μ_0 and innovation terms that are realized prior to period i, so equation (45) implies that the covariance is zero.

(ii) From equation (43), we obtain

$$\mathbb{E}[\mu_{i+h}\,\varepsilon_i] = (1-\alpha)^{h+1}\,\underbrace{\mathbb{E}[\mu_{i-1}\varepsilon_i]}_{=0} + \sum_{j=0}^{h}(1-\alpha)^j\,\mathbb{E}[\varepsilon_i\varepsilon^{\mu}_{i+h-j}]. \tag{46}$$

The first term is zero because of property (i). For the second term, equation (45) implies,

$$\sum_{j=0}^{h}(1-\alpha)^j\,\mathbb{E}[\varepsilon_i\,\varepsilon^{\mu}_{i+h-j}] \;=\; (1-\alpha)^h\left(\sigma^2_\varepsilon \;-\; \tilde{\gamma}\right). \tag{47}$$

□

Proof of Lemma 2. We have, by definition,

$$\mathbb{C}ov(r_{i+h}, r_i) = \mathbb{E}[(r_{i+h}-\underbrace{\mathbb{E}[r_{i+h}]}_{=0})(r_i-\underbrace{\mathbb{E}[r_i]}_{=0})] = \mathbb{E}[r_{i+h}r_i] = \mathbb{E}[(p_{i+h}-p_{i+h-1})(p_i-p_{i-1})]. \tag{48}$$

Now, using, $p_i = \mu_i + p^*_i$, we obtain,

$$\mathbb{C}ov(r_{i+h}, r_i) = \mathbb{E}[(\mu_{i+h} - \mu_{i+h-1} + p^*_{i+h} - p^*_{i+h-1})(\mu_i - \mu_{i-1} + p^*_i - p^*_{i-1})]. \tag{49}$$

Then exploit equation (11) to establish,

$$\begin{aligned}
\mathbb{C}ov(r_{i+h}, r_i) &= \mathbb{E}\Big[\big(-\alpha\mu_{i+h-1} + \underbrace{\varepsilon^{\mu}_{i+h} + \varepsilon^*_{i+h}}_{=\varepsilon_{i+h}}\big)\big(-\alpha\mu_{i-1} + \underbrace{\varepsilon^{\mu}_i + \varepsilon^*_i}_{=\varepsilon_i}\big)\Big] \\
&= \mathbb{E}\Big[\big(-\alpha\mu_{i+h-1} + \varepsilon_{i+h}\big)\big(-\alpha\mu_{i-1} + \varepsilon_i\big)\Big] \\
&= \alpha^2\,\mathbb{E}[\mu_{i+h-1}\mu_{i-1}] - \alpha\mathbb{E}[\mu_{i+h-1}\varepsilon_i] - \alpha\underbrace{\mathbb{E}[\varepsilon_{i+h}\mu_{i-1}]}_{=0} + \underbrace{\mathbb{E}[\varepsilon_{i+h}\varepsilon_i]}_{=0},
\end{aligned} \tag{50}$$

where we exploited Lemma A, and the fact that ε_i is i.i.d.

It follows that,

$$\mathbb{C}ov(r_{i+h}, r_i) = \alpha^2\,\mathbb{E}[\mu_{i+h-1}\mu_{i-1}] - \alpha \underbrace{\mathbb{E}[\mu_{i+h-1}\varepsilon_i]}_{=(1-\alpha)^{h-1}\left(\sigma^2_\varepsilon-\gamma\right)}. \tag{51}$$

The first term is given in equation (43). The two terms may then be evaluated using equation (13) and Lemma A to yield,

$$\begin{aligned}\mathbb{C}ov(r_{i+h}, r_i) &= (\sigma^2_{\varepsilon^*} + \sigma^2_{\varepsilon} - 2\gamma)\alpha^2 \frac{(1-\alpha)^h}{1-(1-\alpha)^2} - \alpha(1-\alpha)^{h-1}\left(\sigma^2_{\epsilon} - \gamma\right) \\ &= \underbrace{\alpha\frac{(1-\alpha)^{h-1}}{(2-\alpha)}}_{=\psi(h-1)} \sigma^2_{\varepsilon^*}\left((1+\lambda-2\tilde{\gamma})(1-\alpha) - (2-\alpha)(\lambda-\tilde{\gamma})\right) \\ &= \psi(h-1)\sigma^2_{\varepsilon^*}\left(1-\alpha(1-\tilde{\gamma})-\lambda\right)\end{aligned} \tag{52}$$

with $\tilde{\gamma} = \frac{\gamma}{\sigma^2_{\varepsilon^*}}$ and $\lambda = \frac{\sigma^2_{\epsilon}}{\sigma^2_{\varepsilon^*}}$.

□

Kalman Filter and Derivatives

Based on the state space formulation provided in Section 4.2, $\partial\ell(Y,\theta)/\partial\theta'$ and $\partial^2\ell(Y,\theta)/\partial\theta\partial\theta'$ are given by

$$\partial\ell(Y,\theta)/\partial\theta' = -\frac{1}{2}\sum_{i=1}^{T}\left\{(s_i^{-2}(1-\nu_i^2 s_i^{-2}))Ds_i^2 + 2\nu_i s_i^{-2} D\nu_i\right\} \tag{53}$$

$$\begin{aligned}\partial^2\ell(Y,\theta)/\partial\theta\partial\theta' = -\frac{1}{2}\sum_{i=1}^{T}\big\{&-(Ds_i^2)'((1-\nu_i^2/s_i^2)s_i^{-4})Ds_i^2 - 2(Ds_i^2)'s_i^{-4}\nu_i D\nu_i \\ &+(Ds_i^2)'s_i^{-6}\nu_i^2 Ds_i^2 + 2(\nu_i s_i^{-2}\otimes I_3)\partial^2\nu_i/\partial\theta\partial\theta' \\ &-2(\nu_i s_i^{-2}\otimes(D\nu_i)'s_i^{-2})Ds_i^2 + 2(D\nu_i)'s_i^{-2}D\nu_i\big\},\end{aligned} \tag{54}$$

yielding

$$\mathbb{E}[\partial^2\ell(Y,\theta)/\partial\theta\partial\theta'] = -\mathbb{E}[2(D\nu_i)'s_i^{-2}D\nu_i - (Ds_i^2)'s_i^{-4}Ds_i^2]. \tag{55}$$

Up to initializations Ds_i^2 and $D\nu_i^2$ can then be computed based on the Kalman filter recursions (e.g., Kalman (1960), Kalman (1963))

$$\nu_i = r_i - FX_{i|i-1} \tag{56}$$

$$X_{i|i-1} = GX_{i-1|i-2} + K_{i-1}\nu_{i-1} \tag{57}$$

$$K_i = GP_iF's_i^{-2} \tag{58}$$

$$P_i = GP_{i-1}L_{i-1}' + \Sigma_w \tag{59}$$

$$s_i^2 = FP_iF' \tag{60}$$

$$L_i = G - K_iF \tag{61}$$

with derivatives

$$D\nu_i = -X_{i|i-1}DF - FDX_{i|i-1} \tag{62}$$

$$Ds_i^2 = (F \otimes F)DP_i + 2(FP_i)DF \tag{63}$$

$$\begin{aligned} DK_i = {} & (s_i^{-2}FP_i)DG + (s_i^{-2}F \otimes G)DP_i \\ & + (s_i^{-2} \otimes GP_i)DF - (s_i^{-2} \otimes GP_iF's_i^{-2})Ds_i^2 \end{aligned} \tag{64}$$

$$DX_{i|i-1} = GDX_{i-1|i-2} + (X_{i-1|i-2} \otimes I_3)DG + K_{i-1}D\nu_{i-1} + (\nu_{i-1} \otimes I_3)DK_{i-1} \tag{65}$$

$$DL_i = DG - (F' \otimes I_3)DK_i - (I_3 \otimes K_i)DF \tag{66}$$

$$\begin{aligned} DP_i = {} & (L_{i-1} \otimes G)DP_{i-1} + (L_{i-1}P_{i-1} \otimes I_3)DG \\ & + (I_3 \otimes GP_{i-1})\mathcal{K}_{33}DL_{i-1} + 2N_3(\Sigma_w^{1/2} \otimes I_3)D\Sigma_w^{1/2}, \end{aligned} \tag{67}$$

where $\Sigma_w = \Sigma_w^{1/2}(\Sigma_w^{1/2})'$. Moreover, $N_m = \frac{1}{2}(I_{m^2} + \mathcal{K}_{mm})$ and $\mathcal{K}_{mn}$ is defined as commutation matrix satisfying $\mathcal{K}_{mn}vec(C) = vec(C')$ with C denoting a $m \times n$ matrix.

Tables and Figures

Parameter	Mean	Std.	Skew.	Ex.Kurt.	$Q_{0.05}$	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	$Q_{0.95}$
Liquidity group 1 (20 most liquid assets from NASDAQ 100)									
α	0.902	0.095	-1.501	12.637	0.772	0.852	0.904	0.957	1.040
σ_*^2 ($\times 10^{-8}$)	2.332	4.426	6.408	58.953	0.204	0.520	1.052	2.286	8.385
$F_\alpha(\gamma, \lambda)$	-0.009	0.138	3.266	31.820	-0.178	-0.081	-0.018	0.049	0.166
corr(r_{t+1}, r_t)	0.014	0.088	-0.201	1.470	-0.133	-0.042	0.015	0.070	0.156
γ_{max}	0.988	0.107	0.930	6.658	0.832	0.925	0.984	1.046	1.151
λ_{max}	0.295	0.120	8.459	100.714	0.201	0.248	0.282	0.320	0.391
Liquidity group 2									
α	0.907	0.101	-1.321	12.594	0.769	0.856	0.910	0.964	1.052
σ_*^2 ($\times 10^{-8}$)	2.163	4.598	6.834	64.933	0.156	0.404	0.847	2.015	8.018
$F_\alpha(\gamma, \lambda)$	0.002	0.154	3.424	29.002	-0.180	-0.074	-0.009	0.059	0.190
corr(r_{t+1}, r_t)	0.006	0.093	-0.273	2.552	-0.148	-0.050	0.007	0.064	0.157
γ_{max}	0.997	0.116	1.094	7.302	0.830	0.931	0.992	1.055	1.172
λ_{max}	0.304	0.139	7.604	76.942	0.200	0.252	0.287	0.326	0.406
Liquidity group 3									
α	0.915	0.098	-0.883	12.493	0.777	0.865	0.918	0.969	1.055
σ_*^2 ($\times 10^{-8}$)	1.388	3.459	10.257	141.462	0.151	0.344	0.618	1.230	4.359
$F_\alpha(\gamma, \lambda)$	0.012	0.152	3.691	31.522	-0.166	-0.063	0.000	0.066	0.193
corr(r_{t+1}, r_t)	-0.001	0.092	-0.382	4.229	-0.150	-0.056	-0.000	0.055	0.145
γ_{max}	1.006	0.114	1.347	9.665	0.844	0.941	1.000	1.061	1.174
λ_{max}	0.309	0.138	7.540	75.339	0.206	0.257	0.292	0.330	0.408
Liquidity group 4									
α	0.920	0.100	-1.067	11.113	0.775	0.871	0.921	0.975	1.063
σ_*^2 ($\times 10^{-8}$)	1.251	3.423	11.224	164.292	0.122	0.294	0.547	1.097	3.735
$F_\alpha(\gamma, \lambda)$	0.023	0.166	3.712	28.051	-0.164	-0.055	0.005	0.075	0.214
corr(r_{t+1}, r_t)	-0.008	0.095	-0.236	3.108	-0.163	-0.063	-0.004	0.047	0.144
γ_{max}	1.015	0.120	1.301	7.749	0.845	0.949	1.005	1.069	1.193
λ_{max}	0.318	0.157	6.952	61.399	0.207	0.262	0.295	0.336	0.421
Liquidity group 5 (20 least liquid assets from NASDAQ 100)									
α	0.923	0.092	-0.869	12.505	0.784	0.902	0.920	0.960	1.061
σ_*^2 ($\times 10^{-8}$)	1.477	3.485	9.652	129.026	0.122	0.339	0.668	1.352	4.866
$F_\alpha(\gamma, \lambda)$	0.029	0.168	4.445	34.487	-0.146	-0.017	0.000	0.058	0.215
corr(r_{t+1}, r_t)	-0.012	0.090	-0.243	6.199	-0.165	-0.048	-0.000	0.015	0.130
γ_{max}	1.019	0.117	1.799	11.659	0.862	0.984	1.000	1.053	1.193
λ_{max}	0.321	0.164	7.055	61.352	0.214	0.282	0.292	0.325	0.421

Table 7: The table reports summary statistics for the model parameter estimates (estimated under restriction $\gamma = 0$) obtained from T = 10 min local intervals with returns sampled at frequency Δ= 2 sec. The descriptive statistics are based on the estimates from all intra-daily intervals over the first 61 trading days of 2014. The results are reported separately for five groups of stocks sorted by the average daily mid-quote revisions. To prevent extreme estimates significantly distort the results, local intervals where parameter estimates of σ_*^2 and λ_{max} exceed their 99.5th percentiles are excluded from the sample.

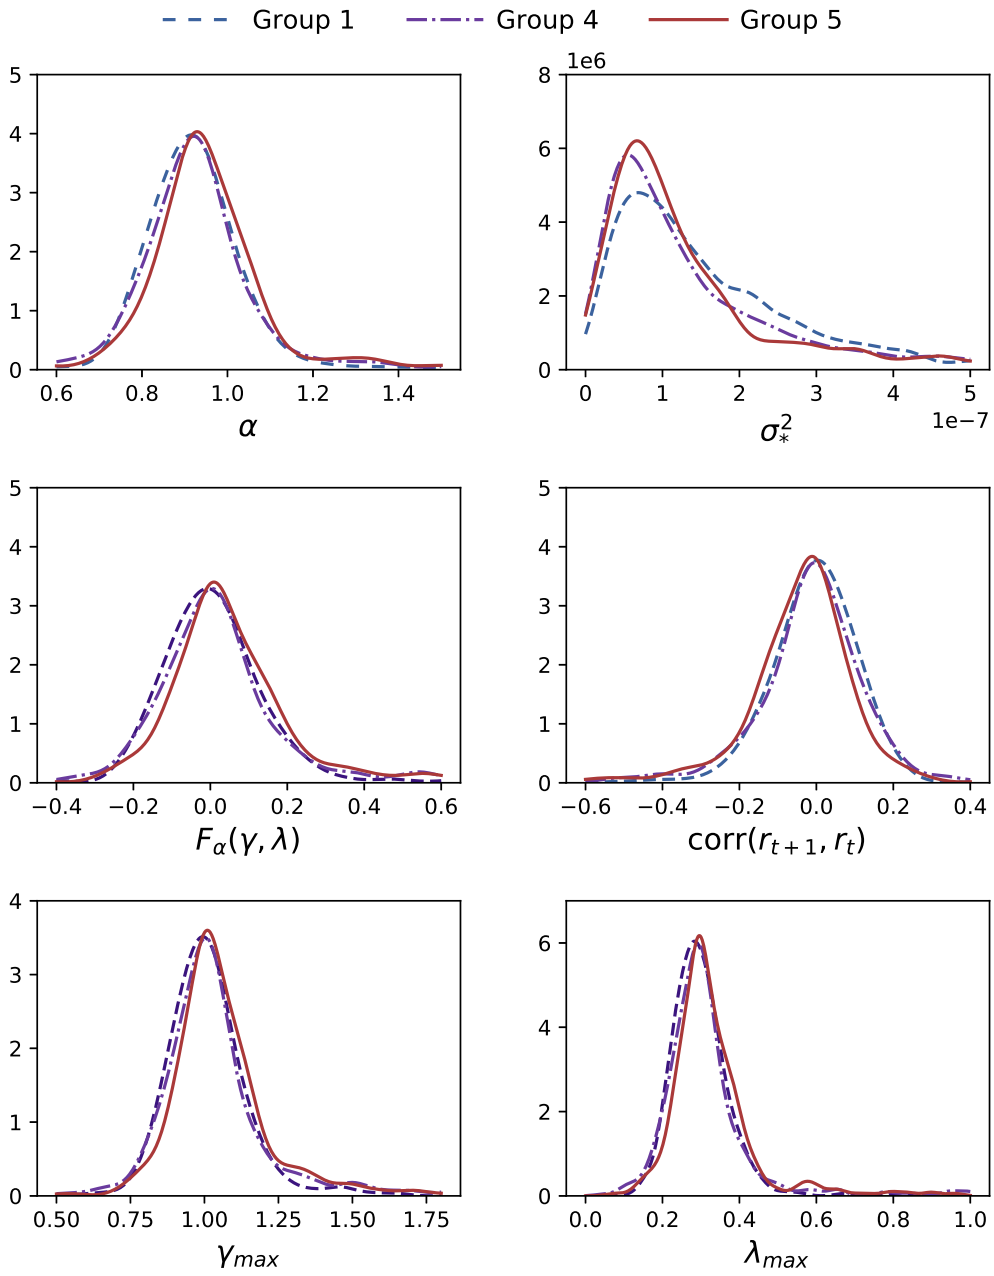

Figure 8: Empirical distributions of the model parameter estimates (estimated under restriction $\gamma = 0$) obtained for T = 10 min local intervals between 9:30 and 9:40 ET (first ten minutes after market opening) with returns sampled at frequency Δ= 2 sec. Resulted densities are shown for the stocks from 1st, 4th and 5th groups of stocks sorted by the average number of daily mid-quote revisions (with 20 stocks in each group) for the first 61 trading days in 2014. In total, each depicted density curve is constructed with parameter estimates from 1,220 local intervals.

Parameter	Mean	Std.	Skew.	Ex.Kurt.	$Q_{0.05}$	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	$Q_{0.95}$
Liquidity group 1 (20 most liquid assets from NASDAQ 100)									
α	0.917	0.133	-0.135	10.946	0.765	0.853	0.914	0.975	1.093
σ_*^2 ($\times 10^{-8}$)	16.687	13.445	1.527	2.385	2.765	6.969	12.481	22.281	43.548
$F_\alpha(\gamma, \lambda)$	0.023	0.196	3.507	26.337	-0.181	-0.077	0.001	0.078	0.258
corr(r_{t+1}, r_t)	-0.007	0.115	-1.317	6.587	-0.190	-0.065	-0.001	0.067	0.158
γ_{max}	1.013	0.145	1.731	9.534	0.828	0.928	1.001	1.072	1.233
λ_{max}	0.320	0.176	6.054	47.309	0.201	0.251	0.293	0.338	0.449
Liquidity group 2									
α	0.925	0.170	1.436	10.860	0.744	0.848	0.915	0.982	1.122
σ_*^2 ($\times 10^{-8}$)	18.774	15.512	1.403	1.484	3.227	7.424	13.626	24.893	51.677
$F_\alpha(\gamma, \lambda)$	0.042	0.253	3.186	15.924	-0.197	-0.079	0.003	0.090	0.445
corr(r_{t+1}, r_t)	-0.016	0.152	-2.075	9.814	-0.226	-0.075	-0.002	0.069	0.174
γ_{max}	1.027	0.188	2.260	10.127	0.813	0.926	1.002	1.082	1.358
λ_{max}	0.338	0.230	4.825	27.025	0.191	0.249	0.293	0.343	0.589
Liquidity group 3									
α	0.942	0.185	2.477	11.235	0.758	0.857	0.919	0.986	1.249
σ_*^2 ($\times 10^{-8}$)	15.742	14.464	1.790	2.849	2.798	6.149	10.733	18.949	50.559
$F_\alpha(\gamma, \lambda)$	0.059	0.293	3.739	18.339	-0.181	-0.070	0.003	0.091	0.556
corr(r_{t+1}, r_t)	-0.028	0.176	-2.399	9.867	-0.287	-0.075	-0.003	0.061	0.163
γ_{max}	1.041	0.220	2.924	12.721	0.826	0.935	1.003	1.083	1.395
λ_{max}	0.347	0.251	4.660	24.831	0.197	0.253	0.293	0.344	0.690
Liquidity group 4									
α	0.931	0.182	1.550	10.277	0.731	0.854	0.919	0.979	1.187
σ_*^2 ($\times 10^{-8}$)	15.899	15.410	1.726	2.494	2.198	5.551	10.206	20.353	53.455
$F_\alpha(\gamma, \lambda)$	0.063	0.296	3.000	12.367	-0.205	-0.073	0.006	0.090	0.555
corr(r_{t+1}, r_t)	-0.024	0.167	-1.882	8.702	-0.275	-0.075	-0.005	0.063	0.185
γ_{max}	1.040	0.212	2.079	8.168	0.804	0.932	1.006	1.082	1.475
λ_{max}	0.357	0.275	4.258	19.816	0.184	0.251	0.296	0.343	0.793
Liquidity group 5 (20 least liquid assets from NASDAQ 100)									
α	0.958	0.171	1.864	9.320	0.758	0.883	0.935	1.012	1.257
σ_*^2 ($\times 10^{-8}$)	14.704	14.093	1.921	3.450	2.004	5.789	9.805	17.316	47.607
$F_\alpha(\gamma, \lambda)$	0.081	0.252	3.488	19.957	-0.167	-0.035	0.030	0.128	0.531
corr(r_{t+1}, r_t)	-0.048	0.157	-1.947	8.429	-0.314	-0.105	-0.026	0.030	0.147
γ_{max}	1.062	0.192	2.444	11.760	0.843	0.967	1.028	1.117	1.399
λ_{max}	0.359	0.226	4.682	27.520	0.205	0.271	0.309	0.366	0.656

Table 8: The table reports summary statistics for the model parameter estimates (estimated under restriction $\gamma = 0$) obtained from T = 10 min local intervals with returns sampled at frequency Δ= 2 sec. The descriptive statistics are based on the estimates from intra-daily intervals between 9:30 and 9:40 ET (first ten minutes after market opening) over the first 61 trading days of 2014. The results are reported separately for five groups of stocks sorted by the average daily mid-quote revisions. To prevent extreme estimates significantly distort the results, local intervals where parameter estimates of σ_*^2 and λ_{max} exceed their 99.5th percentiles (with respect to the entire sample) are excluded.

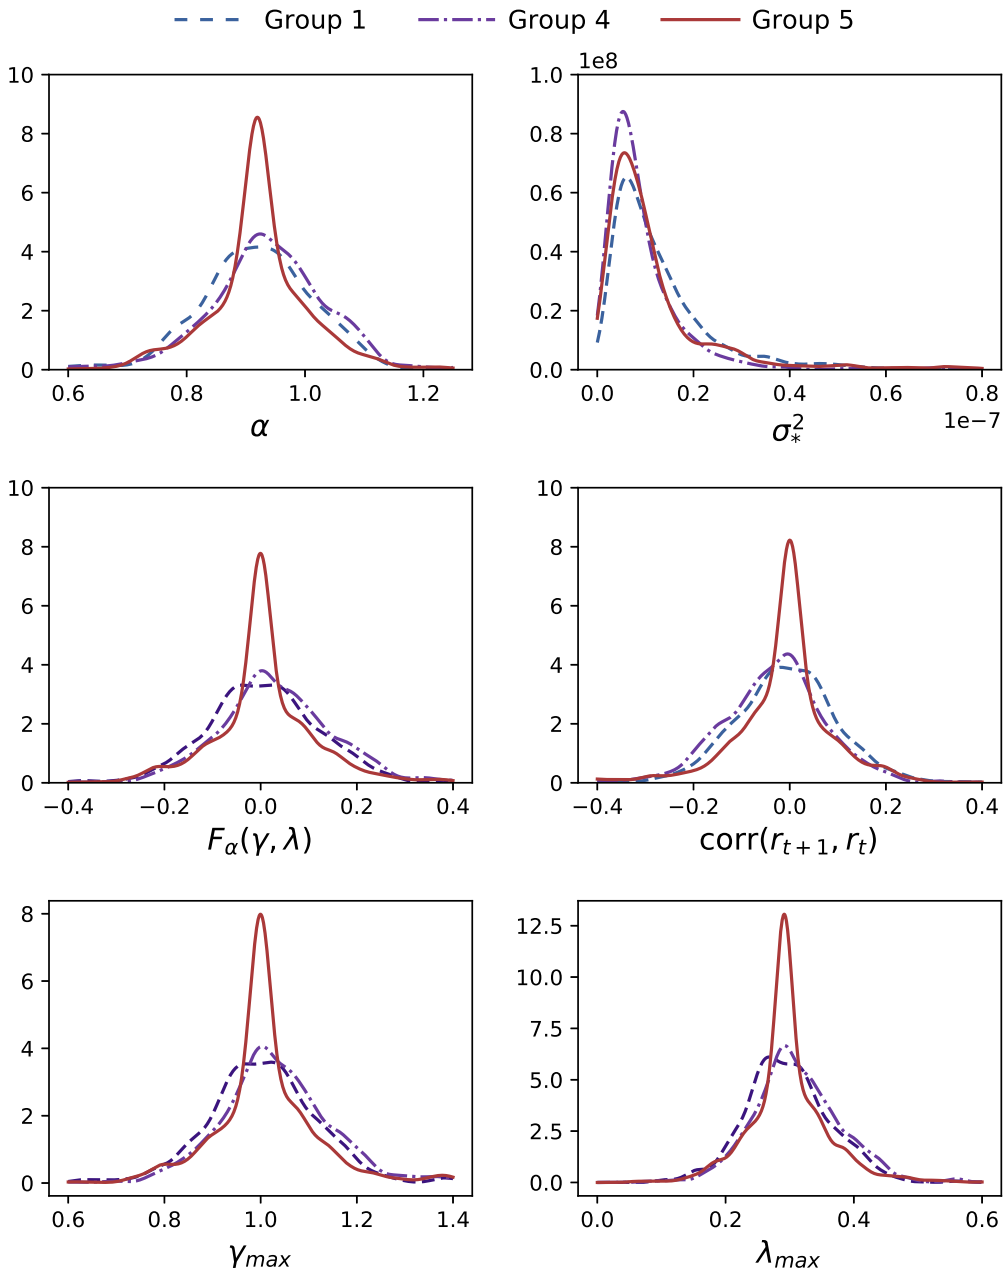

Figure 9: Empirical distributions of the model parameter estimates (estimated under restriction $\gamma = 0$) obtained for T = 10 min local intervals between 15:50 and 16:00 ET (last ten minutes of trading) with returns sampled at frequency Δ= 2 sec. Resulted densities are shown for the stocks from 1st, 4th and 5th groups of stocks sorted by the average number of daily mid-quote revisions (with 20 stocks in each group) for the first 61 trading days in 2014. In total, each depicted density curve is constructed with parameter estimates from 1,220 local intervals.

Parameter	Mean	Std.	Skew.	Ex.Kurt.	$Q_{0.05}$	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	$Q_{0.95}$
Liquidity group 1 (20 most liquid assets from NASDAQ 100)									
α	0.912	0.130	-1.010	11.286	0.752	0.856	0.919	0.980	1.074
σ_*^2 ($\times 10^{-8}$)	1.415	1.562	5.035	41.219	0.310	0.578	0.996	1.653	3.748
$F_\alpha(\gamma, \lambda)$	0.038	0.239	3.654	20.261	-0.184	-0.067	0.007	0.088	0.256
corr(r_{t+1}, r_t)	-0.009	0.111	-0.681	5.468	-0.178	-0.073	-0.006	0.058	0.162
γ_{max}	1.021	0.155	1.735	8.910	0.827	0.938	1.006	1.081	1.232
λ_{max}	0.338	0.241	5.471	33.308	0.198	0.256	0.297	0.344	0.446
Liquidity group 2									
α	0.922	0.123	-1.455	9.684	0.753	0.863	0.927	0.987	1.085
σ_*^2 ($\times 10^{-8}$)	1.608	2.326	6.304	64.191	0.238	0.496	0.896	1.755	5.318
$F_\alpha(\gamma, \lambda)$	0.038	0.206	3.065	19.889	-0.178	-0.062	0.020	0.095	0.264
corr(r_{t+1}, r_t)	-0.015	0.107	0.159	1.584	-0.188	-0.079	-0.016	0.053	0.157
γ_{max}	1.025	0.142	0.738	5.780	0.833	0.942	1.018	1.087	1.241
λ_{max}	0.332	0.199	5.842	40.306	0.200	0.258	0.303	0.348	0.447
Liquidity group 3									
α	0.939	0.116	-0.819	6.782	0.778	0.884	0.944	1.005	1.097
σ_*^2 ($\times 10^{-8}$)	1.045	1.309	7.117	76.149	0.232	0.483	0.743	1.171	2.625
$F_\alpha(\gamma, \lambda)$	0.072	0.217	3.495	18.523	-0.134	-0.033	0.039	0.121	0.338
corr(r_{t+1}, r_t)	-0.037	0.103	-0.334	1.985	-0.206	-0.101	-0.033	0.029	0.117
γ_{max}	1.052	0.142	1.460	5.548	0.874	0.969	1.036	1.111	1.286
λ_{max}	0.357	0.226	5.167	30.733	0.220	0.273	0.315	0.362	0.558
Liquidity group 4									
α	0.932	0.121	-0.565	6.768	0.753	0.878	0.935	0.998	1.093
σ_*^2 ($\times 10^{-8}$)	0.977	1.147	7.388	85.064	0.226	0.453	0.693	1.146	2.411
$F_\alpha(\gamma, \lambda)$	0.071	0.234	3.255	15.092	-0.156	-0.037	0.029	0.117	0.367
corr(r_{t+1}, r_t)	-0.032	0.110	-0.291	3.274	-0.200	-0.097	-0.024	0.032	0.142
γ_{max}	1.048	0.152	1.401	4.809	0.851	0.966	1.026	1.107	1.317
λ_{max}	0.360	0.244	4.728	24.834	0.207	0.272	0.309	0.362	0.565
Liquidity group 5 (20 least liquid assets from NASDAQ 100)									
α	0.922	0.114	-0.487	10.431	0.756	0.890	0.920	0.963	1.078
σ_*^2 ($\times 10^{-8}$)	1.237	1.468	4.286	28.755	0.209	0.488	0.812	1.350	3.586
$F_\alpha(\gamma, \lambda)$	0.035	0.211	4.275	29.208	-0.173	-0.028	0.000	0.065	0.269
corr(r_{t+1}, r_t)	-0.012	0.105	-0.678	4.006	-0.192	-0.055	-0.000	0.024	0.151
γ_{max}	1.022	0.142	1.876	9.733	0.837	0.974	1.000	1.060	1.240
λ_{max}	0.329	0.205	6.253	45.569	0.203	0.276	0.292	0.328	0.455

Table 9: The table reports summary statistics for the model parameter estimates (estimated under restriction $\gamma = 0$) obtained from T = 10 min local intervals with returns sampled at frequency Δ= 2 sec. The descriptive statistics are based on the estimates from intra-daily intervals between 15:50 and 16:00 ET (last ten minutes of trading) over the first 61 trading days of 2014. The results are reported separately for five groups of stocks sorted by the average daily mid-quote revisions. To prevent extreme estimates significantly distort the results, local intervals where parameter estimates of σ_*^2 and λ_{max} exceed their 99.5th percentiles (with respect to the entire sample) are excluded.

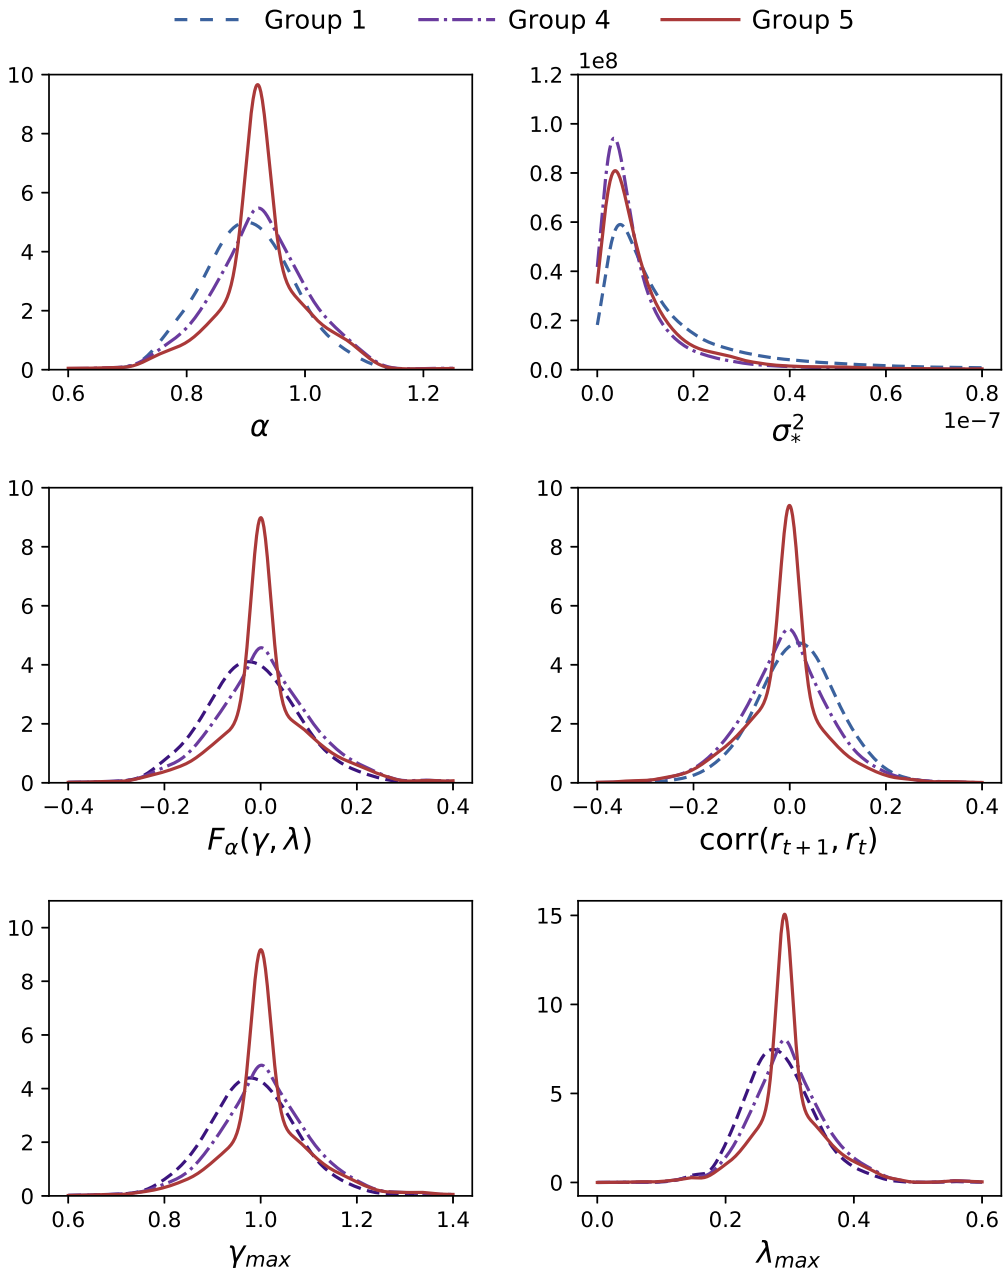

Figure 10: Empirical distributions of the model parameter estimates (estimated under restriction $\gamma = 0$) obtained for T = 10 min local intervals between 9:50 and 15:40 ET (excluding first and last 20 minutes of a trading session) with intra-interval sampling at Δ= 2 sec frequency. Resulted densities are shown for the stocks from 1st, 4th and 5th groups of stocks sorted by the average number of daily mid-quote revisions (with 20 stocks in each group) for the first 61 trading days in 2014. In total, each depicted density curve is constructed with parameter estimates from 42,700 local intervals.

Parameter	Mean	Std.	Skew.	Ex.Kurt.	$Q_{0.05}$	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	$Q_{0.95}$
Liquidity group 1 (20 most liquid assets from NASDAQ 100)									
α	0.901	0.093	-1.665	12.387	0.773	0.851	0.903	0.955	1.037
σ_*^2 ($\times 10^{-8}$)	1.845	2.842	6.410	72.207	0.197	0.499	0.993	2.054	6.148
$F_\alpha(\gamma, \lambda)$	-0.012	0.132	2.982	29.679	-0.179	-0.082	-0.019	0.047	0.160
corr(r_{t+1}, r_t)	0.015	0.087	-0.081	0.598	-0.129	-0.040	0.016	0.071	0.156
γ_{max}	0.986	0.104	0.748	5.383	0.831	0.924	0.982	1.043	1.146
λ_{max}	0.293	0.113	8.596	107.238	0.200	0.248	0.281	0.319	0.388
Liquidity group 2									
α	0.906	0.098	-1.637	11.713	0.770	0.856	0.909	0.962	1.049
σ_*^2 ($\times 10^{-8}$)	1.679	2.888	6.504	71.626	0.152	0.388	0.793	1.781	5.965
$F_\alpha(\gamma, \lambda)$	-0.000	0.148	3.273	28.591	-0.180	-0.075	-0.010	0.058	0.184
corr(r_{t+1}, r_t)	0.008	0.091	-0.064	0.781	-0.144	-0.048	0.008	0.065	0.157
γ_{max}	0.995	0.112	0.899	5.841	0.830	0.930	0.991	1.053	1.166
λ_{max}	0.302	0.133	7.783	82.328	0.200	0.251	0.286	0.325	0.403
Liquidity group 3									
α	0.914	0.094	-1.419	10.254	0.777	0.865	0.917	0.968	1.052
σ_*^2 ($\times 10^{-8}$)	1.011	1.724	13.343	337.298	0.147	0.332	0.587	1.108	3.082
$F_\alpha(\gamma, \lambda)$	0.009	0.144	3.347	28.886	-0.167	-0.064	-0.001	0.064	0.187
corr(r_{t+1}, r_t)	0.000	0.089	0.013	0.982	-0.147	-0.054	0.000	0.055	0.146
γ_{max}	1.003	0.109	0.911	5.824	0.843	0.940	1.000	1.059	1.169
λ_{max}	0.306	0.130	7.755	81.795	0.206	0.257	0.292	0.329	0.404
Liquidity group 4									
α	0.919	0.097	-1.410	9.660	0.776	0.871	0.920	0.974	1.061
σ_*^2 ($\times 10^{-8}$)	0.885	1.498	14.685	438.713	0.116	0.282	0.515	0.985	2.710
$F_\alpha(\gamma, \lambda)$	0.021	0.160	3.670	28.733	-0.163	-0.055	0.004	0.074	0.208
corr(r_{t+1}, r_t)	-0.007	0.092	0.002	1.117	-0.160	-0.062	-0.004	0.047	0.143
γ_{max}	1.013	0.116	1.129	6.475	0.846	0.949	1.004	1.068	1.188
λ_{max}	0.316	0.151	7.177	66.409	0.208	0.262	0.295	0.335	0.418
Liquidity group 5 (20 least liquid assets from NASDAQ 100)									
α	0.922	0.089	-1.332	10.733	0.786	0.903	0.920	0.958	1.059
σ_*^2 ($\times 10^{-8}$)	1.109	1.916	10.889	226.273	0.118	0.324	0.625	1.222	3.361
$F_\alpha(\gamma, \lambda)$	0.027	0.163	4.443	34.739	-0.145	-0.015	0.000	0.055	0.209
corr(r_{t+1}, r_t)	-0.011	0.087	0.061	4.449	-0.162	-0.046	-0.000	0.013	0.129
γ_{max}	1.018	0.114	1.641	10.345	0.863	0.986	1.000	1.051	1.189
λ_{max}	0.320	0.161	7.187	63.698	0.215	0.283	0.292	0.324	0.421

Table 10: The table reports summary statistics for the model parameter estimates (estimated under restriction $\gamma = 0$) obtained from T = 10 min local intervals with returns sampled at frequency Δ= 2 sec. The descriptive statistics are based on the estimates from intra-daily intervals between 9:50 and 15:40 ET (excluding first and last 20 minutes of a trading session) over the first 61 trading days of 2014. The results are reported separately for five groups of stocks sorted by the average daily mid-quote revisions. To prevent extreme estimates significantly distort the results, local intervals where parameter estimates of σ_*^2 and λ_{max} exceed their 99.5th percentiles (with respect to the entire sample) are excluded.