Research Statement

With the introduction of electronic trading, security prices are quoted and traded at greater speed and precision. This opens possibilities for asking new and more complex empirical questions. At the same time, inference using such tick-by-tick data becomes more susceptible to the impact of market microstructure. For both these reasons, new econometric methodology is called for.

My primary research interest is the econometrics of high frequency financial data. I work with tick-by-tick financial data for the purpose of developing sound statistical methodology, to help better understand the data process underlying the markets, and hence improve regulatory monitoring and risk management.

My early work on volatility estimation is the first to provide consistent and rate efficient volatility estimator, in the presence of market microstructure [7, 10]. I also studied how to relate multiple security price processes with the purpose of reducing hedging error [8].

More recent work is as follows. I have

(a) extended the findings on volatility estimation to more general settings ([2, 3, 4]; Sections C.4-C.5), including a new and highly general central limit theory for high frequency volatility and related quantities;

(b) designed and analyzed a covariance estimator for non-synchronously traded securities, again in the presence of microstructure noise ([1]; Section C.1). This is the first paper to find a consistent estimator of covariance in this setting;

(c) introduced empirical model selection to high frequency data, which has direct application to the hedging of securities, including both options trading and making portfolios neutral with respect to the market, or volatility ([6, 13]; Section C.2);

(d) developed a local likelihood approach for the purpose of analyzing continuous semi-martingales ([5]; Section C.3). This provides a general framework for creating and analyzing estimators in non-transparent situations. It also helps improving efficiency of existing estimators;

(e) studied the behavior of price jumps, and endogeneity and irregularity in trade times ([3,15,16]; Section C.4), leading to a greater understanding and better analysis of the causes of heavy tailed returns;

(f) further explored volatility forecasting ([14]; Section C.5).

Overall, Section A provides the setting of my research, Section B summarizes my early work. All recent findings are presented in Section C.

A Setting

The price $\{S_t\}$ of a financial instrument is usually thought of as additive on the logarithmic scale. Specifically, if $X_t = \log S_t$,

$$dX_t = \mu_t dt + \sigma_t dW_t, \tag{1}$$

where W is a standard Brownian motion, and σ_t (or σ_t^2) is the *volatility* of $\{S_t\}$. Model (1) is typically referred to as the continuous-time semi-martingale model for asset prices. Since the coefficients (μ_t, σ_t) can be random in various manners, (1) covers a wide range of modeling schemes for asset prices and at the same time it satisfies the "no arbitrage law" from continuous-time financial theory.

Volatility is arguably the most central concept in finance. First, volatility tells an investor how risky a security is, through, say, the Sharpe ratio. Also, it gives money managers a tool to select portfolios and assess their performance. Second, volatility enters in the Black-Scholes option pricing formula, which still serves as a benchmark in options trading practice; it also helps setting the hedging strategy. Third, central banks and risk monitoring agencies use volatility-related measures to control the risk that a financial institution bears, through value-at-risk and similar quantities.

Precise measurement of volatility is therefore of great importance. In this respect, high frequency data have created a revolution. In principle, it is now possible to estimate daily (or intra-day) volatility with a fair amount of precision, because according to results in probability theory, the sum of squared returns approximates the integrated volatility over the same time period. One can write this as follows: if a trading day starts from time zero and ends at time T, if the process X_t is observed (through, say, transaction prices) at a large number of times $0 < t_1 < ... < t_n = T$, then

$$\sum_{i=1}^{n}(X_{t_{i+1}} - X_{t_i})^2 \approx \int_0^T \sigma_t^2 dt \qquad (2)$$

The left hand side is usually called *realized volatility* (or *realized variance*). Other time periods than a day can also be used.

The high frequency approach complements the two earlier methods used by the financial industry: calibration (using options implied volatility) and econometric time series methods. These earlier methods typically assume rather specific models or functional forms for σ_t, while this is not necessary when exploiting the relationship (2).

B Early Work

B.1 Market Microstructure

As far as volatility estimation with high frequency data is concerned, a common practice is to use the realized volatility $\sum_{i=1}^{n}(X_{t_{i+1}} - X_{t_i})^2$. In practice, however, this estimator runs into the problem of market microstructure. When a transaction happens at time t_i, the observation $\log S_{t_i}$ is not the X_{t_i} given by (1), but rather

$$\log S_{t_i} = X_{t_i} + \epsilon_i, \qquad (3)$$

where X_t follows (1) and ϵ_i is a contamination or error representing market microstructure. Market microstructure could exist in the form of bid-ask spread, discreteness of price change, gradual response of prices to a block trade, and many others. Just as its name suggests, market microstructure has a more pronounced role when the data is observed more finely.

The microstructure noise must be of small magnitude, otherwise it would give rise to obvious arbitrage opportunities. Despite its small size, the effect of the noise will swamp the signal when it comes to estimating volatility. An illustration of this is given in Figure 1.

In reaction, a standard approach in the finance literature is to construct a volatility estimator based on a sparse subset of the data, even if a data series of finer sampling frequency is at hand. In other words, the usual financial practice throws away most of the available data. It is argued in my research that this customary way of estimating volatility is suboptimal and arbitrary. My approach is to incorporate the microstructure noise explicitly into the estimation procedure.

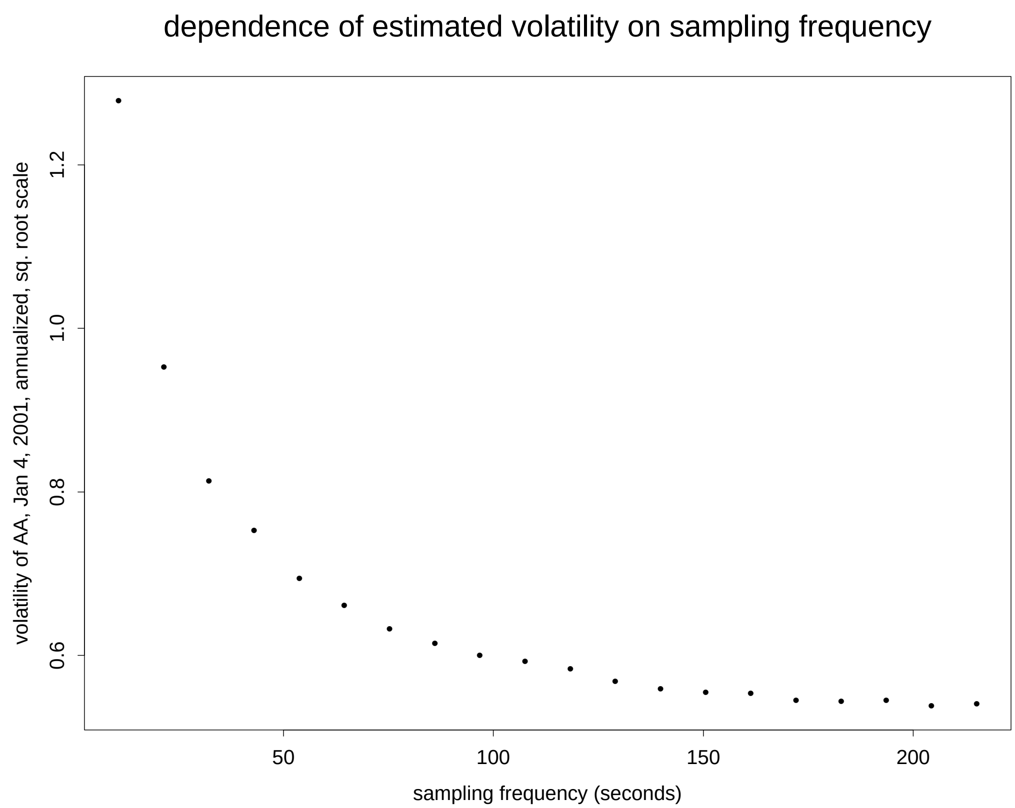

Figure 1. Realized volatility $\sum_{i=1}^{n}(X_{t_{i+1}} - X_{t_i})^2$ when X is sampled at different frequencies, where X is the log transaction price for Alcoa Aluminum on January 4, 2001. The shape of the curve stays the same for other days and other securities. – The plot illustrates that realized volatility diverges at higher sampling frequency (i.e. smaller sampling intervals, as shown by X-axis). This is the opposite of what probability theory predicts.

The two main papers developing this methodology are:

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

Zhang, L. (2006). Efficient estimation of stochastic volatility using noisy observations: A multi-scale Approach, *Bernoulli*, **12 (6)**, 1019-1043.

In Zhang, Mykland and Aït-Sahalia (2005) (A tale of two time scales), it was found that the usual realized volatility mainly estimates the magnitude of the noise term rather than anything to do with volatility. Based on this finding, a methodology was provided to quantify and correct the effect of noise on the nonparametric assessment of integrated volatility. This was a major breakthrough. The approach is built on aggregating the observations on two different sampling scales: a slow-scale (low frequency) estimator achieves the purpose of increasing signal-to-noise ratio, while the usual ("fast scale") realized volatility can be used as a de-biasing device. It was found that the best results can be obtained by combining the estimators from two different scales. The final estimator (*two-scale realized volatility*, TSRV) is asymptotically normal, and converges at the rate of $n^{-1/6}$. The TSRV was the first consistent estimator for integrated volatility when microstructure is present in asset prices.

Zhang (2006) proposed a class of multi-scale realized volatility (MSRV) estimators. MSRV greatly improved the efficiency upon TSRV while maintaining all the desirable properties, such as unbiasedness and the consistency. In particular, MSRV converges at the rate of $n^{-1/4}$, which is the best attainable rate even in parametric inference for constant σ. This is a surprisingly high level of efficiency. The almost-parametric precision for the MSRV makes it essentially cost free to use a general non-parametric specification (1) instead of a parametric model. Zhang (2006) was the first to provide rate efficient (MSRV) estimators of volatility, when microstructure noise is present.

The MSRV is also more robust than the TSRV in that the size of subsamples in the former is bigger. This eases implementation for those data with moderate sample sizes, and reduces the impact of edge effects. In fact, an important feature of the MSRV estimator is its correct handling of edge effects, that is, how to handle the truncation of data at the beginning and end of the trading interval.

The development in the two papers remains valid in the case where there is leverage effect (W_t and σ_t^2 can be dependent). The body of theory is thus very general, it covers all continuous semimartingales. The model given by (1) and (3) is a *hidden semimartingale model.* Central limit theorems for these estimators were also derived in the papers.

B.2 Multiple Securities: Analysis of Variance

While volatility estimation for a single security is important, this is only the first step in understanding the use of high frequency data in financial markets. One natural question is how to relate derivative prices with those of the underlying security. For this purpose, I developed an analysis of variance (ANOVA) for multiple security data. The results are applicable to suitably subsampled data (say, using 5-minute returns as data), or when there is no microstructure.

We have two processes X_t and Y_t. For example, X_t can be the discounted price of a stock, and Y_t can be the discounted price of a market traded option on this stock. One can write their relationship as

$$dY_t = f_t dX_t + dZ_t, \tag{4}$$

where Z_t must be uncorrelated with X_t. One can think of f_t as representing the hedge in X_t against fluctuations in Y_t, while Z_t is the hedging error.

In the simplified setting of the Black-Scholes-Merton (BSM) model, the f_t is given by the theory, and Z_t is zero. In the more realistic world, traders and/or quants need to either (1) hedge empirically, letting f_t be estimated by nonparametric regression, or (2) use another hedging methodology (BSM delta hedge, or an f_t provided by a more elaborate model). In both cases, it is important to be able to estimate the volatility of the ideal (least volatile) Z_t, to monitor performance, and in the case (2) also to assess the adequacy of the trading algorithm.

Questions of this type were studied in

Mykland, P.A., and Zhang, L. (2006). ANOVA for diffusions and Ito processes, *Annals of Statistics*, **34 (4)**, 1931-1963.

The setting is that X_t and Y_t are observed at high frequency. It is shown that one has a decomposition that parallels standard analysis of variance. Specifically if we call $\int_0^t f_s dX_s$ the ideal hedge, the decomposition leads to

$$\text{volatility of } Y \;=\; \text{volatility of ideal hedge} \;+\; \text{volatility of } Z.$$

An estimator is given for the integrated volatility of Z_t. It is shown that the estimator of volatility of Z_t behaves as if Z_t were observed at the trading times.

Since the volatility of Z is a measure of hedging error, one can also use the ANOVA procedure to test whether a financial derivative can be fully hedged in another asset.

C New Findings

C.1 Estimation of Covariance: Asynchronous Transactions

In portfolio management, a central question concerns how to reliably assess the correlation or the covariance between assets prices. This question becomes more challenging in the context of high frequency data. First of all, high frequency time series are often contaminated with microstructure noise. Existing covariance estimators, which ignore the microstructure in the data, become heavily biased when the available data are denser. Secondly, high frequency price series of different assets are often not synchronized. For example, IBM rarely trades at the exact same milli-second as Intel.

Previous Tick Estimators, Epps Effect, and Microstructure. For the traditional previous tick covariance (PTCV) estimator, their distributional properties are studied in:

Zhang, L. (2009). Estimate covariation: Epps effect and microstructure noise. Forthcoming in *Journal of Econometrics*.

In this work it was found that the bias in the traditional covariance estimator is directly related to cumulative nonsynchronicity between the trading times of two different assets, whereas overall time discreteness and microstructure noise contribute to asymptotic variance. Interestingly, when using a subsample with size $o(n)$ (n is total sample size), the effect of nonsynchronicity vanishes asymptotically. The results provide an analytical interpretation of the "Epps effect", documenting how the interpolation-based correlation estimate varies with sampling frequency.

A strong form of these results is shown: there is a set of hypothetical times (same for both assets) so that the suitably subsampled PTCV estimator behaves as if it were observed synchronously at these hypothetical times. This finding has several implications. One is that optimal choice for subsample size is easily established. A deeper implication of this finding concerns the two and multiscale estimators. It is shown in the paper that when using suitable previous tick observations, such as "refresh times", TSRV has exactly the same asymptotic distributional properties as if there were no asynchonicity. The same principle should apply to the multi-scale estimator.

In conclusion – the two- and multiscale estimators automatically remove not only noise, but also asynchronicity.

Local Likelihood Based Estimation. The question of efficient estimators based on asynchronous data is a subtle one, in that it raises the question of whether to use observation times. In the absence of microstructure, this question is investigated in Section 6.3 of

Mykland, P.A. and Zhang, L. (2009a). The econometrics of high frequency data. Forthcoming in *Statistical Methods for Stochastic Differential Equations*, M. Kessler, A. Lindner, and M. Sørensen, eds. Chapman & Hall/CRC Press.

The approach is to set up a locally parametric estimate of covariance, using the observed recording times. Consistency and asymptotic normality is established. It is also shown that this estimator is more efficient than the Hayashi-Yoshida estimator.

C.2 Statistical Trading and Hedging: Local Linear Estimates, and Empirical Model Selection

Trading using High Frequency Regression. The vector of securities prices is high dimensional, including not only stock and bond prices, but also the prices of derivative securities, foreign exchange rates, commodities prices, etc. It will commonly be the case that trading is done while hedging in other securities, with regression model

$$dY_t = \sum_{k=1}^{p} f_t^{(k)} dX_t^{(k)} + dZ_t. \tag{5}$$

The situations covered can be widely different depending on need. For example, if the market portfolio is one of the covariates, then the portfolio becomes market neutral. Another relevant covariate is to let X be a measure of overall volatility (such as VIX), thereby making the portfolio volatility neutral. The latter is particularly interesting in view of current findings suggesting that a large fraction of hedge fund returns are due to a selling of volatility. In a very different setting, Y can be the price of an option to be hedged in underlying securities X.

The formulation (5) was previously used (in the one-covariate case) for the study of ANOVA in Mykland and Zhang (2006). The current development addresses the hedging problem in a more comprehensive fashion. This involves having several securities forming a vector of X's, perhaps an a priori unknown number. Also, the intention was to find a good estimator of the trading position f without making model assumptions.

This program is implemented in

Zhang, L. (2009). Implied and realized volatility: Empirical model selection. (In revision for *Annals of Finance*),

and highlights of the work are described in the following.

Trading with Pre-Selected X Processes. My approach is to use high frequency local linear regression. The high frequency data structure yields that the trading coefficients have desirable properties (consistency and rate efficiency) without parametric assumptions. X can have any dimension.

In the case where the methodology is used to hedge options, the approach yields an empirical version of trading with the minimal martingale measure, as introduced and studied by Föllmer, Sonderman, and Schweizer. Also for the case of options, the regression can be used to compare implied and realized volatilities, with potential for statistical arbitrage between these two. In this formulation, the number of factors in the regression is closely related to the factors in the term structure of volatility. In particular, under a one-factor framework, the options "delta" becomes an additive modification of the (implied) Black-Scholes delta.

Model Selection in High Frequency Regression. In order for this methodology to be fully nonparametric, there is a need for empirical model selection in the regression. The paper proposes two such methods, based on F-testing and on Cross Validation.

For the *F-test* procedure, I define a cumulative F-statistic based on computing the conventional F-statistic within local blocks, and then aggregating the results. I show that under nonparametric assumptions, the resulting statistic can be compared with an F-based threshold, with p-values that are asymptotically accurate. A stepwise procedure (either forward, backward, or both) can now be applied to select the most statistically significant portfolio. The only difference from the classical stepwise methods of linear regression is that the statistic is cumulative over time, and so p-values have to be computed using the high-frequency-based development in the paper.

An alternative approach is to estimate the integrated volatility of the residual process using *Cross-Validation*. Instead of computing a straight estimate of residual volatility (as in the earlier ANOVA paper), I introduced a cross-validated integrated residual variance $\widehat{\langle Z, Z\rangle}_T^{CV}$, as follows.

Within blocks of size M, cross-validated residuals are formed by $\widehat{\Delta Z}_{t_{n,j}}^{CV} = \Delta Y_{t_{n,j}} - \sum_{k=1}^{p} \hat{f}_{(-j)}^{(k)} \Delta X_{t_{n,j}}^{(k)}$ where $(\hat{f}_{(-j)}^{(1)}, ..., \hat{f}_{(-j)}^{(p)})$ is the least squares estimate in the block i containing j $(Mi < j \leq M(i+1))$ using all observations in the block except observation no. j. The cross-validated sum of squares is then CVSS $= \sum_{j=1}^{n} \left(\widehat{\Delta Z}_{t_{n,j}}^{CV}\right)^2$. A computationally efficient method for computing the residual $\widehat{\Delta Z}_{t_{n,j}}^{CV}$ is given in Weisberg's *Applied Linear Regression* (1985), so computation does not require the fitting of n models. Finally, a high frequency limit analysis shows that to get asymptotic unbiasedness, one takes $\widehat{\langle Z, Z\rangle}_T^{CV} = \frac{M-p-2}{M-p-1}$CVSS. The paper proceeds to give the asymptotic variance of $n^{1/2}(\widehat{\langle Z, Z\rangle}_T^{CV} - \langle Z, Z\rangle_T)$.

The advantage of this approach is that we avoid the stepwise aspect of the F-based procedure. A good choice of covariates is one which gives a small cross-validated integrated residual variance.

C.3 Locally Parametric Inference

A substantial amount of work in the area of high frequency data relies on the approximation that the volatility is locally constant. In the paper

Mykland, P.A., and Zhang, L. (2009b). Inference for continuous semimartingales observed at high frequency, *Econometrica*, **77 (5)**, 1403-1445, we investigate the implications of such approximation, and to what extent the ensuing error can be corrected. The main result of the paper is that if the total number of n returns of the process is divided into blocks of size M, and if volatility is held constant at its initial value throughout each block, then this gives rise to a probability distribution Q_n which is *contiguous* (asymptotically mutually absolutely continuous) to the original probability distribution P. (In both cases, this is when viewing the probabilities as measures on the observables X_{t_i}, not the whole process $\{X_t\}$.)

This contiguity makes it straightforward to analyze the discrepancy between P and Q_n, and, in fact, to compensate for it with an ex post adjustment involving asymptotic likelihood ratios. These adjustments are derived and documented. One can therefore "think parametrically" inside each block, and then adjust for any resulting errors after asymptotic analysis. Several examples of estimation are provided: powers of volatility, leverage effect, and integrated betas.

The advantages to thinking parametrically is three-fold, as illustrated by examples in Section 4 of the paper:

• *Efficiency:* In the case of quantities like $\int_0^T |\sigma|_t^r dt$, there can be substantial reduction in asymptotic variance (see Section 4.1).
• *Transparency:* Section 4.2 shows that the analysis of integrated betas reduces to ordinary least squares regression. Similar considerations apply to other examples, such as realized quantiles and ANOVA. Hence, the theoretical development becomes much easier using the blocking method.
• *Definition of new estimators:* In the case of the leverage effect, blocking is a *sine qua non.* Blocking is also helpful for defining new estimators in the multivariante context, where there is less guidance on how to find good estimators.

The case of the leverage effect is one where the benefits of the methodology are starkly clear, and it is worth dwelling on this case for a paragraph (Section 4.3 of the paper). Define in each block $(\tau_{n,i}, \tau_{n,i+1}]$

$$\hat{\sigma}^2_{\tau_{n,i}} = \frac{1}{\Delta t_n (M_n - 1)} \sum_{t_{n,j} \in (\tau_{n,i}, \tau_{n,i+1}]} (\Delta X_{t_{n,j}} - \overline{\Delta X}_{\tau_{n,i}})^2 \text{ and} \tag{6}$$

$$\overline{\Delta X}_{\tau_{n,i}} = \frac{1}{M_n} \sum_{t_{n,j} \in (\tau_{n,i}, \tau_{n,i+1}]} \Delta X_{t_{n,j}} = \frac{1}{M_n}(X_{\tau_{n,i+1}} - X_{\tau_{n,i}})$$

The natural estimator of leverage effect $\langle \sigma^2, X \rangle_T$ would be

$$\widetilde{\langle \sigma^2, X \rangle}_T = \sum_i (\hat{\sigma}^2_{\tau_{n,i+1}} - \hat{\sigma}^2_{\tau_{n,i}})(X_{\tau_{n,i+1}} - X_{\tau_{n,i}}),$$

but it is asymptotically biased. This bias appears in the ex post adjustment involving asymptotic likelihood ratios. The correct, asymptotically unbiased estimator, involves a factor of two:

$$\widehat{\langle \sigma^2, X \rangle}_T = 2\widetilde{\langle \sigma^2, X \rangle}_T.$$

Moving Windows. The above theory is couched in terms of non-overlapping blocks. In fact, it is shown in Section 6.2 of Mykland and Zhang (2009a) that by using a small M-large M interplay, one can generalize the results to moving windows. This is desirable as moving windows are both more natural and more efficient. The estimation of $\int_0^T \sigma_t^p dt$ for general p is here developed in the moving window context, including expressions for the asymptotic variance.

C.4 Further Exploration of the Efficient Price Process

Early papers in high frequency data focused on only two types of returns: (1) continuous returns over a regular sized interval (of, say, five minutes), and (2) jumps. In the ANOVA paper, I sought to go beyond this by allowing time spacings to be irregular, involving concepts such as the quadratic variation of time (Proposition 1 of the paper). The result already showed that the regularity of spacings can substantially affect the asymptotic variance of the realized volatility $\sum_{i=1}^{n}(X_{t_{i+1}} - X_{t_i})^2$.

The importance of going beyond regular spacing is on the one hand to characterize the genesis of heavy-tailed returns. As reported below, I am evolving towards a decomposition of the causes of heavy-tailedness into

- time-varying volatility;
- endogeneity of trading times; and
- price jumps.

The distinction between these three is soft, as particularly discussed in connection with jumps.

Another reason for scrutinizing the trading times is to clean up the conditions for central limit theorems in high frequency data. There are a large number of CLT conditions in the literature, and a broad theory is called for. I am pleased that such a theory has been derived in a recent joint work.

Jump-Ito Processes. In

Zhang, L. (2007). What you don't know cannot hurt you: On the detection of small jumps. Manuscript,

I investigate the impact of the jump component in asset returns on high frequency estimation.

In contrast to (1), the model in this case takes the form

$$X_t = X_0 + \int_0^t \mu_s ds + \int_0^t \sigma_s dW_s + J_t. \tag{7}$$

The additional term $\{J_t\}$ is a Lévy process evolving only through jumps. The function of model (7) is manifold. First, it can take care of rare extreme events such as big changes in securities prices during devaluations and crashes. Second, it covers the possibility that (finitely or infinitely) many small jumps constitute a major part of the evolution in the stock price. This is a popular alternative model for the price process.

The main finding is that it is difficult to distinguish continuous evolution and many small jumps. Specifically, it is shown that one can approximate the process (7) by a continuous process where the volatility σ_t is slightly higher, while removing from J_t all jumps below a certain threshold. The approximation does not affect the behavior of inference or trading.

The approximation in Zhang (2007) has a wide range of applications:

- It shows that very small jumps can be ignored for the purposes of trading (by boosting σ_t).

- It establishes that there can be big biases in estimating σ_t when there are many small jumps.

• It facilitates analysis. For example, the paper finds the asymptotic bias of the multipower estimator for a range of processes that have previously been intractable.

• The approximate process is also much easier to simulate, and the validity of such simulation results is established in the paper.

Endogeneity of Trading Times. We are here back to the continuous model (1). In joint work,

Li, Y., Mykland, P., Renault, E., Zhang, L. and Zheng, X. (2009). Realized volatility when sampling times can be endogenous,

it is shown that if trade times are endogenous, unexpected and non-standard phenomena can occur. The main result in the paper is that even absent microstructure, the realized volatility can be asymptotically biased. Specifically, if $IV = \int_0^T \sigma_t^2 dt$, one can obtain

$$n^{1/2}(RV - IV) \xrightarrow{\mathcal{L}} N(b, a^2)$$

where $b \neq 0$. It should be emphasized that this is different from inconsistency, in that the bias is of the same magnitude as the asymptotic variance. When setting intervals for the integrated variance, these should be centered at $RV + \hat{b}/n^{1/2}$, where $\hat{b}$ is an estimator of b.

An extensive data analysis suggests that the bias does seem to occur in real stock market data. There are also theoretical underpinnings. For example, if transactions occur when the price process hits a barrier (as in limit order), then the bias b will typically be present.

There is a soft interface between endogeneity and jumps, in that a lot of behavior attributed to jumps can also occur under endogenous times.

A General Central Limit Theory for Volatility Type Objects. There has been a need for a general set of conditions for central limit theorems in high frequency data. The typical problem is that the model for the observation times has typically been restrictive (equidistant observations), and this causes problems when there are multiple securities, possibly observed asynchronously.

The search for conditions on times have been carried out in papers by Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008); Hayashi, Jacod, and Yoshida (2009); and Phillips and Yu (2009). Now in Mykland and Zhang (2009a), a general set of conditions has been arrived at. The conditions on spacings are that

$$\sum_{i=0}^{n-1} (t_{n,i+1} - t_{n,i})^3 = O_p(n^{-2}). \tag{8}$$

This covers most contingencies, such as times that arrive as a Poisson type process. In this case, it is shown that the asymptotic variance of the realized volatility doubles relative to its value when times are equidistant. The conditions also covers most endogenous times. Note also, for example, that if each of K securities satisfy condition (8) individually, the refresh-times will satisfy the same condition.

Limit results based on (8) are the core of Sections 3 and 4 of this paper, and we believe that this development will become a gold-standard on which to base further theoretical developments in the area of high frequency data.

C.5 Further Study of Microstructure

In my early collaborative work on TSRV and MSRV, the microstructure noise is taken to be iid and independent of the efficient price process. In follow-up work, we extended our findings under more feasible and practical conditions. The paper

Aït-Sahalia, Mykland, and Zhang (2009). Ultra high frequency volatility estimation with dependent microstructure noise. Forthcoming in *Journal of Econometrics*,

explores the question of volatility estimation when the microstructure noise follows a strong-mixing process. Essentially, the dependence structure in the noise process increases the asymptotic variance in the estimator, but it does not alter the convergence rate nor the optimal order of the sampling grid. Meanwhile, the quality of the normal approximation is explored in

Zhang, Mykland and Aït-Sahalia (2009). Edgeworth expansions for realized volatility and related estimators. Forthcoming in *Journal of Econometrics*,

which provides the explicit form of the Edgeworth expansion for the TSRV and some other estimators.

I have also worked on volatility forecasting with both inter- and intra-day information. In

Kang, Z.X., Zhang, L., and Chen, R. (2009). Forecasting return volatility in the presence of microstructure noise,

a vector fractionally integrated autoregressive and moving average (VARFIMA) model was proposed to jointly model different measures of daily volatility. This approach captures the long memory properties in the volatility process. Also, it successfully incorporates the rich information from options market and from intra-day data (through TSRV) while modeling the dynamic relations. In an out-of-sample comparison, the VARFIMA model significantly outperforms the existing volatility models in the literature.

References

Publications

[1] Estimating covariation: Epps effect, microstructure noise. Zhang, L. (2009), forthcoming in *Journal of Econometrics.*

[2] Edgeworth expansions for realized volatility and related estimators. Zhang, L., Mykland, P.A., and Aït-Sahalia, Y. (2009), forthcoming in *Journal of Econometrics.*

[3] The econometrics of high frequency data. Mykland, P.A. and Zhang, L. (2009a). forthcoming in *Statistical Methods for Stochastic Differential Equations*, M. Kessler, A. Lindner, and M. Sørensen, eds. Chapman & Hall/CRC Press.

[4] Ultra high frequency volatility estimation with dependent microstructure noise. Aït-Sahalia, Y., Mykland, P.A. and Zhang, L. (2009), forthcoming in *Journal of Econometrics.*

[5] Inference for continuous semimartingales observed at high frequency. Mykland, P.A. and Zhang,

L. (2009b), *Econometrica*, **77 (5)**, 1403-1445.

[6] Inference for volatility-type objects and implications for hedging. Mykland, P.A. and Zhang, L. (2008), *Statistics and Its Interface*, **1**, 255-278.

[7] Efficient estimation of stochastic volatility using noisy observations: A multi-scale approach. Zhang, L. (2006), *Bernoulli*, **12 (6)**, 1019-1043.

[8] ANOVA for diffusions and Ito processes. Mykland, P.A. and Zhang, L. (2006), *Annals of Statistics*, **34 (4)**, 1931-1963.

[9] Comment on "Realized variance and market microstructure noise" by P.R. Hansen and A. Lunde. Aït-Sahalia, Y., Mykland, P.A. and Zhang, L. (2006), *Journal of Business and Economic Statistics*, **24 (2)**, 162-167.

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

[11] Discussion of paper "A selective overview of nonparametric methods in financial econometrics" by Jianqing Fan, Mykland, P.A. and Zhang, L. (2005), *Statistical Science*, **20 (4)**, 347-350.

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

In Revision, Submitted, Working Papers

[13] Implied and realized volatility: Empirical model selection. Zhang, L. (2009). (In revision for *Annals of Finance*)

[14] Forecasting return volatility in the presence of microstructure noise. Kang, Z.X., Zhang, L., and Chen, R. (2009).

[15] Realized volatility when sampling times can be endogenous. Li, Y., Mykland, P., Renault, E., Zhang, L. and Zheng, X. (2009).

[16] What you don't know cannot hurt you: On the detection of small jumps. Zhang, L. (2007).