Review of Futures Markets

Volume 19 Special Edition

Research funding for this Special Edition was made possible by a generous donation from the Clearing Corporation Charitable Foundation to the Institute for Financial Markets.

IFM THE INSTITUTE FOR FINANCIAL MARKETS

In Cooperation with The IFM

HIGH-FREQUENCY TRADING: METHODOLOGIES AND MARKET IMPACT

Frank J. Fabozzi, Sergio M. Focardi, and Caroline Jonas*

This paper discusses the state of the art of high-frequency trading (HFT), its requisite input, high-frequency data (HFD), and the impact of HFT on financial markets. The econometrics of HFD and trading marks a significant departure from the econometrics used when dealing with lower frequencies. In particular, ultra HFD might be randomly spaced, requiring point process techniques, while quantities such as volatility become nearly observable with HFD. At high frequency, forecasting opportunities that are different from those present at lower frequencies appear, calling for new strategies and a new generation of trading algorithms. New risks associated with the speed of HFT emerge. The notion of interaction between algorithms becomes critical, requiring the careful design of electronic markets.

In this paper, we discuss the state of the art of high-frequency trading (HFT) and important issues related to the econometric analysis of high-frequency data (HFD) and the impact of HFT on financial markets. The econometrics of HFD is different from standard econometric analysis employed in the analysis of lower frequency data. In particular, time series of HFD might be randomly spaced, thereby requiring the techniques of point processes. Many quantities such as volatility become nearly observable. At high frequency, forecasting opportunities that are different from those present at lower frequency appear, calling for a new generation of trading algorithms. As we explain in this paper, this results in the emergence of new risks related to the speed of HFT. The notion of interaction between algorithms becomes critical, requiring the careful design of electronic markets.

---

*Frank J. Fabozzi CFA (the corresponding author) is a professor in the Practice of Finance at the Yale School of Management. E-mail: fabozzi321@aol.com.
Sergio M. Focardi is a professor of finance at the EDHEC Business School. E-mail: sergio.focardi@edhec.edu.
Caroline Jonas is a partner at The Intertek Group (Paris). E-mail: cjonas@theintertekgroup.com.

Acknowledgements: This survey paper on high-frequency data and high-frequency trading is based on a review of the literature and conversations with 13 academics prominent in econometrics and market microstructure and three representatives from major exchanges. The academics have been identified throughout the paper; the exchange personnel whom we interviewed are not disclosed at their request. The authors wish to thank all those who shared their insights and experience. This paper was prepared under a grant provided by The Institute for Financial Markets.

Keywords: high-frequency data (HFD), high-frequency trading (HFT)
JEL Classification: G10, G12

Figure 1. Algorithmic Trading Adoption by Asset Class.

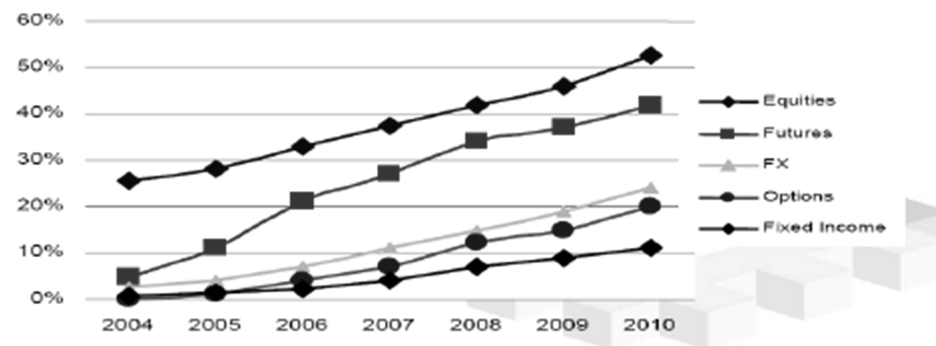

Source: Aite Group estimates.

I. DEFINING HIGH-FREQUENCY TRADING

Although there is no universally accepted definition of HFT, among its defining characteristics are the fact that investments are held for very short periods of time and typically (but not necessarily) positions are not carried overnight. How to quantify these characteristics is a matter of debate. Kearns, Kulesza, and Nevmyvaka (2010) define high-frequency traders (HFTers) as those traders who hold positions between 10 milliseconds and 10 seconds. However, the U.S. Securities and Exchange Commission (SEC) adopts a somewhat less precise definition, defining HFTers as professionals acting in a proprietary capacity and able to generate a large number of trades per day.

HFT is a form of trading that leverages high-speed computing, high-speed communications, tick-by-tick data, and technological advances to execute trades in as little as milliseconds. A typical objective of HFTers is to identify and capture (small) price discrepancies present in the market. They do so with no human intervention, using computers to automatically capture and read market data in real-time, transmit thousands of order messages per second to an exchange, and execute, cancel, or replace orders based on new information on prices or demand.

High-speed trading strategies use computerized quantitative models (i.e., algorithms) that identify which type of financial instrument (for example, stocks, options, or futures) to buy or sell, as well as the quantity, price, timing, and location of the trades. In this paper, we focus on the equity market and equity futures and options. While algorithmic trading is now used in many asset classes, its origin is in equities and, still today, the share of trades based on algorithms is highest in the equity market (see Figure 1).

It is widely estimated that HFT was responsible for 40 to 70% of all trading volume in the U.S. equities market in 2009, roughly double its share just four years earlier; it is estimated to represent about 35 to 40% of all trading volume in European equities.

In practice, HFT is engaged in by a wide variety of entities including proprietary desks, hedge funds, and institutional investors. Nevertheless, it is estimated that high-frequency transactions in the U.S. equities markets are initiated by just 2% of the 20,000 trading firms in the United States, that is to say, by some 400 firms (see Clark 2010). Many of these firms are privately held proprietary trading firms or hedge funds. The biggest players in HFT are reported to include the electronic market-makers Getco, Tradebot, Citadel, and QuantLab; hedge funds such as D.E. Shaw, SAC Global Advisors, and Renaissance Technologies; and the proprietary trading desks of Goldman Sachs, Morgan Stanley, and Deutsche Bank. The technology goal of HFTers is to reduce latency (i.e., delay) in placing, filling, confirming, or cancelling orders; the business goal is typically to profit from small arbitrage opportunities present at short time horizons. Trading strategies differ and include electronic market-making and statistical arbitrage.

A. Setting the Stage for HFT

A number of factors have combined with technology to lead to an explosion in (algorithmic) trading activity. First, the 2001 decimalization of U.S. capital markets coupled with smaller tick sizes led to an explosion in market data volumes. Chakravarty, Harris, and Wood (2001) analyzed the effect of decimalization in the transition period and found a significant increase in trading volumes after decimalization. They note that the SEC expected a 139% increase in the number of trades due to decimalization. Second, the cost of trading has dropped. This was a consequence of several decisions, including the 1998 SEC decision to authorize electronic exchanges to compete with the traditional exchanges. It is estimated that while in the 1990s the New York Stock Exchange (NYSE) and Nasdaq accounted for 80% of trading volume in securities they listed, as much as 60 to 70% of trading in their listed companies is now dispersed on as many as 50 competing trading venues, for the most part fully electronic. Third, an increase in derivatives products and exchange-traded funds (ETFs) has led to an explosion in trading volumes. Angel, Harris, and Spatt (2010) report that equity trading volumes tripled in recent years, going from about 3 billion shares per day in 2003 to nearly 10 billion shares per day in 2009. According to data from the NYSE, average daily volume on U.S. stock exchanges was up 164 percent in 2009 compared to 2005 (see Duhigg 2009).

At the same time, high-performance computing systems, advanced trading technology, and low-latency messaging middleware and feed handlers have reduced the time necessary to execute market orders. Angel et al. (2010) cite data from Thomson, according to which the speed of execution for small market orders has gone from about 25 seconds for NYSE-listed firms and 5 seconds for Nasdaq-listed firms in September 2001 to about 2.5 seconds in August 2009 (see Figure 2).

According to Eric Bertrand of NYSE Technologies (see Bertrand 2009), the capacity as measured by order messages per day has gone from one million in 1995 to hundreds of millions in 2009. During the same period (i.e., 1995–2009), throughput as measured by messages per second has gone from 20 to over 100,000 and latency from one second to one thousandth of a second (i.e., one millisecond). At the same time, network and data distribution speeds have gone from 64 kb per second to 10-100 Mb per second. Bertrand foresees order messages per day going to billions, messages per second to millions, latency to millionths of a second (i.e., microseconds), and network and data distribution speeds to a gigabyte per second.

Figure 2. Market Order Execution Speed.

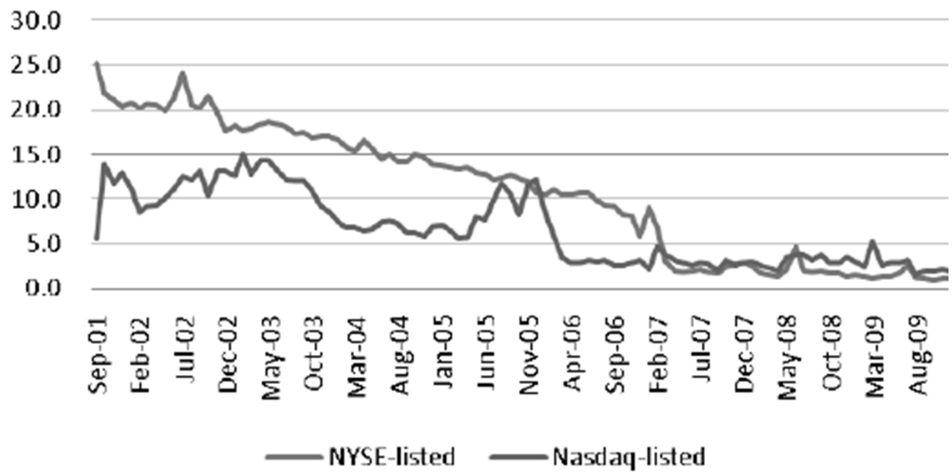

Note: Evolution of market order execution speeds as measured in seconds, concerning NYSE-listed and Nasdaq-listed firms during the period Sept 2001–August 2009 (from Angel et al., p. 22).
Source: Rule 605 data from Thomson for all eligible market orders (100-9999 shares).

To further reduce latency, HFTers are placing their trading servers at the trading venues to be close to the exchange matching engines. This is commonly referred to as co-location. In her March 2010 *Chicago Fed Letter* Carol Clark, a financial markets and payments system risk specialist in the Chicago Federal Reserve's financial market group (see Clark 2010) remarks that it is estimated that for each 100 miles the server is located away from the matching engine, 1 millisecond of delay is added to the time it takes to transmit trade instructions and execute matched trades or to access the central order book where information on buy/sell quotes and current market prices is warehoused.

The NYSE is completing construction of a nearly 400,000-square-foot data center facility in Mahwah, New Jersey, where it hopes to attract in co-location large Wall Street banks, traditional brokerages, and hedge funds. The center's 40-gigabyte-per-second standard hardware will allow it to handle up to a million messages a second; new trading technology will reduce latency to 10 microseconds. Meanwhile, work is proceeding at the NYSE Euronext to design an ultra-low latency core network that will support 50-microsecond roundtrips.

II. ECONOMETRICS FOR HFT AND ULTRA HFT DATA

As mentioned above, daily closing price data typically used in past efforts at modeling financial markets are not sufficient for engineering HFT strategies; the latter calls for the use of HFD, data taken at intraday frequencies, typically minutes. Data relative to each transaction, or tick-by-tick data, are called ultra high-frequency data (UHFD). HFD and UHFD might be considered the fuel of HFT.

In this section we will first discuss questions related to the handling of (U)HFD and then discuss separately the modeling of HFD and UHFD. We will do so because, from an econometric perspective, there is a distinction between the methods and research objectives of HFD and UHFD. Both HFD and UHFD require econometric methodologies different from those employed at lower frequencies.

A. Data Handling Issues

(U)HFD are routinely provided by electronic exchanges, albeit at a possibly high price. Data currently available include tick-by-tick data and order-book data. A "tick" includes information at a given time, the "time stamp." The sequence and content of the ticks might depend on the time of observations and on the exchanges that are observed. Significant differences between the ticks of different exchanges might be due to technology, exchange structure, and regulation. Order-book data availability is not the same on all exchanges. Some exchanges offer complete visibility on the order book while others offer only partial visibility. Still other exchanges "flash" the order book only for a short period of time, for example, a fraction of a second.

HFD and UHFD present significant problems of data handling. (See Brownlees and Gallo 2006 for a review of the challenges.) Both HFD and UHFD need to be filtered as errors and outliers might appear in a sequence of ticks. Bauwens and Giot (2001) and Oomen (2006), among others, deal with many aspects related to data cleansing. Brownlees and Gallo (2006) analyze the question of cleansing data from the NYSE's Trades and Quotes (TAQ) files. Boehmer, Grammig, and Theissen (2006) discuss problems related to synchronizing data from the TAQ and from the NYSE's order book.

Falkenberry (2002) reports that errors are present both in automatic and semiautomatic trading systems. He reports that, as the speed of transactions increases, errors become more frequent. The first task in data cleansing is therefore the elimination of erroneous data. However, it is also important to deal with outliers and with data that are not compatible with normal market activity. Methods for eliminating outliers are described in Boehmer et al. (2006).

In addition, HFD are not simply observed but imply some form of interpolation in order to represent prices. In fact, by the nature of the trading process, the truly "primitive" observations, that is, tick-by-tick data or UHFD, are an irregularly spaced time series given that trading and quotes occur at random times. For example, the frequency of UHFD for individual assets varies within a wide range of values in function of the observed processes (i.e., trades). In his study of HFT activity relative to 120 stocks traded on the NYSE, Brogaard (2010) found trading frequencies ranging from eight transactions per day for the lesser traded stocks to 60,000 transactions per day, or roughly two transactions per second on average, for the most heavily traded stocks.

If we want to construct regularly spaced sequences of HFD, we must use a methodology to determine a price in moments when there are no transactions. Methods include linear interpolation between the two closest observations or using the previous or ensuing observation. If data have a high frequency, these two methods yield similar results. For rarely traded securities, different methods might result in significant differences.

B. Better Econometrics with (U)HFD?

The availability of (U)FHD has been welcomed as a major advance with the potential of revolutionizing the study and the practice of econometrics. The expectation is that with (U)HFD, market participants can significantly improve the estimation of parameters used in continuous-time finance and "observe" quantities such as covariances or volatility as opposed to having to treat them as hidden variables.

However, it has become clear that there are significant limitations in the use of HFD in general. As we will discuss, limitations come mainly from two sources. First, due to market microstructure effects, the behavior of prices at time horizons of the order of seconds is different from the behavior of prices at time horizons of minutes or longer, thus introducing basic limitations in the use of HFD. Second, it is difficult to compute correlations and covariances between assets that trade at significantly different frequencies.

There are possibly different models at different time scales; a single model that is valid at every time scale and in every time window, if it exists at all, would be too difficult to create and to estimate. The usual assumption is that prices follow a jump-diffusion process.

Jump-diffusion processes allow to describe with some accuracy the statistical uncertainty of financial quantities. Thus, a jump-diffusion model of prices allows a reasonable representation of the statistical characteristics of the uncertainty of the distribution of returns and of co-movements between returns. However, the deterministic drifts can be estimated only with limited precision, and they depend on the data sample employed. Jump-diffusion processes do not allow one to make accurate forecasts based on trends and drifts. If we estimate jump-diffusion processes on different samples of past data, we obtain intrinsically different estimates of drifts although the estimates of volatilities and covariances can be made reasonably coherent. Therefore, although the use of HFD represents a significant step forward in the estimation of some financial quantities, it does not allow us to formulate universal laws.

Let us now look at the limitations in the use of (U)HFD. From a purely statistical point of view, estimates improve with a growing number of samples. Therefore, it would seem reasonable to use all available (U)HFD. However the behavior of prices at very high frequencies is not the same as the behavior of prices at lower frequencies. In fact, assuming that prices are modeled as jump-diffusion processes, as the length of sampling intervals approaches the length of trading intervals, micro structure effects introduce biases. These biases reduce the accuracy of forecasts.

Actually, as described in Aït-Sahalia and Mykland (2003), we can identify several different effects that limit our ability to estimate continuous-time models. First, the inevitable discreteness of samples, both in time and price, introduces biases in estimation. These are the first effects studied in the literature on estimating continuous-time models. Second, the randomness of spacing, which introduces biases that, following Aït-Sahalia and Mykland, are at least as large as the discreteness effects. Third, there are many microstructure effects, possibly exchange-dependent, which are generally accounted for as “noise” in the observation of prices. A number of papers have analyzed the theoretical and empirical optimal sampling frequency at which prices should be sampled to estimate the covariance matrix of diffusion processes.[1]

There is no consensus as to including noise in the observation of prices. Ionut Florescu, assistant professor of mathematics in the Department of Mathematical Sciences of the Stevens Institute of Technology, remarks that the paradigm of noisy observations is typical of physics and engineering, but he suggests that it does not really apply to finance. Professor Florescu says, “A price of a trade is not a noisy observation: We introduce noise only as a mathematical idealization.” His research effort is focused on estimating continuous-time models starting from “true” observations.

C. Using UHFD in Econometrics

The econometrics of UHFD is interested in representing the process of the random arrival of trades. The latter is important to HFTers because there are relationships between the volume of trades and prices. The econometric study of UHFD cannot be performed with the usual methods of time series analysis, given that the latter assume observations at fixed time intervals. The problems associated with and methods applicable to UHFD are specific to randomly sampled data. An early model of nonsynchronous data is Lo and MacKinlay (1990). Bauwens and Hautsch (2006a) and Hautsch (2004) provide overviews of the modeling of randomly spaced financial data.

Trades are events of random magnitude that occur at random times. The times at which trades take place are a sequence of strictly increasing random variables. The number of trades $N(t)$ in any given interval is also a random variable. Processes of this type are referred to as point processes.

Point processes are continuous-time processes given that an event[2] might occur at any moment; they are well known mathematical constructs in the field of insurance where claims of unpredictable magnitude occur at random times. The simplest point process is the Poisson process, which is characterized by the following properties:

- The number of events in any given interval of time is a random variable that follows a Poisson distribution: $e^{-\lambda} \dfrac{\lambda^k}{k!}$

1. See, among others, Zhang, Mykland, and Aït-Sahalia, (2005), Aït-Sahalia, Mykland, and Zhang (2005), Bandi and Russell (2006), Bandi and Russell (2008), and Bandi, Russell, and Zhu (2008), Voev and Lunde (2007).
2. We use the term “event” not in the sense of probabilistic events but to denote something that occurs at a given time, for example, a trade.

- The number of events in any given interval of time is independent from the number of events that occurred in any previous interval.

- The distribution of the time between two consecutive events follows an exponential distribution whose density is: $\lambda e^{-\lambda t}$

The parameter λ is called the intensity of the process. Poisson processes are characterized by constant intensity. The Poisson process is the point-process equivalent of the Brownian motion: It implements the notion of total uncertainty as regards the moment when the next event will occur. If a queue is described by a Poisson process, the probability that an event will occur in any future interval is unrelated to the time elapsed since the last event. For example, if a Poisson process describes the passage of a bus, a passenger waiting for the bus would have always the same probability to catch a bus in any next period independently of how long he/she has been waiting for the bus.

The Poisson process is a parsimoniously parameterized process with attractive mathematical properties, but it is too simple to describe the arrival times of trades. In fact, the time intervals between trades, referred to as the durations between trades, are not independent but exhibit autocorrelation phenomena. In order to represent autocorrelations, we need to generalize Poisson processes to allow for time-varying intensity. Point processes where the intensity is a separated process are called Cox processes.

Engle and Russell (1998) introduced a particular Cox process that they called an Autoregressive Conditional Duration (ACD) process. ACD processes are the point process equivalent of ARCH/GARCH models insofar as they allow autoregressive intensity. The original ACD has been generalized and extended in many different ways, for example in Bauwens and Veredas (2004) and Bauwens and Hautsch (2006b). McAleer and Medeiros (2008) and Pacurar (2008) provide a summary of theoretical and empirical work done on the ACD models. The ACD model and its generalizations are now widely used in the study of intra-trade durations.

D. The Econometric Study of HFD

While the econometrics of UHFD is mainly interested in representing the process of the random arrival of trades, the econometrics of HFD is principally interested in estimating covariances, which are fundamental data for any investment process. As described above, HFD are data taken at fixed intraday frequencies, typically from a few minutes to less than an hour. When raw data are prices in the form of ticks, HFD are recovered using some form of data aggregation and interpolation.

Although HFD are classical time series, they are typically modeled as continuous-time models, typically jump-diffusion processes, sampled at finite intervals. The underlying reasoning is that HFD tend to a continuous-time process if the observation frequency grows. Intuitively, one might think that a jump is a large discontinuity so that a jump-diffusion process simulates large movements such as crashes. However, mathematically this is not the case. A discontinuity is a point where the left and right limits of a path do not coincide regardless of the size of the difference. Therefore, a jump-diffusion process is a rather abstract mathematical concept that is useful to provide a better fit to the distribution of returns found empirically, but it is not necessarily related to big jumps in price processes.

Mathematically, if we sample a continuous-time process with time intervals that tend to zero, many quantities estimated on the sampled process will tend to an average of the true parameters of the process. For example, if we compute a covariance matrix on a given interval using an increasing number of points, the empirical covariance matrix will tend to the average of the theoretical instantaneous covariance. It should be noted that the above is a theoretical property of jump-diffusion processes sampled at frequencies that tend to infinity. Therefore, we can state that volatilities and covariances estimated with high frequency intra-day data tend to the true volatilities and covariances only if we assume that price processes are jump-diffusion processes. If they are not, the above property might not hold.

1. Applying HFD to the Measurement of Volatility

With the above caveat, assuming prices are jump-diffusion processes, one of the major applications of HFD is the measurement of volatility. When prices and returns are observed at time intervals of days or weeks, volatility is a hidden variable typically modeled with ARCH/GARCH models. When HFD are available, volatility is considered to be almost observable. This is because with HFD we have sufficient intraday data to estimate daily volatility as an average of the instantaneous volatility. Though it is conceptually wrong to say that volatility can be observed with HFD, it is nevertheless possible to make very precise estimates of the average volatility over short intervals where volatility does not change much. A number of papers have discussed the measurement of volatility at high frequency.[3]

The problem of forecasting volatility remains. Because observed daily volatility changes significantly from day to day, there is the need to forecast volatility. A general class of models for forecasting volatility, the Multiplicative Error Model, was introduced in Engle (2002) and extended in Cipollini, Engle, and Gallo (2006). For a comparison of different methods used to forecast volatility, see Brownlees and Gallo (2007).

From the above, it is clear that the interest in HFD is related to the fact that they make available a much larger quantity of data with respect to daily observations, and they do so without stretching the observation period. Dacorogna et al. (2001) observed that, on average, one day of HFD contains as many data as 30 years of daily data. Today, in some markets, this estimate can be multiplied 10 times. Therefore, it would seem reasonable to consider that HFD allow estimating richer models with more parameters. However, this advantage might have limitations given that we have to capture an intraday dynamics that is not needed when we model daily data. In other words, it is questionable if HFD aid us in understanding data at longer time

3. See, among others, Andersen, Bollerslev, Diebold and Labys, (2001), Andersen et al. (2003), Andersen, Bollerslev, and Meddahi (2002), Bandi and Phillips (2003), Barndorff-Nielsen and Shephard (2002a, b), Barndorff-Nielsen and Shephard (2004), Hansen, Lunde, and Voev (2007), and Ghysels, Santa-Clara, and Valkanov (2006).

horizons. For example, daily volatilities change and need to be forecasted; in addition very short-term movements are generated by microstructure effects.

Commenting on how HFD can be used for forecasting longer time horizons, Ravi Jagannathan, Chicago Mercantile Exchange/John F. Sandner Professor of Finance and a Co-Director of the Financial Institutions and Markets Research Center at Northwestern University, remarks:

> HFD does help forecast at longer time horizons, but not very long. HFD do help for forecasting one week ahead, but not one year ahead. HFD poses an enormous challenge: If price moves between bid/ask, microstructure noise dominates. You need to filter out more microstructure noise. For example, if you look at what happened 6 May 2010 and observe HFD, it will not tell you much about what might happen next week.

The question is primarily empirical, but there are also theoretical considerations. The problem can be stated as follows. Suppose there is a true price process $p(t)$, which we assume is generated by a jump-diffusion mechanism. This model includes a time-dependent instantaneous covariance matrix p_t. Suppose we can observe the true process only at discrete points p_t in a given interval. It can be demonstrated (see Barndorff-Nielsen and Shephard 2002a,b) that if the frequency of observations tends to infinity, then the empirical covariance tends to the integral of the instantaneous covariance.

However, if we assume that our observations are contaminated by market microstructure noise, then estimates of the covariance matrix are negatively biased. Aït-Sahalia and Mykland (2003), Aït-Sahalia, Mykland, and Zhang (2005), Bandi and Russell (2006, 2008), Bandi, Russell, and Zhu (2008) determine the optimal sampling rate in the presence of microstructure effects.

Professor Jagannathan observes that, in the case of volatility measurements:

> If markets are frictionless, that is, if there are no microstructure effects, the higher the frequency, the better the measurement of values as volatility. However, in rare or severe events, HFD are of no help; microstructure — the way people trade, the strategies used, lack of knowledge of what the others are doing — becomes more important. These effects are particularly severe for illiquid stocks. To make use of HFD, you have to have people trade at high frequency. If people trade at high frequency, you have observations. The econometrician can understand what is going on.

E. Different Pricing Theories for Different Data Frequencies?

We observed above that there is a big difference in the frequency of trading at the level of individual assets and that HFT has exacerbated this phenomenon in that most HFT is concentrated in a small number of stocks. Given this difference, and given the importance of HFD on pricing theories, we might ask if we need different pricing theories for assets that are heavily traded and assets that are not. The question can be reformulated as understanding what impact, if any, HFT has on price processes.

Frederi Viens, Professor of Statistics and Mathematics and Director of the Computational Finance Program at Purdue University, offers an initial response:

> It is my guess is that HFT impacts price processes in a big way. As far as I am aware, financial mathematics people have not yet found a way to explain how to price equities under microstructure noise without arbitrage, and therefore I would venture to say that high-frequency-traded stocks can still be priced using standard frequency methods, but there will be some uncertainty in the pricing due to the microstructure noise. I am not aware of any way to perform equity and option pricing in an arbitrage-free way on UHFD without having to resort to saying that microstructure noise exists. However, if one such way would exist, it would automatically imply that there should be two distinct pricing theories depending on the frequency of trading. That would be a most uncomfortable situation. My guess is that microstructure noise is real, so that we simply have to deal with it, that is to say, account for the added uncertainty in our prices. Theoretically, this added uncertainly goes against the possibility of arbitrage opportunities. Since, in practice, the contrary is true, a balance will only be achieved when enough people have access to and the ability to work with UHFD.

When discussing the relationship of HFD and long-term behavior, there are actually two distinct problems: the problem of the model itself and the problem of noise. Professor Viens observes:

> The problem with HFD as it relates to longer-term trends is that the market microstructure which is visible using HFD may or may not have any bearing on the longer term trends. This is still being hotly debated in academia. We are quite a way from being able to provide definite answers on this debate, and my guess is that the connection between the two will be relevant in some markets, and irrelevant in others. … One theoretical example where the two are linked is the case of self-similar markets, particularly ones where stochastic long memory occurs because of so-called fractional Gaussian noise. From my experience with real data, I can say that there is no evidence of any markets with such a self-similarity property. In other words, I have first-hand evidence showing that important long-term market parameters, such as stochastic long memory for volatility series, cannot be estimated using UHFD or even HFD.

F. Benefits of (U)HFD

In general, the more data that are available, the happier the statistician is. For econometricians and financial modelers, the availability of (U)HFD is beneficial to understanding what happens to prices intraday and might help shed light on financial econometrics in general. Eric Ghysels, Bernstein Distinguished Professor of

Economics at the University of North Carolina's Kenan-Flagler Business School, says:

> HFD allow us to improve estimation of certain parameters or models used in various financial applications ranging from derivative pricing to asset allocation. HFD also allow us to improve upon existing market-based measures or to construct new ones. Prominent examples include volatility and correlation. HFD and UHFD also allow us to study certain phenomena related to the actual trading process — topics that could not be studied without such data. Examples here are abundant and relate to the so-called market microstructure literature.

(U)HFD are also a challenge for the econometrician or modeler. Nikolaus Hautsch, who holds the Chair of Econometrics at the Center for Applied Statistics and Economics at Humboldt University in Berlin, comments:

> HFD are affected by a lot of noise, lots of data with no information content. What matters is the ratio between the signal to noise. The signal-to-noise ratio must be greater than 1. If not, we have more noise than signal, and no gain. In the very beginning, the role of noise was overlooked. Over the past four, five years, we have gained a better understanding of this.

We will now take a closer look at what academics to whom we spoke identified as specific benefits related to the availability and use of (U)HFD.

1. Better Understanding of Market Microstructure and the its Impact on Modeling

Academics we interviewed agreed that (U)HFD are useful in gaining an understanding of phenomena that occur intraday and the microstructure that causes them. Chester Spatt, the Pamela R. and Kenneth B. Dunn Professor of Finance and Director of the Center for Financial Markets at Carnegie Mellon University's David A. Tepper School of Business, comments:

> There is information in small bids, small grains that might be significant as they reflect opinions. But not all that shows up in trading is information; it might be a question of micro market structure friction. (U)HFD is very interesting as it allows us to understand the trading process, to drill down. Using only daily data, one cannot understand the fundamentals of the trading process, the motors of decision processes of traders in different contexts. For example, to what extent does an intermediary's inventory influence his decisions?

The expectation is that the availability of (U)HFD will allow better design of exchanges. Valeri Voev, assistant professor of finance at the University of Aarhus (Denmark), says, "HFD is beneficial in studying the design of markets, to decide on market microstructure issues such as an order-driven or a quote-driven market, the role of specialists, etc., in an effort to design better markets."

The analysis of HFD and the study of market microstructure go together, in the sense that, while HFD reveal microstructure, it is also true that understanding microstructure offers a better understanding of HFD. As remarked by Professor Ghysels:

> The modeling of HFD is dependent on the exchange from which they are generated. Are there implications for price discovery and risk management? This is a topic that has been widely studied in the market microstructure literature, notably how price discovery takes place under various trading mechanisms. Part of this literature relies on the different time series characteristics of prices under alternative trading rules.

Professor Hautsch concurs, adding:

> We definitely need to take into consideration the structure of the market place where the data is generated, for example, a market-maker or electronic exchange. The dynamics are different, the levels of noise are quite different, the tick sizes are quite different. Some markets, for example electronics markets, create a lot of noise. If one does not take these factors into consideration, one gets spurious results, strange outcomes.

Professor Florescu says, "(U)HFD offer an unparalleled opportunity to study the trading process and implement learning with artificial intelligence as machines are pitched one against the other and against humans."

2. Improved Measurement of Phenomena at Lower Frequencies, Including Volatility, Covariance, and Risk

Academics whom we interviewed agreed that (U)HFD can also enhance an understanding of lower frequency phenomena, because (U)HFD allow one to model observed quantities and not only hidden quantities. Volatility is a case in point. Though we need to forecast volatility, our forecasts are based on models of observed volatility. Luc Bauwens, professor of finance at the Catholic University in Louvain (Belgium), enumerates:

> First, many useful theoretical pricing models are formulated in continuous time. With UHFD especially, these models can be estimated much better than with less highly frequent data. Second, UHFD data allow to measure volatilities of returns — say daily volatilities — much more precisely than without these data — say when only daily data are available — through "realized volatilities." Third, risk and liquidity can be measured in real time with UHFD.

Professor Bauwens adds:

> In all these areas, much progress is still to be made. From an econometric point of view, UHFD are interesting because they pose a number of issues that have not been much studied earlier by statisticians in the field of finance. There are many open questions in the analysis of time-dependent

data that are irregularly spaced and when the time dependences are complex, for instance, beyond the conditional mean.

According to Professor Voev:

> We can benefit from HFD as many traditional markets use daily returns. Daily squared returns are very noisy. For example, if observations at the beginning and end of the day are the same, then daily returns information shows zero fluctuations versus if there were fluctuations during the day. We can get big performance gains if we use more frequent intraday data because we obtain more statistical precision. We need to know the true volatility ex post. With HFD can get very precise ex post measure of volatility. HFD are a good starting point to measure and understand volatility.

However, just how to use HFD might not be so obvious. In fact, HFD permit the precise measurement of past data but rely on forecasting to extrapolate these measurements. Professor Voev comments, "Evidence is pretty clear that the HFD offer better measurement but it is still not clear that we can optimize the use of this information. When talking about multivariate data volatility, we need to come up with models that allow forecasting matrices."

However, estimating covariances between data at different frequencies is a significant obstacle. According to Professor Hautsch:

> Over the last 10 years, in the literature, the use of HFD has led to more and more efficient estimates of the daily co-variance. However, there are potential problems when we estimate quantities relative to data with different frequency. Assets with high/low liquidity are a big problem if one tries to correlate assets that trade thousands of times a day and assets that trade three times a day. This creates biases. It is a statistical problem that needs to be resolved.

3. Improved Estimation of the Returns Distribution

Having thousands of observations of returns available, one can perform a precise estimate of the return distribution. Of course, there is a caveat: If daily returns are required, we need to project high frequency returns onto daily returns. Doing so requires models of the time evolution of returns and precise measurements of autocorrelations. Still, Professor Voev observes, "We obtain a much better design of the whole returns distribution based on thousands of trades per day."

4. Better Understanding of Liquidity

The study of liquidity is a notoriously difficult problem. Its very definition presents difficulties. The availability of HFD, and more recently the diffusion of HFT, allows one to shed more light on phenomena related to liquidity. Professor Hautsch observes:

> The relationship between liquidity and volatility is very difficult. We cannot understand it well from data 10 years or more back because liquidity then played a completely different role from that it plays today. All work on market microstructure [when markets were populated by market-makers] is no longer relevant. We have a paradigm change, a fundamental change in markets.

5. Discovering New Facts

Professor Hautsch points to the role (U)HFD plays in discovering new facts and theories:

> HFD are interesting in that they need new econometric models to take into account specific properties of data. Properties have changed quite recently given the enormous liquidity in the markets. This raises new statistical problems. The challenge is to manage higher dimensions of data: many characteristics, different markets, limit-order book data. HFD allow one to build better large-scale models, make better estimations of correlations, better estimations of (high-dimensional) co-variance.

6. Improved Market Efficiency

Academics also agree that HFD (as well as HFT) has improved market efficiency. Professor Viens comments:

> From my standpoint as a mathematician and statistician working in quantitative finance with tools from stochastic analysis, I can only say that the more HFD, and especially UHFD, become available to a wider audience — including the ability to analyze such data thanks to increasing computational speed — the more efficient the market should become.

III. HIGH-FREQUENCY TRADING

HFT has become the subject of intense debate; it is feared that the use of computerized programs and high-speed computers and communications networks that characterize HFT might create new risks and allow HFTers to realize profits at the expense of bona fide but less sophisticated investors.

Not everyone agrees. Bernard Donefer, Distinguished Lecturer in Information Technology in Financial Markets at Baruch College and Associate Director at Subotnick Financial Services Center, comments, "HFT itself is nothing more than what has already been done, just off the exchange floor and faster." Intuitively, one can question if HFT is necessary for allocating capital efficiently to manufacturing or service firms whose investment process has long time horizons, often in the range of years. On the other hand, the econometrician's view that financial price processes are continuous-time processes can only welcome a development that brings the reality of trading closer to the ideal of a continuous-time stochastic process.

Clearly there are different views and different interests. While HFTers identify and exploit profit opportunities and academics remark that market quality defined, for example, by the size of spreads, has improved, large institutional investors fear that they are paying a tribute to HFTers for keeping markets efficient.

This has lead to the creation of "dark pools," trading venues open only to specific classes of investors, for example, large institutional investors, where members can trade anonymously and with the expectation that any market inefficiency will ultimately profit themselves rather than being taken by intermediaries. Dark pools, estimated by sources to represent 7 to 8% of all U.S. equity trading, are themselves open to debate because of the lack of transparency.

In this section we will discuss the following issues:

- Is HFT a niche trading strategy or the future of equity markets?
- What phenomena do HFT strategies exploit to earn a profit?
- What is the impact of HFT on the price discovery process, on prices?
- What is the quality of the liquidity provided by HFT?
- What are the benefits of HFT?
- Does HFT introduce new risks?
- Is any new regulation needed to limit these risks?
- Who profits from HFT?

A. Niche Trading Strategy or the Future of Equity Markets

HFT, or the ability to exploit profit opportunities with trading strategies characterized by holding periods of a few minutes and without carrying positions overnight, is a recent phenomenon. However, the market conditions enabling HFT were created little more than a decade ago. As mentioned above, HFT was enabled by a combination of factors including the 2001 decimalization of U.S. equity markets, the advent of the electronic exchange, advances in computer and communications technology, the availability of more data, and new modeling techniques. These factors, combined with the objective of large institutional investors to optimize the trading of large orders, led to algorithmic trading. Algorithmic trading is based on computerized quantitative models and is used by large investors to reduce market impact. This is typically done by spreading large orders over many small transactions, thereby contributing to an increase in the volume of trading, a prerequisite for HFT. Algorithmic trading is not necessarily executed at high frequencies, but HFT is dependent on the development of algorithms. In addition, the ability to access directly the electronic book at the exchanges created new trading opportunities.

A representative from a major options exchange in the United States comments:

> The world of HFT would likely not exist in its present form if not for

> decimalization which allowed for finer pricing. When the market traded in 16ths, 8ths, spreads were very high; there was no capability to provide a better market. Since decimalization, the bid-ask spread has been reduced. This led to a reduction of the overall cost of access to stock or option prices. In the options market, this cost reduction has been multiplied thanks to penny stock trading.

Is HFT a niche market? The answer is two-pronged. On one side, HFTers are a small highly specialized type of trader characterized by the use of advanced information technology and modeling techniques and short time horizons. On the other side, HFTers cannot exist in isolation: They need a robust flow of trades as a main source of profit. HFT, as well as other market participants such as hedge funds, came into being to make a profit by exploiting regularities and inefficiencies in a flow of orders that already existed.

Different markets and different geographies have different populations of HFTers. The share of trades executed by HFTers depends on how HFTers are defined. It is widely accepted that in the U.S. equity market, HFT is responsible for 40 to 70% of all trades. In a study based on tick-by-tick data from Nasdaq and adopting a widely used definition of HFT, Brogaard (2010) finds that, in 2009, well above 70% of all trades can be attributed to HFTs. One source at a U.S. options exchange observes:

> Seventy percent is routinely accepted for market share of HFT in U.S. equity markets, but it depends on how you qualify participants. For example, market makers are intrinsically HFTers. In the equity options markets, I would put HFT market share at around 30 percent. Most HFTers in the options market tend to be very, very small because arbitrage opportunities are very small.

First developed in U.S. equity markets, HFT has now spread to other markets. The big players are present internationally, sources explained. However, HFTers' share of all trading in equity markets in Western Europe and Canada was estimated to be anywhere from one third to one half their estimated share of the U.S. equity market. A representative from a major North American exchange remarks, "The Canadian market has not been overwhelmed by HFT. I would estimate it to be 20–25 percent of all equity trading volume in Canada."

We asked participants if, as short-term arbitrage opportunities are exploited and disappear, HFT will also disappear. Professor Hautsch comments:

> There will always be a need to have a certain level of HF strategies, HFT to ensure efficiency. As for opportunities for statistical arbitrage, I believe that we will see the introduction of new instruments, new assets, new trading platforms. These will create micro arbitrage opportunities. It might be that in some markets, arbitrage opportunities will go to zero. But people will keep on using HFT, if not for micro arbitrage, to exploit optimal trade execution.

The representative of a large North American exchange comments, "We expect to see a blurring of lines between traditional players and HFTers as more traditional players access HF technology." We view this as blurring the lines between traditional assets managers and "quants," where the former have to some extent adopted quantitative methods for at least some parts of their investment management process.

B. Phenomena HFT Strategies Exploit to Earn a Profit

An important question, both from the practical and academic points of view, is what type of strategies HFTers use. As strategies are proprietary, there is very little direct knowledge of strategies employed. We can only make general comments and infer strategies from observing HFD. A first observation is that, given the speed of trading, HFT strategies are based on information that changes rapidly. Therefore, it is unlikely that these strategies are based on fundamental information on stocks or on macroeconomic data.

We can divide trading strategies at high frequency into three major categories. The first is based on trading on news, exploiting a time advantage in placing orders before the market reacts to news. This involves automatic text reading and analysis and modeling techniques that relate news to price movements.

The second type of trading strategy is based on revealing small price discrepancies between different markets or between different assets that should theoretically have the same price. Assuming that prices will realign rapidly, HFTers issue orders with low latency to exploit any arbitrage opportunity. This type of strategy is based on the ability to gather and analyze data, and then issue orders very rapidly before the market realigns. Exploiting arbitrage opportunities clearly entails assessing the cost of the trade that is about to be made. If the cost of a trade exceeds the size of the potential profit from arbitrage, then the trade is not executed. Wing Wah Tham, assistant professor of financial econometrics at the Erasmus School of Economics, observes, "Due to uncertainty in implementing trades, arbitrage strategies are not without risk even in the presence of arbitrage opportunities." Kozham and Tham (2010) use HFD to study the role of execution risk due to crowded trades in financial markets.

The third type of trading strategy is based on making short-term forecasts based on the econometric properties of data. The most likely econometric properties to enter into a HFT strategy are prices, trading volumes, and information related to past trades. A special type of forecast is based on knowledge of the flow of incoming orders. In fact, the knowledge that large orders are coming is a type of information that traders have always exploited to their advantage.

Trading based on the knowledge that large orders are coming is called "front running." If and how this knowledge can be acquired is a subject of debate. In the last 10 years, large long-term investors have invested in techniques to optimize the execution of large orders. As discussed above, one such technique, algorithmic trading, allows one to split large orders into a flow of small orders, thereby matching a flow of opposite orders and reducing market impact.

Secrecy is crucial to the success of algorithmic trading. If it is known in advance that a large order flow is coming, the benefits of algorithmic trading are reduced.

Large investors therefore dislike methods and techniques that reveal their order flow in advance. Barring any illegal disclosure of information, HFTers rely on issuing immediate-or-cancel orders to search for and access all types of undisplayed liquidity, including liquidity in dark pools. They do so in the space of milliseconds. This technique is called "pinging." Whether or not pinging should be banned (or somehow restricted) is now being debated.

In practice, strategies are implemented via trading rules that automatically issue orders when particular patterns of information are detected. While HFTers are often put into various categories, sources we interviewed remarked that the strategies used by HFTers have evolved over the years. A representative from a major North American exchange observes: "We see different strategies coming up. In the early stages, HFTers were mostly rebate takers, predatory. Now there is a more diverse range of strategies. Early adopters worked out inefficiencies in market; now there is the need for more effective strategies."

The perception from academia is similar. Professor Hautsch remarks, "It is hard to observe different strategies from raw data, but from conversations with HFTers, it is clear that over the past three, four years, strategies have changed dramatically."

Brogaard (2010) undertook a systematic exploration of HFT strategies based on tick data from the Nasdaq for 120 stocks for the period 2008–2009. He finds that most HFT strategies are based on short-term reversals. This opinion was shared by sources from academia and the exchanges that we interviewed. A source at a North American exchange observes, "HFTers do not use long-term mean-reverting models; they are looking for arbitrage on intra-day mean reversion. They are different from the market makers who take positions."

While little is known about the trading strategies adopted by HFTers, we do have information on a number of "stylized facts" about returns at very short time horizons, in particular, on the probability distribution of orders and the autocorrelation of orders at very short time horizons (see, for example, Dacorogna et al. 2001). However, HFTers work on strategies typically tested over periods of at most two years. While the broad lines of trading strategies are known, the details are proprietary. It is likely that hundreds of technical HFT rules are used and continuously adapted.

C. Impact of HFT on the Price Discovery Process and on Prices

The question of the impact of HFD on the price discovery process and on prices is a multifaceted question that is not easy to define theoretically. This is because it requires a comparison of the actual outcome with some hypothetical outcome in the absence of HFT. Nevertheless, there is a consensus that HFT impounds information faster and impacts some market parameters. Earlier studies analyzed the impact of decimalization on market quality (see, for example, Chakravarty et al. 2001 and Bessembinder 2003).

Terry Hendershott, at the Haas Finance Group at the University of California-Berkeley, observes, "If you consider the actual price as having fundamental information plus noise, HFD has no long-term fundamental information, but HFT can help get short-term information into prices faster."

Brogaard (2010) analyzes the impact of HFT on market parameters such as volatility and the bid-ask spread. He employs the now widely used methodology for analyzing market quality introduced in Hasbrouck (1993). In the sample that he analyzed (HFD from Nasdaq on 120 stocks for the period 2008–2009), he concludes that volatility did not increase and the bid-ask spread was reduced. On these points, there seems to be agreement. HFTers have not produced an increase in volatility, as many had feared, and have generally had a beneficial effect on parameters that define market quality such as the bid-ask spread.

One problem in analyzing the impact of HFT on the bid-ask spread is to separate the impact of HFT and that of decimalization and other changes introduced in the U.S. equity markets over the last decade and a half. An industry source, who confirms having seen a reduction in the bid-ask spread due to the activity of HFTers, remarks:

> If you look at the Canadian equity market, it's easier to separate the impact of HFT from that of decimalization. Decimalization was introduced in Canada in 1996 while HFT in Canada is relatively new, having started only as of late 2008–2009. It is possible to see a tightening of the spreads that occurred at the different time periods.

If we measure price efficiency in terms of parameters such as bid-ask spread, HFT has increased market efficiency. However, as HFTers trade against each other using algorithms that are in general based on technical rules that have nothing to do with fundamentals, we can ask if HFT might cause prices to depart from fundamentals. James MacIntosh, investment editor of the *Financial Times*, remarks that fundamental information is no longer reflected in stock pricing (see MackIntosh 2010). He suggests that pricing is now driven by market sentiment and possibly by the increase in trading on trends and patterns.

One market fact that can possibly be ascribed to HFT is the observed increase in correlation. Professor Voev comments:

> There is recent evidence that HFT is leading to more correlation, a fact that has serious implications for diversification. This is making it more difficult to diversify with index tracking or exchange-traded funds. There are now thousands of algos trading indexes, moving prices. Is price momentum dominated by traders trading indexes?

Professor Bauwens comments that while HFT has improved market efficiency overall, there is the possibility that it can cause artificial price trends:

> Finance theory holds that prices reflect past information but is not precise on how this works. My conjecture is that HFT has in most cases increased the speed at at which prices adjust to reflect new information; thus, it has led to increased efficiency. However, it has also been noted that correlation between intraday returns of stocks has increased without apparently much

> reason, and this may be caused by HFT driven by econometric models disconnected from fundamentals.

The action of HFTers has probably reduced volatility. Nevertheless, some sources mentioned that while volatility is down in normal times, HFT might lead to volatility spikes. Professor Voev remarks:

> We now have faster channels of market fear, uncertainty. Is HFT causing this or is it just a question of faster channels, with HFT facilitating fast channeling of emotions, fear? In normal times, HFT brings smoother adjustment to new levels versus discrete moves which are more volatile. But in more extreme circumstances, it can lead to spikes in volatility.

Commenting on the impact of HFT activity on volatility, an industry source says, "It (is) hard for us as an exchange to evaluate the impact of HFT on markets. HFT has probably had a dampening effect on volatility as the bid-ask spread is constantly narrowing except when all the HTFers turn off their computers. HFTers don't try to make their models fit beyond mean returns."

D. More (or Better) Liquidity with HFT?

It is widely held that HFT provides liquidity to equity markets. However, HFT per se provides liquidity only for a very short time. By the nature of their business, HFTers buy and sell at high frequency. If they do not find a counterparty for a trade in a matter of seconds, orders are cancelled. These are the (in)famous flash trades. Among the academics and industry players we interviewed, opinions were divided as to the nature of liquidity provided by HFTers. Some argue that liquidity provided by HFT is exercisable liquidity; those who question the benefit of HFT liquidity point to its fleeting quality.

Among those defending the utility of the short-term liquidity provided by HFT, the representative of a major North American exchange asks, "Is the liquidity provided by HFT real or phantom? It is tough to answer this given the different strategies employed by HFTers, but it is exercisable liquidity, available for someone to hit, even if it is only there for a short period. Certainly it is real if you have the technology to grab it."

Another industry source took the opposite position, arguing:

> HFT does add liquidity on a very shallow basis on narrow prices for small amounts and for pure retail customers. It is like a discount store that sells handbags at a low price but has only one handbag around to sell. HFT is less a provider of liquidity for larger volumes. Liquidity provided by HFTers is not deep enough, it is fleeting.

Professor Spatt suggests that the nature of today's liquidity is a reflection of changes in trading behavior. He comments:

> The question of traders showing their hands versus HFTers coming out for brief periods of time is the question of how to engage to obtain liquidity.

> The types of tactics used by HFTers leads to cancellation rates that keep exploding. Most orders are now cancelled almost instantaneously. It is not a question of being manipulative; HFTers are just trying to understand the liquidity out there and scale up and trade against it. HFTers (are) also looking for a lack of liquidity. Liquidity provided by HFTers is not an illusion, but it is different from the usual liquidity. The old notion was that traders want everyone else to show their hands without showing their own hand but it does not work that way. You cannot mandate liquidity. You must make it attractive for people to show their hands without the fear of being picked off. If a trader shows impatience, he or she will not get a good price.

E. Do Markets Benefit from HFT?

We discussed above several widely ascribed, but not universally acclaimed, benefits of HFT to equity markets (i.e., a lowering of the bid-ask spreads, reduced volatility, and increased albeit short-term liquidity). However, not everyone agrees. Professor Jagannathan suggests that the benefits of HFT have perhaps not been sufficiently or correctly studied:

> The relative benefit if all trading once at the end of day as opposed to HFT has not been established. When people say markets are better off because of HFT, no one has correctly measured this against benefit of trading at a lower frequency. Think about it. Suppose I know that something is happening and trade. My trade will affect the price at a point in time. Does it really matter if I know the price at exactly the minute rather than at the end of the day? At the fundamental level, HFT will not make us much better off.

Angel et al. (2010) perform a detailed analysis of changes in equity trading over the last 10 years. They conclude that the market quality has improved. But James Angel, co-author of the study and associate professor of finance at Georgetown University's McDonough School of Business, questions if pushing trading ever faster produces a real benefit:

> Market-makers buy on a dip and sell on a rebound. They have made it easier for the long-term investor to trade at lower costs. Cost reductions were realized as computers replaced humans as market-makers. No one would say that pure market-makers have hurt the investor. But how much benefit is there if pricing is made more accurate in seconds as opposed to in minutes? It is debatable.

Professor Spatt comments that the current environment has promoted more competition in the equity markets and that the competition has been beneficial. But he suggests that there is not enough competition in other markets. In particular, he observes that there is inadequate attention on the bond market microstructure.

One benefit that the equity exchanges have seen is increased attention being paid to listed firms, at least the larger of the listed firms. A representative from a major North American exchange remarks:

> The net benefit is that we have a better market with the participation of HFTers. HFTers' entry into the Canadian market led to an influx of new participants in the exchange. As a result there is a diversification of the order flow and of trading strategies. Previously, in Canada, there was a concentration of market participants. A knock-on effect is that, as big names in the U.S. set their sights on Canada, others opened their eyes and began to look at the Canadian market. As liquidity improves, as trading velocity grows, the increased activity on listed shares means that firms that were before screened out by filters that screen out stocks that trade less than 1 million shares a day are now traded. There is a benefit for the firms as this gives them greater access to capital, lowers the cost of capital. What happens on an intraday basis does not have a material impact on the long-term investor if not when the investor wants to get into the market. And when the long-term investor wants to get into the market, he/she finds a buyer/seller. Speculators facilitate the trade; they are a necessary element of the market place.

It might be, however, that the activity of HFTers is keeping some investors away from the equity markets. Spicer (2010) refers to data released in the beginning of September 2010 that show that flows have exited U.S. mutual fund accounts in every week since the May 6th flash crash. He writes that these outflows are fueling speculation that the crash continues to undermine investor confidence. Fabozzi, Focardi, and Jonas (2010) remark that following the 2007–2009 market turmoil, regaining investor confidence is the biggest challenge for all in the financial services industry. Retail investors have seen strong market movements without any fundamental reason for the ups and downs. According to sources for that study, such movements are reinforcing people's perception that markets are casinos and an inappropriate placement for one's savings.

Nevertheless, Professor Jagannathan believes that, if market participants are uneasy about trading in venues where HFTers are active, they can trade elsewhere: "HFTers can trade among themselves and this might keep investors away. People could invent other markets, for example, you could have one auction a week much as the old Dutch auction system. If the activity of HFTers gets really bad, people will invent other things such as dark pools; it is an easy thing to fix."

F. Does HFT Introduce New Market Risks?

Generally speaking, there is little understanding of the highly secretive strategies used by HFTers. A representative of a U.S. options exchange comments:

> If a HFTer does pure arbitrage and is not predatory, not manipulative, there is no problem. The problem is that we do not know. The SEC is now requesting all exchanges to identify HFTers by some formula, for example,

> more than 399 trades/day and to tag trades for analysis. From the exchange's standpoint, it is not possible to tell what the trader is doing as he/she might be doing something in other markets, exchanges. It is hard to tell an elephant from touching one part of the body.

One problem is that data that have been collected by the regulators have not helped to elucidate trading practices. Professor Donefer notes:

> The problem is that regulators have been running at their studies on players, for example, broker-dealers, hedge funds, etc. FINRA [the largest independent securities regulator in the U.S.] has no clue as to the kind of trading being done and the strategies behind it. Regulators should require tagging of orders by algos as opposed to by category of players.

To our knowledge, academic studies have not revealed any evidence of dubious practices by HFTers such as "front running," a strategy based on anticipating the arrival of large orders. The (probabilistic) knowledge of the arrival of large orders is in itself obtained through other practices such as "pinging," which consists of issuing and cancelling orders in the space of a few milliseconds in order to reveal pools of existing liquidity. Nor, to our knowledge, have academic studies produced evidence of market manipulation.

Addressing the question of new risks introduced by HTF, Professor Hendershott remarks:

> I am not sure that we have any evidence so far of new risks, but that does not mean it could not happen. Is the fear that algos create prices causing people to not understand what is the correct price in the market, either intentionally or unintentionally? If someone is causing prices to move in a way as to not reflect information, others can trade against them and make money.

On the other hand, sources agreed that new risks related to technology and speed have been introduced. Professor Angel remarks, "The high-speed world might produce some high-speed risks." HFT can ultimately be described as fast machines trading against other fast machines. Professor Angel adds:

> I do not think HFT makes it easier to manipulate the market. Games to manipulate the markets have been going on for 400 years. If anything, it is now harder to manipulate the market. But the big problem is markets act so quickly now. Can something go wrong? Yes, consider, for example May 6 (2010). There are various risks, such as run-away algos, computer failures, intentional hacking, programming problems. Yes, the system is vulnerable to breakdown, to attack. So you need to have something in place to respond as quickly as possible when computers crash, for example, circuit breakers, for when machines malfunction.

Persons we interviewed believe that the biggest problem with HFT is the possibility of cascading effects (not the creation of bubbles) or system collapse due to the high speed of trading or an excessive number of messages. Professor Donefer, who developed his argument in an article recently published in *The Journal of Trading* (2010), remarks:

> HFT and direct market access represent an additional risk in that all strategies that track markets are pegged to NBBOs. Imagine that one algo goes wild. All other markets see this, reset their prices, and there is a cascading effect. There are too many models based on the same information, too many crowded trades.

Relative to cascading effects, Professor Voev comments:

> When you have computers programmed to trade on price patterns, you might have avalanche effects. Automatic trading can push prices way too low. If markets are efficient, the price bounces back to fundamental values. But in some cases prices do not bounce back because there is general market uncertainty and no one knows what the price should be.

In this sense, protecting the system is more a question of intelligent design of trading than the issuing of rules banning this or that process. Referring to the use of rule-based trading algorithms, Professor Jagannathan comments: "Anything that is mechanical, rule-based, needs oversight rules. Things change as you go along — portfolio insurance, the May 6 flash crash — and you need intelligent rules for trading. If there is a large change in the price, rules should be in place to handle such situations."

Sources pointed to the flash crash of May 6, 2010, when the Dow Jones Industrial Average lost some 700 points before sharply rebounding in the space of just 20 minutes, to argue that the presence of HFTers likely helped the markets bounce back rapidly. Professor Donefer remarks:

> If you look at the flash crash of October 1987, there were market-makers but people walked off the floor, and those that did not risked bankruptcy. Greenspan was just in as head of the Federal Reserve, and ordered the banks to lend money to market-makers to keep them solvent, to help the markets recover. It took one year for markets to recover from that crash. With the flash crash of May 6th and the presence of statistical arbitrageurs, HFTers, the market recovered in matter of less than one day as these people got back into the market. When markets start to crash, risk models take over if the firm's jeopardy is at stake. These firms are no longer the family businesses such as those in the 1987 crash, but corporations. They use more sophisticated risk models. If they see too much capital at risk, they walk away from the markets. But they come back minutes later when profit opportunities are identified. I have no first-hand knowledge of what happened but my perception is that among the players in the May 6th flash crash, there were high-frequency market-makers as Getco, Virtu, and Knight Capital. They all came back into the market right away.

In addition to the risk of cascading effects or technology-related risks due to the speed and messaging typical of HFT activity, sources identified other risks such as increased correlation. Professor Hautsch observes, “HFTers try to exploit statistical arbitrage. This leads to greater correlations across markets, assets, instruments. In turn, diversification effects are weakened, leading to increased risk. Greater efficiency is a good thing but more correlation is a risk: Many nice portfolio models don’t work anymore.”

G. Is New Regulation Needed to Limit These Risks?

Though sources agreed that HFT has introduced new risks related to technology, there was no consensus as to how exchanges or regulators should respond. Some sources were in favor turning off the quant models and keeping only the market-makers or end buyers/sellers going; others suggested the use of circuit breakers. Commenting after the May 6th flash crash and the regulators’ move to bust trades when prices moved far from their value, Professor Angel remarks, “Markets can get into situations, chaotic events in which an algo can push a price far from its value. I favor circuit breakers and then switching to a different market mechanism, shutting all computers as is done at the Deutsche Boerse and then starting all over the morning after with an auction.”

However, not all our interviewees were in favor of circuit breakers. Professor Spatt argues against circuit breakers as they are disruptive of the trading process but is in favor of filters to catch mistakes. Professor Spatt is concerned about the risks created by intervention:

> May 6th was a fiasco but one risk now created is that liquidity won’t arrive because of a lack of clarity in the process given the regulator’s decision to cancel trades whose price movement was more than 60% while trades whose price movement was under 60% were not canceled. People are not under an obligation to keep providing liquidity and will pull back if they don’t understand what the regulator’s response to a situation will be.

On October 1, 2010, the SEC released its report on the May 6th flash crash. The report attributes the crash to a cascade effect following an unusually large trade ($4.1 billion). Two observations can be made.

1. It has been well established that intraday returns are fat-tailed, as are the size of trades and indeed the capitalization of firms. In consequence, one should expect fat-tailed returns even in the absence of cascading effects. As pointed out by our interviewees, the rapid recovery of markets after initial losses provides a positive evaluation of the robustness of the system.

2. Cascading effects can occur again, as our interviewees remarked. However, avoiding cascading or limiting its effects is a question of system design. It might be a very difficult objective to achieve with regulation.

One area of consensus on the need to regulate was on sponsored access. Sponsored (or naked) access gives trading firms using brokers' licenses unfettered access to stock markets. The Boston-based research firm Aite estimates that by 2009 38% of all U.S. stock trading was done by firms using sponsored access to the markets. The fear is that naked access — typically without (adequate) validation of margins — via direct market access may create strong short-term price movements up or down and liquidity crashes.

Professor Hautsch comments, "The problem is not just HFT or direct market access (DMA) but a combination of this together with high leverage, stop orders, naked access, etc. But this does not product bubbles. In normal times, naked access is not a problem but in non normal times, if all the effects come together, it can produce a cascading effect. What is missing is a warning system."

Most sources expect the SEC to act soon on restricting naked access.

H. Who Profits from HFT?

As to who profits from HFT, a first answer, of course, is that HFTers profit from HFT. Early estimates by the Tabb Group put HFT profits in the U.S. equity markets for 2008 at $21 billion, but the figure was subsequently revised downward to $7–9 billion. Perhaps coincidentally, the earlier figure is what Kearns et al. (2010) estimated to be the maximum that an omniscient HFTer could earn on the U.S. equity markets. Nevertheless, it was reported that Citadel realized a $1 billion profit from HFT in 2007.

If the $7–9 billion estimated profits for HFT is close to reality, global profit opportunities on U.S. equity markets appear to be relatively small, but this number should not be surprising: Ultimately, HFT exploits small inefficiencies left after major trends have been exploited. HFT requires very liquid markets. Irene Aldridge, managing partner of Able Alpha Trading LTD, a proprietary firm specializing in HFT, writes that HFT is not profitable in illiquid markets (2010a).

There is some expectation that HFT will be less profitable in the future, at least in U.S. equity markets. Professor Angel remarks:

> Basic statistical arbitrage trading strategy is simple, straight forward, so it is a cut-throat commodity business. To survive, you must be a low-cost producer and do it in scale. There is a lot of competition out there as anyone can buy a computer — they are fairly cheap. The intense competition has pushed margins down to almost zero. HFT will not go away but we will see a shake-out of the less efficient, less intelligent players.

As U.S. equity markets become more efficient thanks to tick-by-tick HFT strategies, sources expect that the diminished returns will see HFTers looking for other sources of profits, including the extension to other asset classes, options markets, and dark space.

Is HFT a zero-sum game in which the HFTers profits are gained at the expense of other, more slow-moving traders? On her web site Aldridge (2010b) writes, "While no institution thoroughly tracks the performance of high-frequency funds, colloquial evidence suggests that the majority of high-frequency managers delivered positive returns in 2008, while 70% of low-frequency practitioners lost money, according to *The New York Times*."

Others suggested that HFTers may have taken the place and profits of other players, such as the market-makers and investment banks. Professor Hendershott comments, "It is possible that HFT firms are not causing a change in the amount of trading profit but are taking the profit for themselves. For example, market-makers and banks used to make about $5 billion a year and now this figure is zero or close to zero."

The exchanges themselves stand to raise transactional and other revenues as they gear up to support HFTers with high-speed computers and communications and co-location facilities. A source at a major North American exchange comments, "Co-location is a very strong source of revenues, customer loyalty, and stickiness." But the revenues come at a cost: The exchanges are beefing up their investment in technology to meet the needs of HFTers.

It is enormously expensive for an exchange to support HFT. Exchanges need to constantly upgrade their architecture to process more messaging. According to industry sources, it is not uncommon for HFTers to send more than one million messages a day and trade only a few contracts. One source comments:

> From a technological point of view what is needed is having the required robustness, constantly upgrading from one gigabyte to 10 gigabyte lines, more and more powerful servers, faster speeds, next generation of computers. But next generation architecture is more and more expensive. We are moving towards software to eliminate latency in the computer reading the software code. Software-on-a-chip servers are priced at $100,000 versus $5,000–7,000 for today's servers. Today we are processing orders at 500 microseconds but racing to do so at single-digit microseconds.

The race for speed has also benefited technology suppliers. One North American source observes, "We have seen a proliferation of technology vendors — hardware, software, middlewear, smart order systems, security… The number of technology suppliers around has tripled over the last 12–18 months."

Sources from the exchanges also identified benefits for firms listed on the exchange. As mentioned above, at least one exchange evaluates that the activity of HFTers has brought more investors to the exchange's listed firms, thereby increasing their access to capital and reducing its cost.

Nevertheless, there is concern that the activity of HFTers is concentrated on a small number of stocks. A representative from a U.S. exchange observes, "We have seen a greater concentration [of trades] in the last two years than in the last 10 years. It is very dangerous for an exchange when there is so much interest in few names, when all investments concentrated around a few names. We lose flexibility."

For the investor at large, retail, or institutional, the benefits are not so clear.

While most sources believe that the cost of trading and bid-ask spreads have been reduced by the activity of HFTers, there is to our knowledge no study that factors in the cost of exchange infrastructure needed to service HFTers and how this cost affects the total price of trading. Professor Hendershott comments:

> A most legitimate concern outside of manipulation is the over investment in technology, for example, end users of assets as Vanguard, Fidelity want to find each other and trade directly. The question is: Is the system such that whatever the end user does, he/she finds a HFTer on the other side of the trade? So instead of selling to another end user, the investor sells to an HFTer which in turn sells to another end user. This would be a bad thing as trading would become more costly and, normally, a buy/sell transaction should be mutual. …HFTers takes some slice; we can try to get around this with dark pools, for example, a call-auction that occurs once a day. It would reduce the role of the HFTers.

IV. CONCLUSION

In this paper, we analyzed high-frequency trading (HFT) and its econometric foundation based on high-frequency data. From this analysis it is possible to argue that HFT is a natural evolution of the trading process, enabled by advances in computer and communications technology and a high-frequency flow of trades due to algorithmic trading by long-term investors. High-frequency traders (HFTers) employ computerized algorithms and fast computers and communications channels to exploit this "raw material."

Empirical analysis has shown that the presence of HFTers has improved market quality in terms of lowering the cost of trading, adding liquidity, and reducing the bid-ask spreads. This improvement in market quality comes at a cost as HFTers make a profit, albeit not a very large profit, as a percentage of trading volume.

Given the short-time nature of HFT and the fact that positions are typically not carried overnight, the potential for market manipulation and for the creation of bubbles and other nefarious market effects seems to be modest. The problems posed by HFT are more of the domain of model or system breakdown or cascading (typically downward) price movements as HFTers withdraw liquidity from the markets. The former poses a challenge of the design of electronic trading facilities. As for the second, solutions have been proposed including slowing down or interrupting the trading process or changing the trading mechanism.

References

Aït-Sahalia, Y. and Mykland, P.A., 2003, The Effects of Random and Discrete Sampling When Estimating Continuous-Time Diffusions. *Econometrica,* 71, 483–549.

Aït-Sahalia, Y., Mykland, P., and Zhang, L., 2005, How Often To Sample A Continuous-Time Process In The Presence Of Market Microstructure Noise. *Review of Financial Studies*, 18, 351-416.

Alridge, I., 2010a, *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems* (John Wiley & Sons, Inc., Hoboken, NJ).

Alridge, I., 2010b, How Profitable Are High-Frequency Trading Strategies? <http://www.ablealpha.com/aaindex.php?option=com_content&view=article&id=132:how-profitable-are-high-frequency-strategies-by-irene-aldridge&catid=1:latest-news&Itemid=50>.

Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P., 2001, The Distribution of Exchange Rate Realized Volatility. *Journal of the American Statistical Association*, 96, 42–55.

Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P., 2003, Modeling and Forecasting Realized Volatility. *Econometrica*, 71, 579–625.

Andersen, T.G., Bollerslev, T., and Meddahi, N., 2002, Correcting the Errors: A Note on Volatility Forecast Evaluation Based on High-Frequency Data and Realized Volatilities. *Working paper, University of Montreal*.

Angel, J., Harris, L. and Spatt, C.S., 2010, Equity Trading in the 21st Century. *Working paper FBE 09-USC, Marshall School of Business.*

Bandi, F.M., and Phillips, P.C.B., 2003, Fully Nonparametric Estimation of Scalar Diffusion Models. *Econometrica*, 71, 241–283.

Bandi, F.M., and Russell, J.R., 2006, Separating Microstructure Noise from Volatility. *Journal of Financial Economics,* 79, 655-692.

Bandi, F.M. and Russell, J.R., 2008, Microstructure Noise, Realized Variance, and Optimal Sampling. *Review of Economic Studies*, 75, 339-369.

Bandi, F.M., Russell, J.R., and Zhu, Y., 2008. Using High-Frequency Data in Dynamic Portfolio Choice. *Econometric Reviews*, 27, 163-198.

Barndorff-Nielsen, O.E. and Shephard, N., 2002a, Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models. *Journal of the Royal Statistical Society B*, 64, 253-280.

Barndoff-Nielsen, O.E. and Shephard, N., 2002b, Estimating Quadratic Variation Using Realized Variance. *Journal of Applied Econometrics* 17, 457-477.

Barndoff-Nielsen, O.E. and Shephard, N., 2004, Econometric Analysis of Realised Covariation: High Frequency Based Covariance, Regression and Correlation in Financial Economics. *Econometrica*, 72, 885-925.

Bauwens, L. and Giot, P., 2001, *Econometric Modelling of Stock Market Intraday Activity* (Kluwer Academic Publishers, Dordrecht).

Bauwens, L. and Veredas, D., 2004, The Stochastic Conditional Duration Model: A Latent Variable Model for the Analysis of Financial Durations. *Journal of Econometrics,* 119, 381-412.

Bauwens, L. and Hautsch, N., 2006a, Modelling Financial High Frequency Data Using Point Processes. *Discussion paper 2006-39, Catholic University of Louvain.*

Bauwens, L. and Hautsch, N., 2006b, Stochastic Conditional Intensity Processes. *Journal of Financial Econometrics*, 4, 450-493.

Bertrand, E., 2009, Technology Song and Dance. *Presentation at Union of Arab Stock Exchanges Conference, Beirut, NYSE Technologies*, November.

Bessembinder, H., 2003, Trade Execution Costs and Market Quality after Decimalization. *Journal of Financial and Quantitative Analysis*, 38, 747-777.

Boehmer, E., Grammig, J., and Theissen, E., 2006, Estimating the Probability of Informed Trading: Does Trade Misclassification Matter? *Journal of Financial Markets*, 10, 26-47.

Brogaard, J.A., 2010, High Frequency Trading and its Impact on Market Quality. *Working paper, Kellogg School of Management, Northtwestern University School of Law.*

Brownlees, C.T. and Gallo, G.M., 2006. Financial Econometric Analysis at Ultra-High Frequency: Data Handling Concerns. *Computational Statistics & Data Analysis*, 51, 2232 – 2245

Brownlees, C.T. and Gallo, G.M., 2007, Comparison of Volatility Measures: A Risk Management Perspective. *Working paper, University of Florence*.

Cipollini, F., Engle, R. F., and Gallo, G. M., 2006, Vector Multiplicative Error Models: Representation and Inference. *NBER Working Paper Series 12690, National Bureau of Economic Research.*

Clark, C., 2010, Controlling Risk in a Lightning-Speed Trading Environment. *Chicago Fed Letter*, March.

Dacorogna, M.M., Gencay, R., Muller, U.A., Olsen, R., and Pictet, O.V., 2001, *An Introduction To High Frequency Finance* (Academic Press, London).

Donefer, B.S., 2010, Algos Gone Wild: Risk in the World of Automated Trading Strategies. *The Journal of Trading*, 5, 31-34.

Duhigg, C., 2010, Stock Traders Find Speed Pays, in Milliseconds. *The New York Times*, 24 July, A1.

Engle, R.F. and Russell, J.R., 1998, Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data. *Econometrica*, 66, 1127-1162.

Engle, R.F., 2002, New Frontiers for ARCH Models. *Journal of Applied Econometrics*, 17, 425-446.

Fabozzi, F.J., Focardi, S.M., and Jonas, C., 2010, *Investment Management After the Global Financial Crisis* (The Research Foundation of CFA Institute, Charlottesville, VA).

Falkenberry, T.N., 2002, High Frequency Data Filtering: A Review of the Issues Associated with Maintaining and Cleaning a High Frequency Financial Database. *TickData Technical Report*.

Friedman, M., 1953, The Methodology of Positive Economics. In *Essays in Positive Economics* (University of Chicago Press, Chicago, IL).

Ghysels, E., Santa-Clara, P., and Valkanov, R., 2006, Predicting Volatility: Getting the Most out of Return Data Sampled at Different Frequencies. *Journal of Econometrics*, 131, 59-95.

Hansen P.R., Lunde, A., and Voev, V., 2007, Integrated Covariance Estimation in the Presence of Noise. *Journal of Financial Econometrics*, 5, 68-104.

Hasbrouck, J., 1993, Assessing the Quality of a Security Market: A New Approach to Transaction-Cost Measurement. *Review of Financial Studies*, 6, 191–212.

Hautsch, N., 2004, *Modelling Irregularly Spaced Financial Data – Theory and Practice of Dynamic Duration Models.* Lecture Notes in Economics and Mathematical Systems, vol. 539 (Springer-Verlag, Berlin).

Kearns, M., Kulesza, A., and Nevmyvaka, Y., 2010, Empirical Limitations on High Frequency Trading Profitability. arXiv:1007.2593v2

Kozhan, R. and Tham, W.W., 2010, How Riskless is 'Riskless' Arbitrage. *Working Paper, January 21.* Available at SSRN: http://ssrn.com/abstract=1540693.

Lo, A. W., and MacKinlay, A. C., 1990, An Econometric Analysis of Nonsynchronous Trading. *Journal of Econometrics*, 45, 181-211.

McAleer, M. and Medeiros, M.C., 2008, Realized Volatility: A Review. *Econometric Reviews*, 27(1–3), 10-45.

MacIntosh, J., 2010, *Financial Times Fund Management,* August 23.

Oomen, R.C.A., 2006, Properties of Realized Variance under Alternative Sampling Schemes. *Journal of Business Economic Statistics*, 24, 219–237.

Pacurar, M., 2008, Autoregressive Conditional Duration (ACD) Models in Finance: A Survey of the Theoretical and Empirical Literature. *Journal of Economic Surveys,* 22, 711-751.

Spicer, J., 2010, SEC Says Flash Crash Report Will Bring Confidence. *Reuters*, September 28.

Voev, V. and Lunde, A., 2007, Integrated Covariance Estimation Using High-Frequency Data in the Presence of Noise. *Journal of Financial Econometrics*, 5, 68-104.

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 the American Statistical Association*, 100, 1394-1411.