

REVIEW OF FUTURES MARKETS

Editor

Mark E. Holder

Advisory Board

Patrick J. Catania, Chair
Educational Research Foundation
Chicago, IL

Philip McBride Johnson
Skadden Arps
Washington, D.C.

Alger B. Chapman
The Cambridge Group
Chicago, IL

James Oliff
CME Group, Inc.
Chicago, IL

Frank J. Fabozzi
Yale University
New Haven, CT

Jay S. Sorkin
CME Group, Inc.
Chicago, IL

Michael Gorham
Illinois Institute of Technology
Chicago, IL

William G. Tomek
Cornell University
Ithaca, NY

Editorial Board

Charles Cao
Pennsylvania State University
University Park, PA

Paul Kofman
University of Melbourne
Melbourne, Australia

Eric C. Chang
University of Hong Kong
Hong Kong

Peter Locke
Texas Christian University
Fort Worth, TX

K. C. Chen
Cal State University
Fresno, CA

Robert McDonald
Northwestern University
Evanston, IL

Alex Frino
University of Sydney
Sydney, Australia

Sotiris Staikouras
Cass Business School
London, UK

Joseph K.W. Fung
Hong Kong Baptist University
Hong Kong

Chris Veld
University of Glasgow
Glasgow, Scotland, UK

William K.H. Fung
London Business School
London, UK

Robert I. Webb
University of Virginia
Charlottesville, VA

Stewart D. Hodges
University of Warwick
Coventry, UK

Jeffrey Williams
University of California-Davis
Davis, CA

David A. Hsieh
Duke University
Durham, NC

Taewoo You
Myongji College
Seoul, Korea

Assistant Editor

Mary K. Aaby

Editorial office: Review of Futures Markets, Department of Finance, P.O. Box 164, Massillon, Ohio 44648 Ohio USA. E-mail: rfm1@rfmjournal.com Website: www.rfmjournal.com.

REVIEW OF FUTURES MARKETS

VOLUME 19 SPECIAL ISSUE

Made Possible by Grants from the Institute for Financial Markets

Copyright: *Review of Futures Markets* (ISSN #0898-011X) is published quarterly, one volume per year.

Subscriptions: \$80 for individuals, \$300 for institutions. Payment must be made in U.S. dollars drawn on a U.S. bank. Subscription inquiries and changes of address should be addressed to Subscriptions, Review of Futures Markets, P.O. Box 164, Massillon, Ohio 44648 USA.

The journal is printed by The Sheridan Press, 450 Fame Avenue, Hanover, Pennsylvania 17331.

Review of Futures Markets is published in cooperation with The Institute for Financial Markets, 2001 Pennsylvania Ave. NW, Washington, DC.

Copyright ©*Review of Futures Markets*. All rights reserved. No part of this publication may be reproduced in any form by any means without the prior written permission of the publisher except as permitted by Sections 107 and 108 of the United States Copyright Law governing copies of articles in this journal made for teaching and research purposes. Address permission requests to Review of Futures Markets, P.O. Box 164, Massillon, Ohio 44648. USA. For more information, contact rfm1@rfmjournals.com.

Manuscripts: All manuscripts should be submitted electronically to rfm1@rfmjournals.com. Guidelines for submission are located at the end of this publication.

Opinions expressed by the authors are their own and do not necessarily represent those of the editors or publisher.

NOTE: The *Review of Futures Markets* has a new website address: www.rfmjournals.com. The new contact information for the journal office is P.O. Box 164, Massillon, Ohio 44648; e-mail rfm1@rfmjournals.com; telephone 330-833-8582.

Editor Mark Holder may be contacted at mholder@kent.edu.

Contact Assistant Editor Mary Aaby at mkaaby@rfmjournals.com.

Table of Contents


Introduction to Special Issue <i>Trish Foshée</i>	5
High-Frequency Trading: Methodologies and Market Impact <i>Frank J. Fabozzi, Sergio M. Focardi, and Caroline Jonas</i>	7
Clearing House, Margin Requirements, and Systemic Risk <i>Jorge A. Cruz Lopez, Jeffrey H. Harris, and Christophe Pérignon</i>	39
Would Price Limits Have Made any Difference to the “Flash Crash” on May 6, 2010? <i>Bernard Lee, Shih-fen Cheng, and Annie Koh</i>	55
Direct Market Access in Exchange-Traded Derivatives: Effects of Algorithmic Trading on Liquidity in Futures Markets <i>Ahmet K. Karagozolu</i>	94

The Industry Resource for Education and Data

The Institute for Financial Markets is an independent, nonprofit foundation. Our mission is to provide unbiased information and education to increase the public understanding of the importance of financial markets to the global economy, and to improve the technical competence of those in the industry.

We provide educational books, courses and web-training to market-participants, as well as, U.S. and international regulators, government agencies, public policy-makers and multi-country member organizations. In addition, the IFM maintains a number of one-of-a-kind data storehouses on exchange trading fees, position limits, contract specifications and volume. These global industry utilities provide market users and firms the information they need to conduct business and to meet their regulatory responsibilities.

Annually, the IFM funds objective studies on high-priority issues; publishes the Review of Futures Markets; and supports the annual Asia-Pacific Futures Research Symposium — all to expand the independent literature of futures, options, and derivatives research.



THE INSTITUTE FOR
FINANCIAL MARKETS

www.theIFM.org | 1.202.223.1528 | info@theIFM.org

THE IFM RELEASES SPECIAL ISSUE OF REVIEW OF FUTURES MARKETS

During this unprecedented period of transformation in global markets and financial reforms in the United States and around the globe, the IFM is pleased to present a special issue of the *Review of Futures Markets*. The research studies in this issue address high-priority issues critical to the trading and clearing of derivatives worldwide. Funding for the studies was made possible by a generous \$2 million gift from The Clearing Corporation Charitable Foundation that established an IFM endowment to fund futures and options research and education.

IFM's decision to fund the independent studies included in this issue were based on the character of the problem, the implications of the proposed research, and the anticipated contribution of the research and findings to derivatives literature — and above all, research that fits within the IFM's nonprofit mission* to provide quality balanced information. Solicitation for researchers was completed through a Call for Papers in 2010. The studies were then selected from the proposal received by the IFM and its Grant Advisory Committee, a group of industry professionals with a deep acumen in the derivatives business.

The findings in this issue were subjected to a rigorous peer-review process that encourages authors to meet quality standards and to avoid the dissemination of unwarranted findings, superfluous claims or interpretations, and personal views. We believe the research contained in this edition can help educate market users, policy makers, regulators, academics, and other stakeholders, while building public understanding and confidence in exchange-traded markets.

We hope you enjoy this complimentary issue of the *Review of Futures Markets*, and we welcome your comments.



Trish Foshée
Executive Vice President

Institute for Financial Markets
2001 Pennsylvania Ave. NW, Suite 600
Washington, DC 20006
www.theIFM.org

*The IFM's mission is to increase the technical competence of those in the global financial markets and the financial services industry through research, educational publications, and industry services. Our vision is to expand the knowledge and skills of individuals, while building public understanding and confidence in the markets. The Institute for Financial Markets, founded in 1989, is a Section 501(c)(3) nonprofit educational foundation. It has no membership and does not engage in any lobbying or political activities.

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)HFD 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 event² 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: $e^{-\lambda} \frac{\lambda^k}{k!}$

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 uncertainty 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 led 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 HFT 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 MacIntosh 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 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 HFTers 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 HFT, 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 also 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-aldrige&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.
- Barndorff-Nielsen, O.E. and Shephard, N., 2002b, Estimating Quadratic Variation Using Realized Variance. *Journal of Applied Econometrics* 17, 457–477.
- Barndorff-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, Northwestern 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.

CLEARING HOUSE, MARGIN REQUIREMENTS, AND SYSTEMIC RISK

Jorge A. Cruz Lopez, Jeffrey H. Harris, and
Christophe Pérignon*

Margins are the major safeguards against default risk on a derivatives exchange. When the clearing house sets margin requirements, it does so by only focusing on individual clearing firm positions (e.g., the SPAN system). We depart from this traditional approach and present an alternative method that accounts for interdependencies among clearing members when setting margins. Our method generalizes the SPAN system by allowing individual margins to increase when clearing firms are more likely to be in financial distress simultaneously.

Recent turmoil in financial markets has heightened the need for well-functioning clearing facilities in derivatives markets, particularly when large market participants are in financial distress and eventually default (Acworth 2009; Pirrong 2009; Duffie and Zhu 2010). In a derivatives exchange, the clearing house is responsible for the clearing function, which consists of confirming, matching, and settling all trades. The clearing house operates with a limited number of clearing firms or futures commission merchants, which are private firms that have the right to clear trades for themselves (i.e., proprietary trading), for their own customers, and for the customers of non-clearing firms.¹

1. While derivatives clearing systems have been developed to deal with exchange-traded futures and options, there is strong pressure to force over-the-counter derivatives to go through similar clearing processes (Acharya et al. 2009; U.S. Congress' OTC Derivatives Market Act of 2009; U.S. Department of Treasury 2009; Duffie, Li, and Lubke 2010). In response, the CME, Intercontinental Exchange, and EUREX have recently created clearing facilities for Credit Default Swaps (CDS).

*Jorge A. Cruz Lopez is a Ph.D. student at Simon Fraser University. E-mail: jacruz@sfu.ca.
Jeffrey H. Harris is a finance professor in the College of Business and Economics at the University of Delaware and visiting professor and the James M. Collins Chair in Finance at Southern Methodist University. E-mail: harrisj@lerner.udel.edu.
Christophe Pérignon (the corresponding author) is an associate professor in the Finance Department of HEC Paris. E-mail: perignon@hec.fr.

Acknowledgements: We are grateful to an anonymous referee and to Dion Bongaerts, Mathijs van Dijk, Joost Driessen, Robert Engle, Thierry Foucault, Laurent Frésard, Eric Jondeau, Hai Lin, Christine Parlour, and Andrew Patton, and also participants at the 2011 Conference on Extreme Dependence in Financial Markets (Erasmus University), 2011 French Business School Finance Conference, and 23rd Australasian Finance and Banking Conference for their useful comments. This paper was prepared under a grant provided by The Institute for Financial Markets and that financial support is gratefully acknowledged.

Keywords: derivatives, tail risk, default risk, extreme dependence, copulas
JEL Classification: G13

In order to mitigate default risk, the clearing house requires clearing members to post margin (i.e., collateral). At the end of each day, the clearing house marks-to-market all outstanding trading positions and adjusts margins accordingly. A problematic situation arises, however, when the daily loss of a clearing firm exceeds its posted collateral. In this case, the firm may decide to default on its obligations, and the clearing house may have to draw on its default fund to compensate the winning counterparties.² Eventually, the clearing house may default as well after its default fund has been exhausted. This scenario, as unlikely as it may appear, is plausible, especially if several large clearing firms are in financial distress and ultimately default. It is also economically significant, because the failure of a clearing house would cause a major systemic shock that could spread default risk throughout the financial system.

Current practice on derivatives exchanges is to set the margin level of a derivative contract in such a way that it leads to a given target probability of a loss in excess of the margin (Figlewski 1984; Booth et al. 1997; Cotter 2001). Similarly, for a portfolio of derivatives, the margin requirement is derived from a distribution of simulated losses associated to the current portfolio positions (e.g., the Chicago Mercantile Exchange's SPAN system). We depart from this traditional view and account for tail dependence in the losses of clearing firms when setting collateral requirements. More specifically, we allow the margin requirements of a particular firm to depend not only on its own trading positions but also, potentially, on other clearing firms' positions. The basic intuition behind this concept is that the collateral requirement for a given clearing firm should increase when it is more likely to experience financial distress at the same time as other clearing firms.

Joint financial distress and defaults are more likely to occur when the trading positions of different clearing firms are similar or when they have similar risk exposures. Conceptually, the main cause of correlated trading across large clearing firms is that they share a common (and superior) information set. This informational advantage leads to similar directional trades. Furthermore, much of the proprietary trading activity on derivatives exchanges consists of arbitraging futures and over-the-counter markets or cash markets (e.g., cash-futures arbitrage of the S&P 500 index and eurodollar-interest rate swap arbitrage). As a result, if large clearing firms exploit similar arbitrage opportunities, they will have similar trading positions. Empirical evidence of correlated trading among large financial institutions is found in many settings, including futures markets. Using data for all Chicago Mercantile Exchange's (CME) clearing firms, for instance, Jones and Pérignon (2010) show that extreme losses by systemically-important clearing firms tend to cluster. This finding suggests that the derivative positions of the largest trading firms can be at times very similar.

Our approach for computing margins can be summarized as follows. We start from the trading positions of each clearing firm at the end of a given day. We then

2. Although exogenous events unrelated to futures losses might also result in default, we do not specifically address these situations.

consider a series of scenarios in which both the level and the volatility of all underlying assets are shocked by an arbitrary amount – in the spirit of *stress testing*. For each scenario, we mark-to-model the clearing firm’s portfolio and compute the associated hypothetical profit-and-loss (hereafter P&L). The standard collateral requirement of each clearing firm is equal to the $q\%$ quantile of the simulated P&L across all considered scenarios. Then for each pair of clearing firms, we compute the coefficient of lower tail dependence from the vectors of hypothetical P&L of both firms. This coefficient is defined as the probability of two clearing members having simultaneous extreme trading losses. We then set the collateral requirement of each clearing firm as a function of the highest coefficient of tail dependence between this firm and every other clearing firm. We show that accounting for interdependencies among clearing members reduces the likelihood of several clearing members being simultaneously in financial distress, as well as, the magnitude of the margin shortfall given joint financial distress, which greatly lowers systemic risk concerns.

Our methodology displays several attractive features. First, it is perfectly compatible with existing risk management techniques in place in derivatives exchanges, such as the SPAN system (Chicago Mercantile Exchange 2009). Second, our methodology can be applied at a daily or even higher frequency. This is important as an increasing number of derivatives exchanges mark-to-market positions twice a day (e.g., EUREX). Third, our approach differs from the “concentration risk” collateral method, which is most typically applied at the individual firm level. For instance, the Chicago Mercantile Exchange’s clearing house monitors concentrations by focusing on the proportion of open interest on a given contract that is controlled by a single clearing firm, and it assigns additional margin to reflect the incremental exposure due to concentration.

In terms of methodology, this paper is at the confluence of two streams of literature. First, we rely on modeling techniques for extreme dependence as in Longin and Solnik (2001), Ang and Chen (2002), Poon, Rockinger, and Tawn (2004), Patton (2008), and Christoffersen et al. (2010). While previous papers focus on stock or hedge fund returns, we show that tail dependence can also be very useful to jointly model clearing members’ P&L on a derivatives exchange. Second, our analysis builds on the recent literature on systemic risk. Adrian and Brunnermeier (2009) introduce the *CoVaR* measure that is the VaR of the financial system conditional on the distress of a given financial institution. Then they estimate the $\Delta\text{CoVaR}(\text{firm } i) = \text{CoVaR}(\text{system}|\text{firm } i) - \text{VaR}(\text{system})$ that captures the marginal contribution of a particular institution to the overall systemic risk. Related studies by Acharya et al. (2010) and Brownlees and Engle (2010) focus on the *Marginal Expected Shortfall* of a given bank, defined as the expected loss of a particular firm conditional on the overall banking sector being in distress. Similar to these papers, we measure, and attempt to internalize, the potentially negative externalities of having interconnected market participants. Although in the same spirit, we use a totally different methodology and focus on margin requirements and the risk that correlated positions pose to the clearinghouse.

The outline of the paper is the following. In Section I, we show how to estimate tail dependence among clearing firm losses. In Section II, we formally describe our methodology to set collateral as a function of tail dependence. We compare the performance of our method to the standard margining system using simulations in Section III. Section IV summarizes and concludes our paper.

I. TAIL DEPENDENCE

In derivatives markets, margins serve as performance bonds to guard against default. In our work, the *performance bond* $B_{i,t}$ represents the margin requirement imposed by the clearing house on clearing firm i at the end of day t , for $i = 1, \dots, N$. This performance bond depends on the outstanding trading positions of the clearing firm. The *variation margin* $V_{i,t}$ represents the aggregate mark-to-market profit or loss of clearing firm i on day t . The *relative variation margin* $R_{i,t}$ is defined as:

$$R_{i,t} = \frac{V_{i,t}}{B_{i,t-1}} \quad (1)$$

Clearing firm i is in financial distress at time t if $R_{i,t} < -1$, or equivalently if $B_{i,t-1} + V_{i,t} < 0$, since in this case the trading loss exceeds posted collateral. In such a situation, the clearing firm may decide to default, which would generate a shortfall in the system that needs to be covered by the clearing house.

By definition, tail dependence measures the probability of two random variables having simultaneous extreme events in the same direction. We define the coefficients of upper and lower tail dependence to quantify the comovement in revenues across clearing firms in extreme market conditions. In our context, the tail dependence structure captures the degree of diversification across clearing firms and the likelihood of having simultaneous financial distress across several clearing firms. The coefficient of *upper tail dependence* of the relative variation margins of clearing firms i and j at time t is defined as:

$$\tau_{i,j}^U = \lim_{\alpha \rightarrow 1} \Pr[R_i \geq F_i^{-1}(\alpha) \mid R_j \geq F_j^{-1}(\alpha)] = \lim_{\alpha \rightarrow 1} \Pr[R_j \geq F_j^{-1}(\alpha) \mid R_i \geq F_i^{-1}(\alpha)] \quad (2)$$

where $F_i(R_i)$ denotes the marginal cumulative distribution function of R_i for $i = 1, \dots, N$, and $a \in (0, 1)$ represents the marginal cumulative distribution level. Likewise, the coefficient of *lower tail dependence* of the relative variation margins of clearing firms i and j at time t is defined as:

$$\tau_{i,j}^L = \lim_{\alpha \rightarrow 0} \Pr[R_i \leq F_i^{-1}(\alpha) \mid R_j \leq F_j^{-1}(\alpha)] = \lim_{\alpha \rightarrow 0} \Pr[R_j \leq F_j^{-1}(\alpha) \mid R_i \leq F_i^{-1}(\alpha)] \quad (3)$$

Because we are primarily concerned with shortfall in the clearing system, we focus on the lower tail and simplify the notation as follows: $\tau_{i,j} = \tau_{i,j}^L$.

We model trading revenue dependence across clearing firms by using a bivariate copula (Patton 2009). A copula is a function that links together marginal probability distribution functions, say $F_i(R_i)$ and $F_j(R_j)$, to form a multivariate probability distribution function, in this case $F(R_i, R_j)$. According to Sklar's Theorem (Sklar 1959), if the marginal distributions are continuous, there exists a unique copula function C such that:

$$F(R_i, R_j) = C(F_i(R_i), F_j(R_j)) \quad (4)$$

Several features of copulas are useful in our context. First, marginal distributions do not need to be similar to each other. Second, the choice of the copula is not constrained by the choice of the marginal distributions. Third, copulas can be used with N marginal distributions. Fourth, the use of copula functions enables us to model the tails of the marginal distributions and tail dependence separately. This last point is very important in our case because in a multivariate setting, the likelihood of an extreme event can increase either because of fatter tails in the marginal distributions or because of fatter tails in the joint distribution function.

A natural candidate that allows us to incorporate tail dependence is the Student t -copula. Let t_v be the univariate Student t probability distribution function with v degrees of freedom. Then, for continuous marginal distributions, $F_i(R_i)$, the bivariate Student t -copula, $T_{\rho, v}$, is defined as:

$$T_{\rho, v}(F_i(R_i), F_j(R_j)) = t_{\rho, v}(R_i, R_j) \quad (5)$$

where $t_{\rho, v}$ is the bivariate distribution corresponding to t_v and $\rho \in [-1, 1]$ is the correlation coefficient between R_i and R_j .

A Student t -copula corresponds to the dependence structure implied by a multivariate Student t distribution. It is fully defined by the correlation of the implicit variables, ρ , and the degrees of freedom, v . The degrees of freedom define the probability mass assigned to the extreme co-movements of the relative variation margins (both positive and negative). In addition, this copula assigns a higher probability to joint extreme events, relative to the Gaussian copula, the lower the degrees of freedom, because a Student t copula with $v_i \rightarrow \infty$ corresponds to a Gaussian copula.

Student t -copulas allow us to readily obtain an estimate of the coefficient of lower tail dependence based on the correlation coefficient and the degrees of freedom (Cherubini, Luciano, and Vecchiato 2004):

$$\tau_{i, j} = 2t_{v+1} \left(-\sqrt{v+1} \sqrt{\frac{1-\rho}{1+\rho}} \right) \quad (6)$$

As can be seen from this equation, two parameters, the correlation coefficient and the degrees of freedom, fully describe the dependence structure of trading

revenues. Intuitively, larger correlations and lower degrees of freedom lead to higher tail dependence.

We implement a two-stage semiparametric approach to estimate the pairwise copulas across all clearing firms. The first stage consists of estimating the empirical marginal distribution of the trading revenues of each clearing firm. The second stage consists of estimating the t-copula parameters, ρ and ν , for every pair of clearing members through maximum likelihood (Genest, Ghoudi, and Rivest 1995).

II. COLLATERAL

In this section, we propose a new way of setting margin requirements for clearing firms. Our approach accounts for both tail risk and tail dependence structure across clearing firms. We consider a derivatives exchange with N clearing firms and D derivatives contracts (futures and options) written on U underlying assets. Let $w_{i,t}$ be the number of contracts in the derivatives portfolio of clearing firm i at the end of day t :

$$w_{i,t} = \begin{bmatrix} w_{i,1,t} \\ \vdots \\ w_{i,D,t} \end{bmatrix} \quad (7)$$

We consider two ways of computing the margin requirement of a clearing firm, which we present in turn below.

A. Standard Collateral Requirement

The standard collateral requirement is applied on a firm by firm basis, without regard to correlations across firms. As in the SPAN system utilized by the CME, we consider a series of S scenarios based on potential one-day ahead changes in the value (ΔX) and volatility ($\Delta\sigma_X$) of the underlying assets, as well as in the time to expiration of the derivatives products. For each of the S scenarios, we reevaluate the portfolio (i.e., we “mark-to-model” its positions) and compute the associated hypothetical P&L or variation margin on the portfolio:

$$\tilde{V}_{i,t+1} = \begin{bmatrix} \tilde{V}_{i,t+1}^1 \\ \vdots \\ \tilde{V}_{i,t+1}^S \end{bmatrix} \quad (8)$$

The *standard collateral requirement*, B , corresponds to the $q\%$ quantile of all simulated P&L across all considered scenarios:

$$Pr[\tilde{V}_{i,t+1}^s \leq -B_{i,t}] = q \quad (9)$$

with $s = 1, \dots, S$. Thus, B accounts for the potential financial distress of a particular clearing firm, but it ignores its interdependence with other clearing members. In

this standard case then, the total collateral collected by the clearing house at time t from all clearing firms is:

$$B_t = \sum_{i=1}^N B_{i,t} \tag{10}$$

B. Tail-Dependent Collateral Requirement

The tail-dependent collateral requirement is based not only on the magnitude of simulated losses (as in the standard collateral requirement) but also on the dependence structure across clearing firms' simulated losses. Our objective is to increase the collateral requirement for each individual firm by an amount proportional to its degree of dependence with other firms, with the increased collateral matching the incremental risk presented to the clearinghouse from potentially correlated losses among clearing members. Consider the portfolios of derivatives contracts of two clearing firms at the end of a given day:

$$W_{i,t} = \begin{bmatrix} w_{i,1,t} \\ \vdots \\ w_{i,D,t} \end{bmatrix} \quad W_{j,t} = \begin{bmatrix} w_{j,1,t} \\ \vdots \\ w_{j,D,t} \end{bmatrix} \tag{11}$$

For each clearing firm, we compute the variation margins generated by the S scenarios described in the previous section:

$$\tilde{V}_{i,t+1} = \begin{bmatrix} \tilde{V}_{i,t+1}^1 \\ \vdots \\ \tilde{V}_{i,t+1}^S \end{bmatrix} \quad \tilde{V}_{j,t+1} = \begin{bmatrix} \tilde{V}_{j,t+1}^1 \\ \vdots \\ \tilde{V}_{j,t+1}^S \end{bmatrix} \tag{12}$$

From Equation (12), we compute $B_{i,t}$ and $B_{j,t}$ as in equation (9). The tail dependence between the clearing firms' simulated relative variation margins is given by:

$$\tilde{\tau}_{i,j,t} = \lim_{\alpha \rightarrow 0} Pr \left[\tilde{R}_{i,t+1} \leq F_{i,t+1}^{-1}(\alpha) \mid \tilde{R}_{j,t+1} \leq F_{j,t+1}^{-1}(\alpha) \right] \tag{13}$$

where $\tilde{R}_{i,t+1} = \tilde{V}_{i,t+1} / B_{i,t}$. With N clearing firms, we end up with $N(N-1)/2$ tail dependence coefficients, which can be presented in a lower diagonal matrix:

$$\begin{bmatrix} \tilde{\tau}_{2,1} & & & & \\ \tilde{\tau}_{3,1} & \tilde{\tau}_{3,2} & & & \\ \vdots & & \ddots & & \\ \tilde{\tau}_{N,1} & \tilde{\tau}_{N,2} & \cdots & \tilde{\tau}_{N,N-1} & \end{bmatrix}$$

For each clearing firm we conservatively set its collateral requirement as a function of the highest coefficient of tail dependence with respect to all other clearing firms:³

$$\tilde{\tau}_{i,t} = \max \left\{ \tilde{\tau}_{i,j,t} \right\}_{j=1, j \neq i}^N \quad (14)$$

$$B_{i,t}^* = B_{i,t} \times e^{\max\{\gamma(\tilde{\tau}_{i,t} - \underline{\tau}); 0\}} \quad (15)$$

where γ is the tail-dependence aversion coefficient and $\underline{\tau}$ is a threshold tail dependence coefficient below which the collateral is not affected, that is: $B_{i,t}^* = B_{i,t}$ if $\tilde{\tau}_{i,t} \leq \underline{\tau}$. Thus, the total collateral collected by the clearing house becomes:

$$B_t^* = \sum_{i=1}^N B_{i,t}^* \quad (16)$$

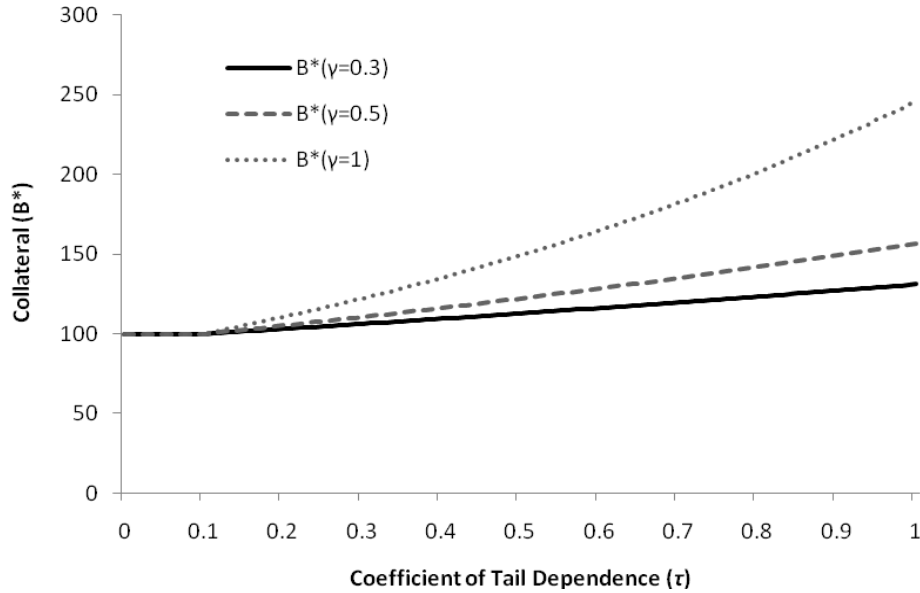
Notice that in the degenerative case where $\gamma = 0$ or if $\tilde{\tau}_{i,t} \leq \underline{\tau}$, we get the standard collateral requirement B . Thus, the standard collateral requirement (i.e., the SPAN system) is a special case of the tail-dependent collateral requirement. In other words, our approach can be seen as a generalized SPAN system. An implication of this result is that $B_t^* \geq B_t$. As an illustration, we plot in Figure 1 the level of tail-dependent collateral for different coefficients of tail-dependence aversion ($\gamma = 0.3, 0.5, 1$), $\underline{\tau} = 0.10$, $B = 100$, and for a tail parameter ranging between 0 and 1. Notice that no additional collateral is required for low coefficients of tail dependence. The required collateral increases with higher tail-dependence aversion, a choice variable for the clearing house, and with higher tail dependence, a parameter that can be estimated from simulated trading revenues.

III. CONTROLLED EXPERIMENT

To demonstrate the difference between the standard and tail-dependent collateral requirements, we consider a derivatives exchange with N clearing firms and two call options written on different underlying assets. Four clearing firms are assumed to be systemically important ($n = 4$) due to their size, so we focus on their margin requirements. Panel A of Table 1 displays the trading positions of these systemically important members in three different states: (1) low tail dependence, (2) moderate tail dependence, and (3) high tail dependence. The first state is obtained by selecting orthogonal trading positions across the systemically important firms. For the remaining states, the level of tail dependence is gradually increased by allowing the second firm to hold a position that progressively resembles that of the first. Notice, however, that the positions of the first, third, and fourth clearing firms remain constant across states. In addition, non-systemically important clearing firms

3. As a nested case, the standard collateral requirement case implies zero tail dependence.

Figure 1. Tail-Dependent Collateral.



Notes: This figure presents the level of tail-dependent collateral B^* as a function of the coefficient of tail dependence τ . The standard collateral requirement B is assumed to be equal to 100 and the threshold tail dependence coefficient $\underline{\tau}$ (below which collateral is not affected) to 0.10. The tail-dependence aversion coefficient γ of the clearing house varies between 0.3 and 1.

are assumed to clear the market in every state. Thus, each option contract is always in zero-net supply.

To simulate the variation margins for each clearing firm ($\tilde{V}_{i,t+1}$), we define S scenarios that combine potential one-day changes in the value of the underlying assets, ΔX_1 and ΔX_2 , with changes in volatility, $\Delta\sigma_{x1}$ and $\Delta\sigma_{x2}$. For each scenario, we mark-to-model the positions using the Black-Scholes model and generate a hypothetical change in the value of the portfolio held by each clearing firm. We then compute the coefficients of tail dependence between the simulated relative variation margins as described in equation (13). Panel B of Table 1 shows the estimated coefficients of tail dependence, and Table 2 shows the parameter values used for this controlled experiment.

Panel C of Table 1 compares three ways of computing collateral. The first two are the standard margin requirement (B) and the tail-dependent margin requirement (B^*) discussed earlier.⁴ The third collateral system aims at being budget-neutral, and it provides a better benchmark against which to compare the tail-dependent margining system because it collects the same aggregate collateral. This *budget-neutral* margin requirement is defined as:

4. See equation (9) for the definition of the standard margin requirement (B) and equation (15) for the definition of the tail-dependent margin requirement (B^*).

Table 1. Controlled Experiment.

	Low Tail Dependence				Moderate Tail Dependence				High Tail Dependence			
	1	2	3	4	1	2	3	4	1	2	3	4
Panel A: Trading Positions												
$d=1$	100	30	-50	-100	100	125	-50	-100	100	95	-50	-100
$d=2$	100	-170	150	-100	100	75	150	-100	100	105	150	-100
Panel B: Tail Dependence Coefficients												
$\tilde{\tau}_{2,j}$.000247908	.	.	.
$\tilde{\tau}_{3,j}$.000	.000	.	.	.000	.000	.	.	.000	.000	.	.
$\tilde{\tau}_{4,j}$.000	.000	.000	.	.000	.000	.000	.	.000	.000	.000	.
$\tilde{\tau}_j$.000	.000	.000	.000	.247	.247	.000	.000	.908	.908	.000	.000

Table 1, continued. Controlled Experiment.

	Low Tail Dependence				Moderate Tail Dependence				High Tail Dependence			
	1	2	3	4	1	2	3	4	1	2	3	4
Panel C: Margins												
B_i	3,849	6,228	4,310	5,319	3,849	3,918	4,310	5,319	3,849	3,851	4,310	5,319
B_i^*	3,849	6,228	4,310	5,319	4,022	4,094	4,310	5,319	4,905	4,908	4,310	5,319
B_i^0	3,849	6,228	4,310	5,319	3,936	4,005	4,397	5,406	4,377	4,380	4,839	5,847
p_i	.050	.050	.050	.050	.050	.050	.050	.050	.050	.050	.050	.050
p_i^*	.050	.050	.050	.050	.041	.038	.050	.050	.007	.007	.050	.050
p_i^0	.050	.050	.050	.050	.045	.044	.046	.046	.022	.022	.026	.033

Notes: Panel A presents the trading positions of the four systemically important clearing firms in two option contracts ($d = 1, 2$) when tail dependence is low, moderate, and high. Panel B displays the tail dependence coefficients among pairs of clearing firms ($\tilde{\tau}_{i,j}$) and the highest coefficient of tail dependence across all pairs ($\tilde{\tau}_i$). Panel C show the standard margins (B_i), the tail dependent margins (B_i^*), and the budget-neutral margins (B_i^0). When computing tail dependent margins, we use a tail-dependence aversion coefficient γ of 0.3 and a threshold tail dependence coefficient $\underline{\tau}$ of 0.1. Finally, the p_i variables denote the probability of a clearing firm being in financial distress: $p_i = \Pr[V_{i,t+1} \leq -B_{i,t}]$, $p_i^* = \Pr[V_{i,t+1} \leq -B_{i,t}^*]$, and $p_i^0 = \Pr[V_{i,t+1} \leq -B_{i,t}^0]$.

Table 2. Controlled Experiment Parameters.

Parameter	Value
A. Market and Clearing Members	
Number of derivatives securities (D)	2
Number of underlying assets (U)	2
Number of systemically important clearing members (n)	4
B. Underlying Assets	
Value of underlying asset 1 at $t = 0$	\$100
Value of underlying asset 2 at $t = 0$	\$100
C. Derivatives Securities	
Strike price of option contract 1	\$100
Strike price of option contract 2	\$100
Time to maturity of option contract 1	1 year
Time to maturity of option contract 2	1 year
D. Margining Systems	
Variation range in the value of the underlying assets	$\pm 50\%$
Variation range in the volatility of the underlying assets' returns	$\pm 50\%$
Number of scenarios for the value of the underlying asset and its volatility (S)	10,000
Quantile for the standard collateral system (q)	5%
Tail-dependence aversion coefficient (γ)	0.3
Threshold tail-dependence coefficient ($\underline{\tau}$)	0.1

$$B_{i,t}^0 = B_{i,t} + \frac{B_t^* - B_t}{n} \text{ for } i = 1, \dots, n \quad (17)$$

where the budget-neutral condition is:

$$\sum_{i=1}^n B_{i,t}^0 = \sum_{i=1}^n B_{i,t}^* \quad (18)$$

and from equation (15) it follows that $B_{i,t}^0 = B_{i,t} = B_{i,t}^*$ when $\tilde{\tau}_{i,t} \leq \underline{\tau}$ or $\gamma = 0$.

The results presented in Panel C of Table 1 show that the three margining systems are equivalent in the low-dependence state and that they diverge as the level of tail dependence increases to 0.247 in the moderate-dependence state, and to 0.908 in the high-dependence state. The equivalence across margining systems in the low-dependence state arises because the tail dependence coefficients are virtually zero; thus, $\tilde{\tau}_{i,t} \leq \underline{\tau}$ and $B_{i,t}^* = B_{i,t}$ for all clearing firms. In other words, when default risk is well-diversified among clearing firms, the tail-dependent margining

system converges to the standard system. On the other hand, in the moderate and high-dependence states, the tail-dependent margin requirement for clearing firms 1 and 2 increases due to their progressively homogeneous trading positions. This homogeneity is captured by the higher coefficient of tail dependence that is incorporated into B^* .


Notice, however, that the standard and tail-dependent margin requirements of firms 3 and 4 remain unchanged across states as their positions stay orthogonal relative to those of the other members. This is not true for the budget-neutral case. The budget-neutral margin requirement increases for all members in the moderate and high-dependence states. This situation arises because the additional collateral that would be collected under the tail-dependent margining system is now collected across all systemically important clearing firms. As a consequence, the budget neutral collateral requirements of firms 3 and 4 increase in the moderate and high dependence state due to the increased tail dependence between firms 1 and 2.

In order to assess the appropriateness of each margining system, we now turn our attention to their relative performance. We simulate changes in the value of the call options by randomly selecting one of the S scenarios. For each margining system, we compute the probability of financial distress across clearing firms, the probability of joint financial distress, and the magnitude of the average margin shortfall given joint financial distress. The bottom part of Panel C in Table 1 shows the probability of financial distress (i.e., the probability that $B_{i,t-1} + V_{i,t} < 0$) across clearing firms. Since the quantile for the standard margining system, q , was set to 5% in the simulation (see Table 2), the standard system has a distress probability of 5% in all scenarios by construction.

Similarly, in the low-dependence state, when $B_{i,t} = B_{i,t}^* = B_{i,t}^0$ for all clearing firms, the probability of financial distress is 5% across margining systems. In the moderate and high-dependence states, however, the distress probability is lower for firms 1 and 2 under the tail-dependent system and lower for all firms under the budget neutral system because more collateral is required relative to the standard case.

At first glance, this result would suggest that the budget neutral system performs better than the alternatives because it reduces the unconditional probability of financial distress across clearing firms. However, Figure 2 shows that the tail-dependent margining system actually provides a better allocation of margin requirements. More specifically, the figure shows that the probability of joint financial distress (i.e., the probability of one or more clearing firms jointly experiencing a loss in excess of their posted margin) is lower under the tail-dependent margining system, particularly when tail dependence is high.

Notice that the probability of joint financial distress increases monotonically with tail dependence under the standard collateral system. Differently, for the tail-dependent system, this probability first increases in the moderate-dependence state and then decreases in the high-dependence state. This result arises due to the value of the tail-dependence aversion coefficient, $\gamma = 0.3$, and the value of the threshold tail dependence coefficient, $\underline{\tau} = 0.1$ (see Table 2), which translates into a slight increase in the required margin for firms 1 and 2 (an additional \$173 and \$176,

Figure 2. Probability of Joint Financial Distress


Notes: This figure presents, for each margining system, the likelihood of several clearing firms jointly being in financial distress, i.e., $B_{i,t-1} + V_{i,t} < 0$, with different levels of tail dependence between clearing firms (low, moderate, and high). The three systems are the standard (B), tail-dependent margin requirement (B^*), and budget-neutral (B°) margin requirement systems. The results are based on 1,000,000 simulations of the actual changes in the underlying asset prices.

respectively) in the moderate tail-dependence state, and a significantly larger increase (an additional \$1,056 and \$1,057, respectively) in the high tail-dependence state. Similar results can be observed in the budget neutral system for the same reasons. A monotonic decrease of the probability of joint financial distress could be obtained for the tail-dependent collateral system if a higher value of γ or a $\underline{\tau}$ of 0 is selected.

Finally, Figure 3 shows that the average shortfall ($B_{i,t-1} + V_{i,t}$), given financial distress, is lower under the tail-dependent margining system in the moderate and high-dependence states. Therefore, we can conclude that the tail-dependent margining system is superior to the other systems because it provides a better allocation of margin requirements. This allocation depends on the composition and homogeneity of the trading positions of the clearing members and it provides better protection against joint negative outcomes.

IV. CONCLUDING REMARKS

In this paper, we present a novel approach to compute margins for a portfolio of derivatives securities. The innovative feature of this method is to account not only for the riskiness of the trading positions of an individual market participant but also for the interdependence between this participant's trading positions and *other* participants' trading positions. Our method is a simulation-based technique that accounts for extreme tail dependence among potential trading losses. Accounting

Figure 3. Average Shortfall Given Joint Financial Distress.

Notes: This figure presents, for each margining system, the average shortfall ($B_{i,t-1} + V_{i,t}$) given joint financial distress with different levels of tail dependence between clearing firms (low, moderate, and high). The three systems are the standard (B), tail-dependent margin requirement (B^*), and budget-neutral (B°) margin requirement systems. The results are based on 1,000,000 simulations of the actual changes in the underlying asset prices.

for interconnections among clearing firms in a derivatives exchange is shown to lower the probability of several clearing members being simultaneously in financial distress (i.e., when losses exceed posted collateral), as well as the magnitude of the margin shortfall given joint financial distress, which decreases systemic risk concerns.

While our simulation analysis focuses on margins for option positions, our method can be applied to any listed derivatives contract such as futures, swaps, or exchange-traded credit derivatives. Furthermore, it is important to realize that our approach should by no means be limited to derivatives exchanges and can also be used to set collateral in any financial network. For instance, our method could be used to set collateral requirements for OTC positions as well.

References

- Acharya, V.V., Engle, R., Figlewski, S., Lynch, A., and Subrahmanyam, M., 2009, Centralized Clearing for Credit Derivatives. In *Restoring Financial Stability: How to Repair a Failed System*, edited by V. Acharya and M. Richardson (John Wiley & Sons).
- Acharya, V.V., Pedersen, L.H., Philippon, T., and Richardson, M., 2010, Measuring Systemic Risk. *Working paper, New York University*.
- Acworth, W., 2009, The Lessons of Lehman: Reassessing Customer Protections. *Futures Industry*, January/February, 33-36.
- Adrian, T. and Brunnermeier, M.K., 2009, CoVaR. *Working paper, Princeton University and FRB of New York*.
- Ang, A. and Chen, J., 2002, Asymmetric Correlations of Equity Portfolios. *Journal of Financial Economics*, 63, 443-494.

- Booth, G.G., Broussard, J.P., Martikainen, T., and Puttonen, V., 1997, Prudent Margin Levels in the Finnish Stock Index Futures Market. *Management Science*, 43, 1177-1188.
- Brownlees, C.T. and Engle, R., 2010, Volatility, Correlation and Tails for Systemic Risk Measurement. *Working paper, New York University*.
- Cherubini, U., Luciano, E., and Vecchiato, W., 2004, *Copula Methods in Finance* (John Wiley & Sons Ltd).
- Chicago Mercantile Exchange, 2009, *CME Clearing Financial Safeguards*.
- Christoffersen, P., Errunza, V., Jacobs, K., and Jin, X., 2010, Is the Potential for International Diversification Disappearing? *Working paper, McGill University*.
- Cotter, J., 2001, Margin Exceedences for European Stock Index Futures using Extreme Value Theory. *Journal of Banking and Finance*, 25, 475-502.
- Duffie, D., Li, A., and Lubke, T., 2010, Perspectives on OTC Derivatives Market Infrastructure. *Federal Reserve Bank of New York, Staff Reports Policy No. 424*.
- Duffie, D. and Zhu, H., 2010, Does a Central Clearing Counterparty Reduce Counterparty Risk? *Working paper, Stanford University*.
- Figlewski, S., 1984, Margins and Market Integrity: Margin Setting for Stock Index Futures and Options. *Journal of Futures Markets*, 4, 385-414.
- Genest, C., Ghoudi, K., and Rivest, L.-P., 1995, A Semiparametric Estimation Procedure of Dependence Parameters in Multivariate Families of Distributions. *Biometrika*, 82, 543-552.
- Jones, R.A. and Pérignon, C., 2010, Derivatives Clearing, Default Risk, and Insurance. *Working paper, HEC Paris*.
- Longin, F. and Solnik, B., 2001, Extreme Correlation of International Equity Markets. *Journal of Finance*, 56, 649-676.
- Patton, A.J., 2008, Are “Market Neutral” Hedge Funds Really Market Neutral? *Review of Financial Studies*, 22, 2495-2530.
- Patton, A.J., 2009, Copula-Based Models for Financial Time Series. In *Handbook of Financial Time Series*, edited by T.G. Andersen, R.A. Davis, J.-P. Kreiss, and T. Mikosch (Springer Verlag).
- Pirrong, C., 2009, The Economics of Clearing in Derivatives Markets: Netting, Asymmetric Information, and the Sharing of Default Risks through a Central Counterparty. *Working paper, University of Houston*.
- Poon, S.H., Rockinger, M., and Tawn, J., 2004, Extreme Values Dependency in International Stock Markets: Diagnostics, Models, and Financial Implications. *Review of Financial Studies*, 17, 581-610.
- Sklar, A., 1959, Fonctions de Répartition à N Dimensions et leurs Marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8, 229-231.
- U.S. Department of Treasury, 2009, *Financial Regulatory Reform: A New Foundation* (Washington, DC).

WOULD POSITION LIMITS HAVE MADE ANY DIFFERENCE TO THE “FLASH CRASH” ON MAY 6, 2010?

Bernard Lee, Shih-fen Cheng, and Annie Koh*

On May 6, 2010, the U.S. equity markets experienced a brief but highly unusual drop in prices across a number of stocks and indices. The Dow Jones Industrial Average (see Figure 1) fell by approximately 9% in a matter of minutes, and several stocks were traded down sharply before recovering a short time later. The authors contend that the events of May 6, 2010 exhibit patterns consistent with the type of “flash crash” observed in their earlier study (2010). This paper describes the results of nine different simulations created by using a large-scale computer model to reconstruct the critical elements of the market events of May 6, 2010. The resulting price distribution provides a reasonable resemblance to the descriptive statistics of the second-by-second prices of S&P500 e-Mini futures from 2:30 to 3:00 p.m. on May 6, 2010. This type of simulation avoids “over-fitting” historical data, and can therefore provide regulators with deeper insights on the possible drivers of the “flash crash,” as well as what type of policy responses may work or may not work under comparable market circumstances in the future. Our results also lead to a natural question for policy makers: If certain prescriptive measures such as position limits have a low probability of meeting their policy objectives on a day like May 6, will there be any other more effective counter measures without unintended consequences?

*Bernard Lee (the corresponding author) is a visiting associate professor (Practice) and Deputy Director, Sim Kee Boon Institute for Financial Economics, at Singapore Management University. E-mail: bernardlee@smu.edu.sg.

Shih-fen Cheng is an assistant professor in the School of Information Systems at Singapore Management University. E-mail: sfcheng@smu.edu.sg.

Annie Koh is an associate professor of finance and Dean, Office of Executive and Professional Education; Academic Director, International Trading Institute, at Singapore Management University. E-mail: anniekoh@smu.edu.sg.

Acknowledgements: The authors thank Ms. Zhan Yingbo of the Quantitative Finance Program at the National University of Singapore for her able research assistance. In addition, the authors gratefully acknowledge the generous financial support by the Institute for Financial Markets, as well as the opportunity to present our findings at the 2011 Asia Pacific Futures Research Symposium. We also want to thank the Monetary Authority of Singapore and the Singapore Exchange for inviting us to discuss our work at a Regulators’ Meeting in conjunction with the 2010 Futures Industry Association Asia Conference. Finally, the authors thank both the Chicago Mercantile Exchange and the Singapore Exchange for providing certain primary sources of data. Any additional errors and omissions remain our own.

Keywords: fill in

JEL Classification: fill in

There are many publicly-available accounts of the market events of May 6, 2010. We will not attempt to repeat those accounts here. We will aim to provide a relatively straightforward summary, for the purpose of setting the proper context of our simulation analysis. Given that we are simply summarizing basic facts for the convenience of our audience, we would like to acknowledge the relevant sources all at once, including the Joint CFTC-SEC Preliminary Report and its corresponding Final Report (CFTC 2010a,b), as well as a research report published by the CME Group shortly after the May 6, 2010 incident (CME Group 2010). In addition, we have benefited from primary sources of data provided by the CME Group as well as the SGX.¹

The trading day of May 6, 2010 started with unsettling political and economic news due to the European debt crisis. Just one day before, the Greek government's debt crisis boiled over into violence on the street of Athens. These factors had weighed on global markets before U.S. trading hours, and the U.S. equity market was down in early trading. At around 2:30 p.m. (all times are shown in Eastern Standard Time), the overall decline suddenly accelerated, after a rush of sell orders. Within a few minutes, both the S&P 500 Index and its June 2010 E-mini futures dropped by more than 5% (shown in Figure 2).


Staff of the Commodity Futures Trading Commission (CFTC) conducted a post-mortem analysis of the top 10 largest longs and shorts. Those analyses suggest that, in most cases, traders with the largest longs and shorts in fact traded on both sides of the market. In other words, there was no obvious one-sided "squeezing" of the market. The CME order books on futures also showed that there were many more sell orders than buy orders from 2:30 to 2:45 p.m. However, the volume of E-mini futures surged to eight times that of SPDRs (after adjustments) between 2:45 and 2:50 p.m. To most traders, this was a clear indication that the futures market was driving the cash market, not the other way around.

The bid-ask of the June 2010 E-mini S&P 500 futures widened considerably at about 2:45 p.m., triggering CME's Globex stop logic functionality. The stop logic functionality aims to prevent the triggering of stop-loss orders that would have resulted in transactions at price levels below the contract's "no-bust range," leading to an avalanche of price declines due to order-book imbalances. This functionality put the market in a "reserve" state when orders could be entered, modified, or cancelled but not concluded. It was, in fact, triggered in the E-mini market at 2:45:28 p.m. for five seconds, precisely when the E-mini contract hit its low of the day. Since futures were not traded during these five seconds, the linkages between the cash and the futures markets would have broken down despite that, in theory, U.S. stock futures that are traded on the CME are supposed to be coordinated with cash equity trading on the New York Stock Exchange (NYSE).

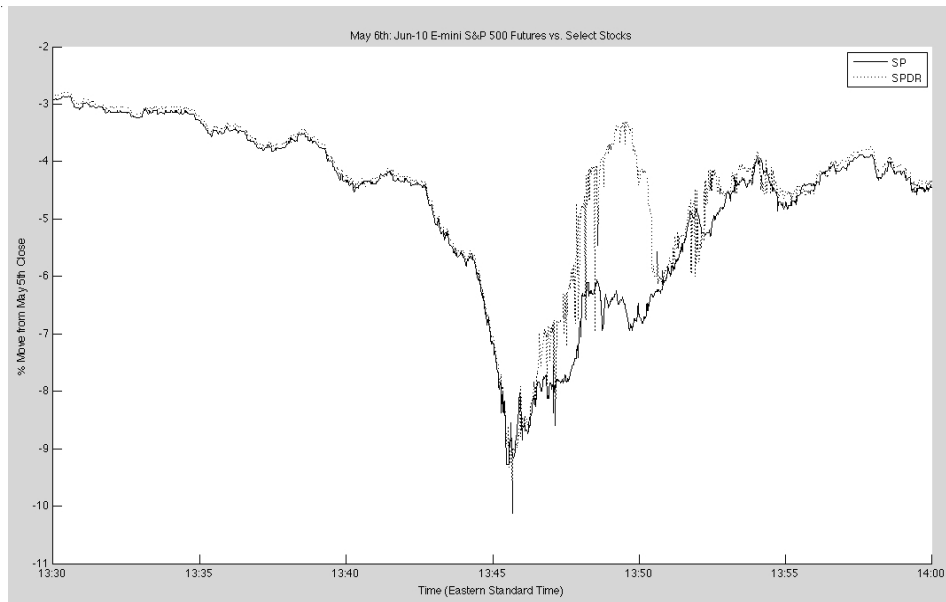
The majority of the single-name stocks had declines consistent with the 5% decline in June 2010 E-mini S&P 500, which traded at its low of 1056 by 2:34:28 p.m. However, three stocks — namely, Proctor and Gamble (PG), 3M (MMM),

1. The authors gratefully acknowledge the help from John Labuszewski of the CME Group as well as that of Sutat Chew from the Singapore Exchange.

Figure 1. DJIA, Dollar-Yen, Nikki and Gold Prices on May 6, 2010.



(Data courtesy of SGX.)

Figure 2. June 2010 E-mini futures on S&P 500 vs. SPDRs.

(Data courtesy of CME Group.)

and Accenture (ACN) — continued to decline even as the E-mini S&P 500 contract hit its low and then began to reverse upward. These three stocks hit their Liquidity Replenishment Points (LRPs) at 2:45:52 p.m., 2:50:36 p.m., and 2:46:10 p.m., respectively, while their lowest trading prices of \$39.37, \$67.98, and \$0.01 were reported at 2:47:15 p.m., 2:45:47 p.m., and 2:47:54 p.m., respectively.

Eventually, Nasdaq announced that it would bust all trades that were more than 60% off the market. Of the U.S.-listed securities with declines of 60% or more away from the 2:40 p.m. transaction prices (resulting in busted trades), approximately 70% were ETFs. This suggested that ETFs as an asset class were affected more than any other categories of securities. One hypothesis is that ETF might have been widely used by investors as inexpensive short hedges and in placing stop-loss market orders.

Several hypotheses were raised by the “CFTC-SEC Preliminary Report to the Joint Advisory Committee on Emerging Regulatory Issues” as to what might have caused the trading experience of May 6, 2010:

1. Disparate trading venues in the United States; this is also known as “market fragmentation.” It refers to the fact that multiple exchanges, alternative trading systems, and private matching networks (dark pools) run by broker-dealers all trade the same stocks in the United States simultaneously. While the overall liquidity may appear substantial, whenever there is a liquidity problem faced by one of the many trading venues containing a fraction of the total liquidity, the manner in which that venue reacts to the problem may initiate an overall chain reaction. Such a chain

reaction may not have happened at all if the total liquidity for each stock can be consolidated into a single trading venue.

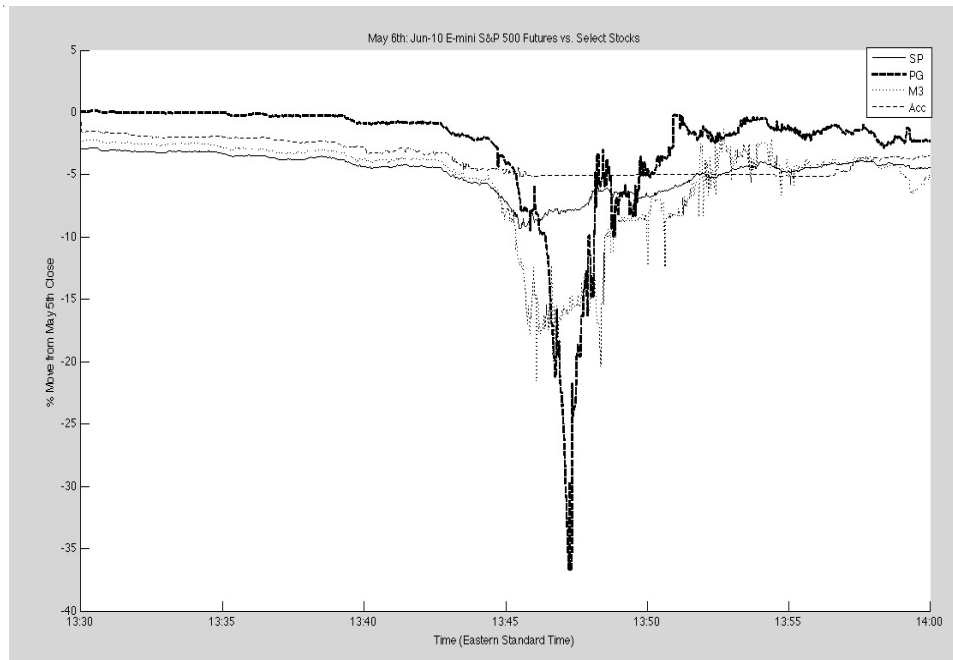
2. *“Liquidity Replenishment Points”* (LRPs) at the NYSE and similar practices. Whenever an LRP is triggered, the NYSE will go into a “go slow” mode and pause momentarily to allow liquidity to enter the market. This may have exacerbated the problem, in that automated trading orders are most likely rerouted to other possible trading venues, resulting in a net loss of trading liquidity at the primary market. This may also have the effect of triggering similar cautionary procedures in parallel trading venues, driving liquidity further from the market.

3. *“Self-Help remedy.”* Two exchanges declared “self help” against NYSE Arca in the minutes prior to 2:40 p.m., after NYSE Arca repeatedly failed to provide a response to incoming orders within one second. Such declarations free the declaring exchanges from their obligations to route unmatched orders to the affected exchange, resulting in additional loss of trading liquidity. For instance, a high bid and a low ask on the same stock appearing on two different exchanges, which could have been matched if there was rerouting, would fail to be matched under such circumstances.

4. *Stop loss market orders.* Some market participants left sell orders much lower than current prices as market orders to sell, primarily as a stop-loss precaution. Those orders were not expected to be executed. In a fast-falling market, these stop-loss market orders might have triggered a chain reaction of automated selling orders, and the sellers would have limited time to reconsider those orders. Typically, such orders would be left by institutional investors, and the quantity involved could be quite substantial as compared to the existing liquidity for a particular stock.

5. *Short sales and stub quotes.* Short sales against stub quotes accounted for more than 70% of the busted trades between 2:45 and 2:50 p.m. and approached a staggering 90% between 2:50 and 2:55 p.m. The fact that stub quotes were never intended to be executed, and that there would be limited (if any) upside to short selling against near-zero bids, suggests that at least some of these short sales were placed in a somewhat automated manner, since it would be unlikely for any experienced human trader to execute such orders.

In Lee, Cheng, and Koh (2010), the authors constructed a simulated market with multiple types of computer agents, including a market maker, systematic traders (deploying several varieties of trend-following strategies, which are among the most common techniques deployed by hedge funds), and “retail-like” investors who place randomized bids and asks in the market in a mean-reverting manner. Unlike traditional market simulations, the evolution of asset prices is the direct result of how these agents are trading against each other as in real markets, and there are no a priori

Figure 3. June 2010 E-mini Futures on S&P 500 vs. PG, MMM, and ACN.

(Data courtesy of CME Group.)

assumptions on asset price distributions. While market simulation is hardly new, the academic contributions of our work are the following:

- (i) We provide a convincing description of market dynamics based on the structure of the market and the type of participants.
- (ii) The resulting price distribution provides a reasonable resemblance of the descriptive statistics of certain commodity markets.
- (iii) Yet the simulation does not contain so many degrees of freedom that it essentially “over-fits” historical data, resulting in limited predictive power and insights.

The key findings from our earlier study include the following:

1. In theory, trend-following is a trading strategy that can be replicated by lookback straddles, which is a traditional “long gamma” strategy. The theoretical strategy is supposed to have unlimited upside but limited downside, much like any option. However, most option pricing theories work under the unrealistic assumptions of infinite liquidity and zero transaction costs. What we have observed is that, as we deliberately withdraw liquidity from the market, the profit-and-loss profiles of the trading strategies will deviate further and further away from the theoretical bounds derived based on option theories.

2. As the percentage of systematic traders in the market exceeds a certain threshold (between 60% and 80%) relative to the total number of market participants, the bids and offers in the market will concentrate on only one side of the market, especially during extreme market movements. Market prices will begin to behave erratically, leading to the eventual breakdown of the market.

3. Finally, any attempt to restore market liquidity by changing the “rules of the game” in the middle of trading is unlikely to produce the desired outcome. The process for market agents to adjust to any new set of rules, as well as subsequently reversing to the original state of the market, appears to cause more problems than it solves by creating significant liquidity disruptions to the market.


The goal of this paper is to determine if the findings from the earlier paper can be used to understand and assess potential regulatory responses, such as those listed in the “Joint CFTC-SEC Preliminary and Final Report.” In particular, the authors contend that the events of May 6, 2010, show a pattern consistent with the type of “flash crash” observed in our earlier study. While some commentators assigned blame to high-frequency trading, our analysis was unable to identify a direct link to high-frequency trading per se, but rather the domination of market activities by trading strategies that are responding to the same set of market variables in similar ways, as well as various pre-existing schemes that modify the “rules of the game” in the middle of trading, resulting in a significant withdrawal of liquidity during extreme market movements. In addition, certain micro-structural safety mechanisms in the market, such as the uneven triggering of circuit breakers by the cash equity, futures, and ETF markets at different times, may have exacerbated the problem.

Furthermore, the triggering of the Liquidity Replenishment Points at the New York Stock Exchange (NYSE), commonly known as “go slow” mode, might have further driven liquidity out of the market when it was needed the most. Only when certain stocks reached “stupid cheap” levels, other investors seized the opportunity to buy and market prices began restoring to levels consistent with fundamental valuations. Moreover, the subsequent cancelling of trades by the NYSE has created a significant worry for market participants (market makers in particular) who can potentially step in to provide much-needed liquidity in similar episodes in the future.

To achieve our objectives, we have constructed nine different simulations in this study, in an attempt to recreate various market conditions for the cascading effects leading to the type of flash crash seen on May 6. Those results allow us to study the potential effects of:

- imposing position limits by traders.
- changing from continuous time auctions to discrete time auctions.
- imposing price limits during a major market dislocation, with different trigger levels.

Figure 4. A Sample Market Structure that Agents Need to Understand.



I. DESIGNING THE SIMULATION PLATFORM

It has been widely speculated that the Flash Crash on May 6, 2010 was caused primarily by two factors: (a) trading venues with different and often inconsistent rules of operations and (b) complex dependency among multiple assets (e.g., between index tracking ETFs and its component stocks). The first factor contributes to the congestion of orders when trading venues slowing down unevenly, while the second factor contributes to the contagion of instability from one asset to other related assets. In order to reconstruct the market conditions leading to the Flash Crash and to evaluate policies that could help preventing similar incidents, we have developed a realistic microscopic financial simulation even though, to the best of our knowledge, no financial simulator can reproduce faithful replications of both features.

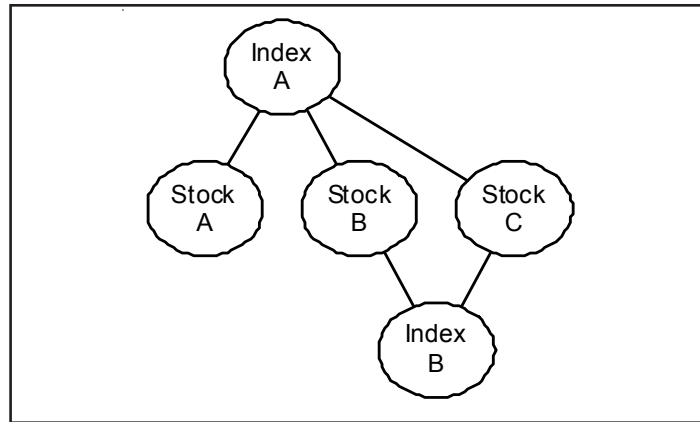
The simulation platform utilized in this paper is derived from the model first introduced in Cheng (2007), and used subsequently for analyzing extreme market conditions in Lee et al. (2010). In the following subsections, we will briefly describe the enhancements necessary for the simulation platform to model the two features mentioned above.

A. Multiple Trading Venues

With any sufficiently generic market engine, introducing multiple trading venues is relatively straightforward: The engine can simply create additional markets according to rules as specified by the user. However, the key challenge of having multiple trading venues is not about creating additional markets but avoiding operational bottlenecks. More specifically, we need to address how we can design a conceptual structure that is understandable by software agents and come up with a reasonable price discovery process under multiple trading venues.

For the software agents that we plan to introduce to the system, they need to recognize the relationship among multiple markets. For example, for the case where a particular asset A is traded simultaneously in two markets, an agent needs to understand that buying and selling A in both markets will directly affect the position of A. In other words, agents in the simulation will need to load a conceptual mapping like the one illustrated in Figure 4. In our simulation design, we allow structural information to be defined compactly and all agents are required to load this same structural information at the starting-up phase. Once such mapping is loaded, an

Figure 5. Introducing Dependencies to Assets.



agent will then be able to keep an aggregated view on position balances through the linkages between markets.

Another important issue that needs to be addressed when introducing multiple trading venues is how prices of the same asset are synchronized across different markets. Take asset A in Figure 4 as an example: An agent intending to establish a long position in asset A needs to decide which market to trade in, since markets 1 and 2 are running independently and may have different prices. Agents certainly may have their own logic in deciding which market to go for; however, to simplify agent design and to emulate real-world trading rules, we assume that all bids and offers submitted by agents will go through a mechanism similar to the National Best Bid and Offer (NBBO) rule implemented in the U.S. stock market. In other words, when picking which market to trade in, an agent will simply pick the market with lowest ask prices (from all markets) when buying and the market with highest bid price when selling. Our assumption is that the updates on best ask/bid prices from all markets will be instantaneous without delay.

The framework presented above will allow us to design arbitrary market structures that suit our needs.

B. Complex Asset Dependency

Another important feature that we want to introduce is to allow assets to be related with each other. For example, the trading price of an index future should be dependent on the prices of all stock components this index future tracks. By allowing such dependencies, we are effectively linking up independent assets. An example of such dependency is illustrated in Figure Y.

Prices of linked assets cannot be directly synchronized, since prices of all assets still need to be determined by the market. Therefore, we need to go through a market mechanism to synchronize these asset prices. In order to achieve such synchronization, we introduce a special agent class called the “Arbitrageur.” Arbitrageurs understand the relationship between assets, and they will trade whenever market prices are significantly out-of-sync.

Taking Index A in Figure 5 as an example: By assuming that Stocks A, B, and C are equally weighted in Index A, we can design the Arbitrageur using the following rules to eliminate any out-of-sync prices:

- If $Bid_{IndexA} \geq (1 + a)\{Ask_{StockA} + Ask_{StockB} + Ask_{StockC}\}$, then the arbitrageur should buy the basket of three stocks and sell the index.
- If $(1 + a)Ask_{IndexA} \leq \{Bid_{StockA} + Bid_{StockB} + Bid_{StockC}\}$ then the arbitrageur should buy the index and sell the basket of three stocks.

The parameter a is introduced to account for market frictions like delays or transaction costs. Arbitrageur will constantly review its holding, and whenever any of the following conditions is met, the Arbitrageur will liquidate its positions:

(1) If the price discrepancy disappears, that is, $Mid_{Index} \approx Mid_{StockA} + Mid_{StockB} + Mid_{StockC}$. The tolerance for being “sufficiently close” for liquidation can be adjusted empirically based on the bid-ask spreads shown in the tradable assets.

(2) If a perfect arbitrage is unsuccessful because of market slippage, we will implement a stop-loss rule to “reverse out” from any yet-to-be completed arbitrage trade based on a time trigger. This will happen when say only three out of the four legs of the arbitrage trade can be executed at the intended prices. This is an important feature to be included in any type of “flood to the gate” scenario, when one or more legs of an arbitrage trade is moving away from its intended price and the Arbitrageur has no choice but to unwind the trade.

(3) If, instead of convergence, an arbitrage trade diverges and creates losses instead of profits, the Arbitrageur will automatically “reverse out” from the arbitrage trade to prevent any run-away negative P&L. This is consistent with real-world practices and is another important feature to be included in any type of “flood to the gate” scenario. The trigger for stop loss is set to 5% initially and will be adjusted empirically based on the actual price behavior shown in the tradable assets.

The above rules for the Arbitrageur can be easily generalized to include an arbitrary number of assets and uneven weights.

II. SIMULATION DESIGN

A. Current Study

As mentioned earlier, we have conducted nine different simulations in this study, in an attempt to recreate various market conditions for the cascading effects leading to the type of flash crash seen on May 6. Those results allow us to study the potential effects of imposing position limits by traders, changing from continuous

time auctions to discrete time auctions, and imposing price limits during a major market dislocation, with different trigger levels.

Specifically, there are the “deltas” from one simulation to the next in the current study:

Simulation 1 → *Simulation 2*: Compressing the action-reaction time from the “go slow” mode in exchange 1 to the “go slow” mode in exchange 2, in order to pinpoint the potential triggering conditions leading to cascading effects. The purpose is to illustrate how market micro-structural issues can make a significant difference to market stability.

Simulation 2 → *Simulation 3*: Imposing position limits by trader, instead of typical position limits by symbols (i.e., per stock trading on each individual exchange).

Simulation 3 → *Simulation 4*: Changing the clearing mechanism from continuous time auction to discrete time auction, which would have negate any trade execution advantages of high-frequency, algorithm-based trading.

Simulation 3 → *Simulation 5*: Simulation 5 is a variant of Simulation 3, in which quotes are not updated during the slowdown.

Simulation 3 → *Simulation 6*: Simulation 6 is a variant of Simulation 3, in which price limits are imposed when prices have dropped by more than 40%, respectively, when compared to the base prices that are sampled from the last done prices every 60 seconds.

Simulation 6 → *Simulation 7*: The trigger level above is set to 30% instead.

Simulation 7 → *Simulation 8*: The trigger level above is set to 20% instead.

Simulation 8 → *Simulation 9*: The trigger level above is set to 10% instead.

B. Technical Descriptions of Market Agents

For each stock, there are two markets in which it can be traded, with one market being roughly twice as large as another market (in terms of initially-available liquidity). Each stock is serviced by a Market Maker (MM) that is willing to provide liquidity by earning a small fee; the Index market, on the other hand, is not serviced by any MM. Besides the Market Maker, there are also Zero Intelligence (ZI) (or “random”) agents, Trend Following (TF) agents, and Arbitrageur (AA) agents, with the latter having been described in detail in Section IB. Both ZI and TF agents are allowed to trade every stock available; however, only ZI agents are allowed to

trade the Index. When trading in the Index market, ZI agents are designed to understand the linkage between index and its stock components. Whenever there are sufficiently large gaps between prices of index and component stocks, the AA agent will be performing arbitrating trades as described in Section IB and pulling the Index back to its fair value in the process. Non-convergence in the Index market is allowed and is one critical element of the market that we intend to model.

We have designated separate agents to emulate automatic stop losses and to generate the initial selling pressure in the Index market similar to the rush of sell orders at around 2:30 p.m. on May 6. A group of four agents (known as Bear Market agents) will automatically start piling in sell orders quickly once the major market slows down, to simulate the initial triggering of sell orders by traders who are likely to interpret the “go slow” mode as highly-negative market sentiments. To trigger automatic stop losses as and when the market suffers significant losses, a group of three agents will constantly monitor the stock prices. When asset price drops to below 60% of initial asset price, these agents (known as Stop-Loss agents) will begin placing large amounts of sell orders. For both groups of agents, the amount of sell orders each agent can issue is capped with a predetermined upper bound.

In all of our simulations, we fixed the agent composition at 18 ZI agents, 27 TF agents, and 9 AA agents, in order to represent a market in which there is significant presence of professional traders using algorithm-based techniques as well as those who are looking for arbitrage opportunities.

III. ANALYSIS OF SIMULATION RESULTS

This section contains a detailed analysis of our nine simulations.

A. Simulation Results

We have conducted nine different types of simulations based on a slowdown on market 1 followed by a slowdown in Market 2. In each case, we have plotted out the price history (for Stocks A, B, and C as well as the Index), the rolling exponentially-weighted volatility based on a λ value of 0.9 and the trading volume of each asset in 30-second buckets. The entire simulation lasted 900 seconds, which is comparable to the most active time period of the “flash crash” on May 6, 2010.

1. Simulation 1

The simulation shown in Figure 6 is based on a slowing down of Market 1 from 120 to 360 seconds and then a slowing down of Market 2 from 240 to 480 seconds. In the first case, we can see that prices collapsed, rolling volatilities spiked, and trading volumes picked up during the interval from 120 to 240 seconds and then during the interval from 400 to 600 seconds. This observation is consistent with our earlier research, in that the real problem appears to be caused by changing the “rules of the game” in the middle of trading, instead of the simple domination of the market by any specific type of traders. Since there are no changes to the fundamental

demand-and-supply balance during the simulation (except for the initial triggering of selling orders by Bear Market agents), the market will function properly once it is stabilized, but the subsequent reversion to normal speed of clearing once again create an imbalance of demand and supply leading to significant price instabilities. In addition, we observe that, in some cases, price actually hit the value of \$1, which is the value of stub quotes left by market-makers.

2. Simulation 2

The simulation shown in Figure 7 is based on a slowing-down of Market 1 from 120 to 240 seconds, and then Market 2 slowed down from 180 to 360 seconds. We are interested in understanding what may happen as and when we push the two slow-down periods closer together, emulating the cascading effects among unstable parallel markets. As expected, we no longer observe two distinct periods of shocks. Even more interesting are the observations that (a) the price-shock periods are compressed; as a result, there really isn't a sufficient time lag for supply and demand conditions in the market to recover from the first price shock before entering the second price shock; (b) prices go through an extended period of instability after the 360th second or the end of the second shock period; and (c) during the time when prices go through an extended period of instability, there continue to be many instances in which the Arbitrageur agents are unable to pull the Index back to its fair value. This is shown in Figure 15. Simulation 2 will be treated as our base scenario for testing other potential policy responses.

3. Simulation 3

The simulation shown in Figure 8 is based on imposing position limits by trader, instead of typical position limits by symbol (i.e., per stock trading on each individual exchange). Although not apparent from the descriptive statistics, the markets in this simulation experienced a significant increase in violent “up and down” shocks, and the price graph clearly shows signs of increased price instability. Readers should note that the type of extreme “up and down” shocks is actually consistent with the type of price movements shown on May 6. Those shocks are not observable with exchange data at the second-by-second level, but the authors have examined internal aggregated client data provided by a broker-dealer at the microsecond level showing exactly that type of extreme “up and down” shocks during the 2:30 to 3:30 p.m. EST period on May 6. The fact that these shocks actually become significantly more pronounced due to the imposition of position limits suggests that position limits are unlikely to have worked as an effective regulatory tool to eliminate “flash crash”-like symptoms.

4. Simulation 4

The simulation shown in Figure 9 is based on changing the clearing mechanism from continuous time auction to discrete time auction, which would have negated

any trade execution advantages of high-frequency, algorithm-based trading. The modified clearing mechanism does not mean that the algorithm-based traders cannot execute trades; it only means that certain traders do not have any speed advantage relative to other market players, so they will profit only when they can come up with a fundamentally superior trading strategy that is not based on more timely execution. Based on both the price graphs and the descriptive statistics, it is not obvious that negating the advantages of high-frequency trading can make any significant difference in maintaining market stability.


5. Simulation 5

The simulation shown in Figure 10 is based on Simulation 3, in which quotes are not updated during the slowdown. This simulation is designed to address the following question: Instead of a total and abrupt stoppage — which is generally considered by the market as a blunt and ineffective tool since it simply delays the resolution to any fundamental imbalances in supply and demand — what would have been another alternative to a simple “go slow” mode? The typical “go slow” mode bears a certain degree of resemblance to discrete time auctions, in that primarily the amount of through-put in the clearing process is slowed down. Therefore, it is natural to ask whether stopping the publishing of quotes will make any difference. Based on both the price graphs and the descriptive statistics, it is not obvious that stopping the publishing of quotes could have made any significant difference in maintaining market stability.

6. Simulations 6, 7, 8, and 9


The simulations shown in Figures 11, 12, 13, and 14 are based on Simulation 3, in which price limits are imposed when prices have dropped by more than 40%, 30%, 20%, and 10%, respectively, when compared to the base prices that are sampled from the “last done” prices every 60 seconds. As a result of imposing this new policy, there are significant decreases in the skewness, kurtosis, and maximum drawdown statistics, with more significant improvements as and when the trigger level is lowered. Readers should note that imposing price limits does not address any fundamental supply and demand imbalances. Such imbalance should result in a natural drop in prices until a new market equilibrium is found, instead of any extreme “up and down” shocks, which rarely result in genuine price discovery and the orderly resolution of excessive demands/supplies. Moreover, there are more extreme “up and down” shocks when the price limit trigger is set either too low (40%) or too high (10%) — that may mean that regulators are either intervening too late (thus not providing any relieves) or needlessly (potentially making the situation worse). The ideal trigger level seems to be between 20% and 30%, which is consistent with the intuitive expectations of some market practitioners. Although we started these simulations by modifying Simulation 3, agent-level position limits are not breached in almost all cases, so that in practical terms Simulation 2 should be considered our true base scenario for these four simulations.

Figure 6. Price, Exponentially-Weighted Volatility and Trading Volume.




Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 1 due to slow-down of Market 1 from 120 seconds to 360 seconds; Market 2 from 240 seconds to 480 seconds. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 7. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 2.




Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 2 due to slow-down of Market 1 from 120 seconds to 240 seconds; Market 2 from 180 seconds to 360 seconds. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 8. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 3.




Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 3 due to imposing position limits by trader, instead of typical position limits by symbols (i.e., per stock trading on each exchange). Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 9. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 4.




Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 4 after changing the clearing mechanism from continuous time auction to discrete time auction, which would negate any advantages of high-frequency trading. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 10. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 5.




Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 5, which is a variant of Simulation 3, but its quotes are not updated during the slowdown. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 11. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 6.



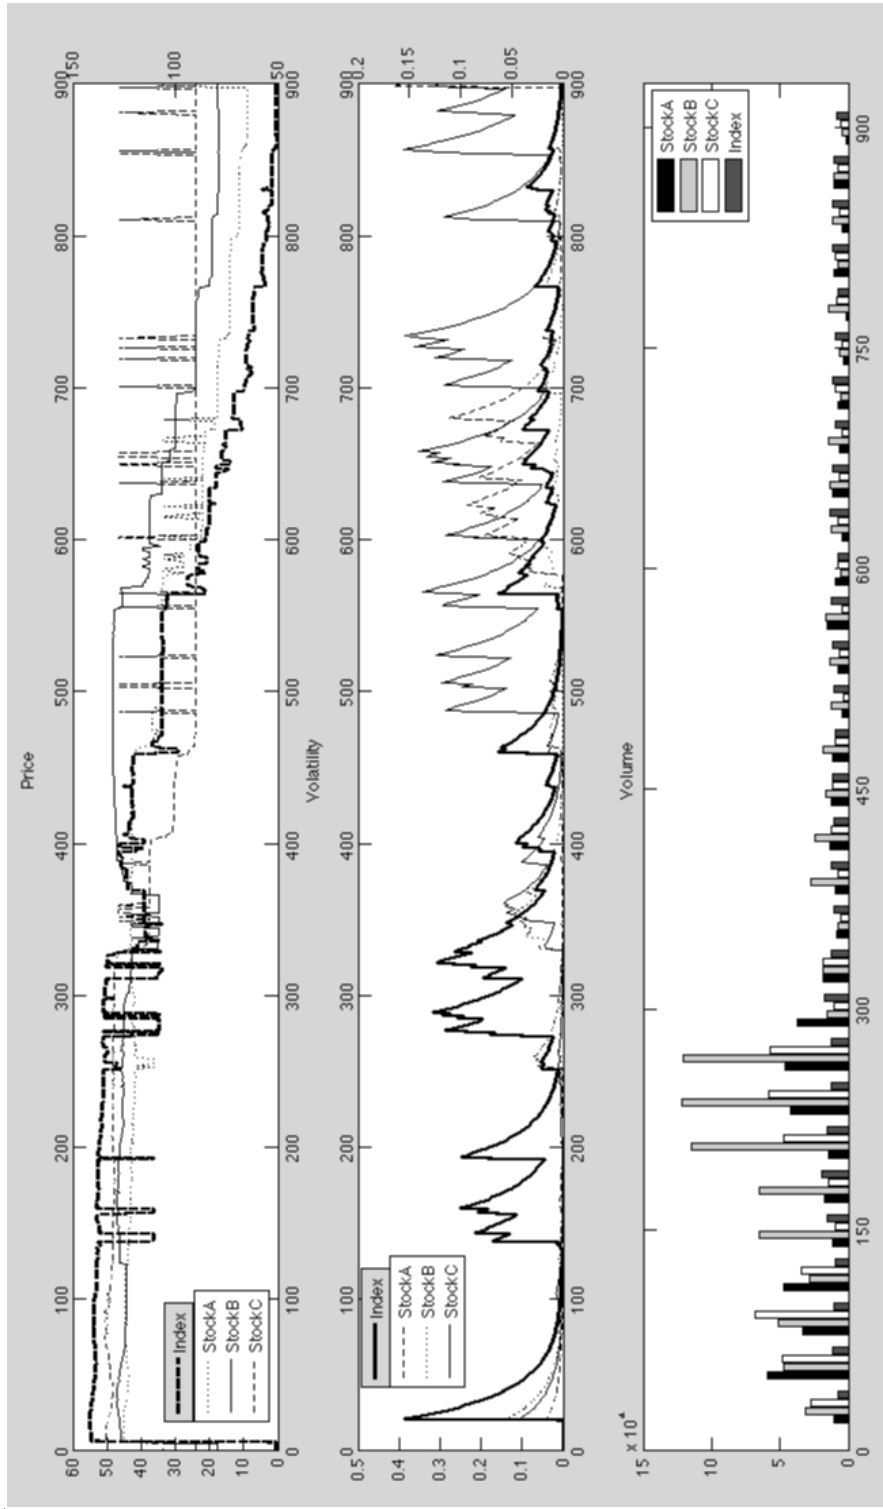
Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 6, which is a variant of Simulation 3, but price limits are imposed whenever prices have dropped by more than 40% when compared to the average of the last 5 trades. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 12. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 7.




Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 7, which is a variant of Simulation 3, but price limits are imposed whenever prices have dropped by more than 30% when compared to the average of the last 5 trades. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 13. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 8.




Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 8, which is a variant of Simulation 3, but price limits are imposed whenever prices have dropped by more than 20% when compared to the average of the last 5 trades. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 14. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 9.



Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 9, which is a variant of Simulation 3, but price limits are imposed whenever prices have dropped by more than 10% when compared to the average of the last 5 trades. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 15. Comparison of Synthetic Fair Value vs. Traded Index Values in Simulation 2.



B. Statistical Analysis

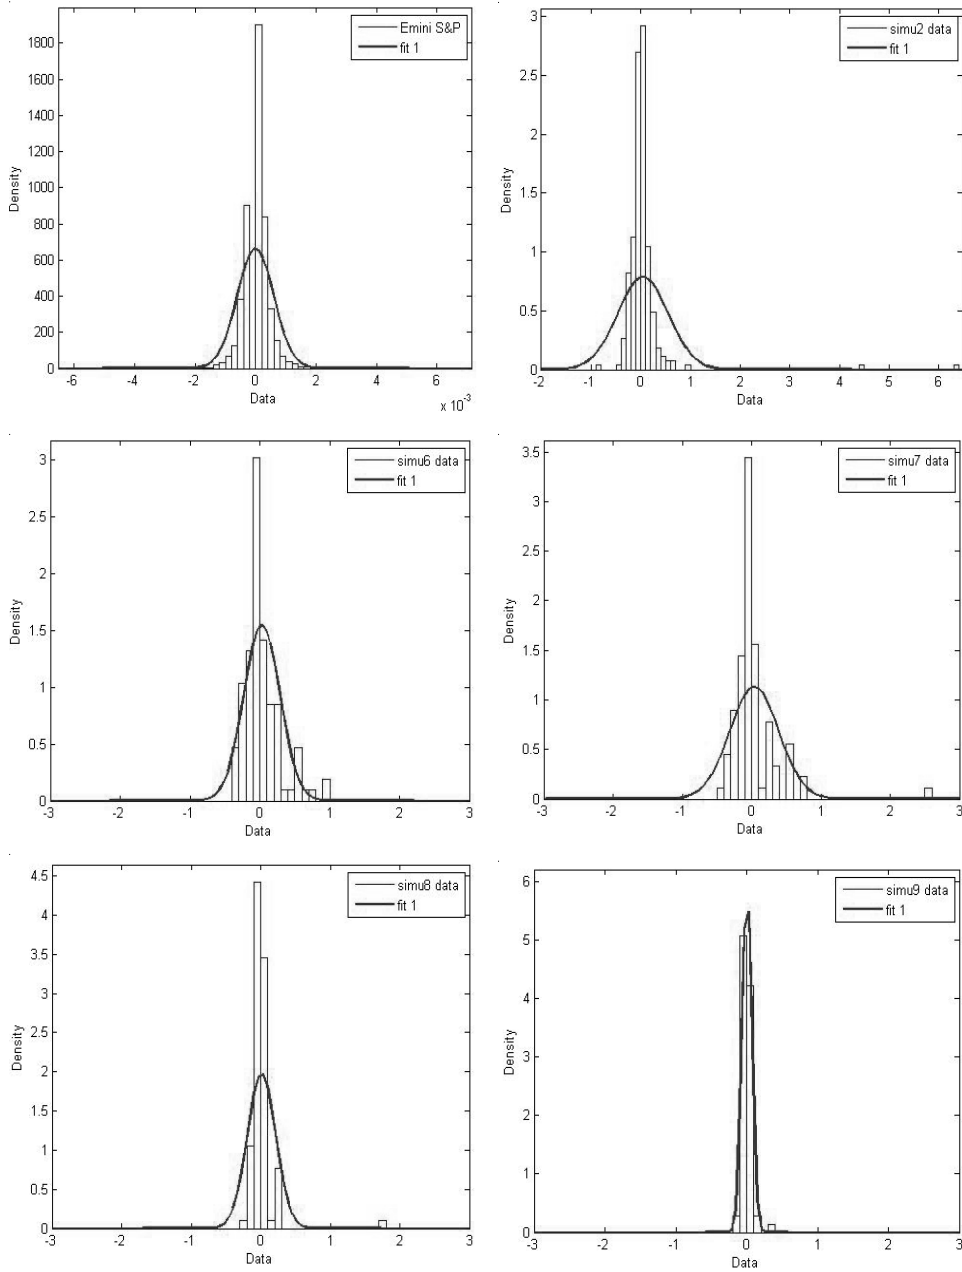
The summary statistics below is computed based on second-by-second data using absolute differences in returns on the Index. Because our simulated Index is composed of only 3 stocks instead of 500 securities in the SPX, the difference in base index values means that computing the proportional differences may produce non-comparable (if not non-sensical) results and in particular unreliable skewness statistics. Skewness and kurtosis are scale invariant, and the simulated skewness and kurtosis appear to be “close enough” when compared to those observed from the SPX E-mini futures market on May 6, 2010. Moreover, the minimum and maximum values of the simulations are roughly about 10 times the size of their corresponding standard deviations. That is not reasonable as compared to real-market returns on May 6, 2010 especially those of single-name stocks. (Refer to our earlier study for a further discussion on the challenges and goals in getting “close enough” when matching moments in simulating extreme market movements.)

The comparison is particularly striking when the outputs of these simulations are lined up side by side against typical fat-tail distributions created by a priori mathematical assumptions. Our assessment is that these simulations have produced price distributions with “reasonable resemblance” of the actual evolution of the prices on SPX E-mini futures from 2:30 to 5:00 p.m. EST on May 6, 2010; changing the observation window within the 30-minute time frame does not result in any dramatic changes to the descriptive statistics on the prices of the SPX E-mini futures.

Figure 16 plots out the comparative return distributions based on the SPX E-mini futures as well as the Index from Simulations 2, 6, 7, 8 and 9.² Readers should

2. To ensure an objective comparison, “zeros” have been deleted from the return distributions, as discussed in Lee et al. 2010.

Figure 16. Comparative Return Distributions.



Comparative return distributions based on the SPX E-mini futures as well as the Index from Simulations 2, 6, 7, 8 and 9.

Table 1. Descriptive Statistics on Stock A as well as the SPX E-mini Futures on May 6, 2010.

Stock	A									Emini			Accen- ture
	1	2	3	4	5	6	7	8	9	S&P*	P&G	3M	
Simulation	776	745	744	735	705	810	830	785	815	1794	1680	1346	656
Observations	45.64	42.28	35.04	43.34	48.16	35.43	34.06	32.92	46.30	1108.80	60.37	81.74	40.83
Mean	5.72	8.37	17.95	5.03	12.90	19.86	12.68	12.68	4.40	17.39	3.20	3.46	0.45
Stddev	0.32	-4.13	-0.66	-7.54	-3.14	-0.19	-1.05	-0.71	-0.33	-1.03	-3.58	-1.82	-0.24
Skewness	1.13	2.04	4.72	2.81	11.52	1.55	2.38	1.83	1.29	3.25	17.98	5.32	1.83
Kurtosis	40.40	1.00	1.00	1.00	1.00	4.90	7.40	9.10	40.30	1056.00	39.37	67.98	40.01
Min	53.90	47.00	52.10	46.10	54.90	58.90	44.00	45.10	52.10	1130.80	62.25	85.49	41.53
Max	13.50	46.00	51.10	45.10	53.90	54.00	36.60	36.00	11.80	74.75	22.88	17.51	1.52
Max - Min	-3.88	-20.09	-28.29	-8.15	-16.43	-10.05	-1.91	-3.65	-2.59	-1.54	-0.69	-1.63	-0.10
CVaR(95%)	0.22	0.98	0.98	0.98	0.98	0.84	0.30	0.48	0.18	0.01	0.17	0.09	0.01
MaxDD	13	38	43	7	21	29	13	18	10	0	3	0	0
 #(DD(>=10%))													

*Average of available bid and ask based on second-by-second data from 14:30:00 to 14:59:59 EST on May 6, 2010.

Table 2. Descriptive Statistics on Stock B as well as the SPX E-mini Futures on May 6, 2010.

Share	B	Emini									P&G	3M	Accer- ture	
		1	2	3	4	5	6	7	8	9				S&P*
Simulation		1	2	3	4	5	6	7	8	9				
Observations		820	804	762	787	729	798	847	807	816	1794	1680	1346	656
Mean		39.15	53.74	40.77	43.97	49.66	23.53	31.05	38.51	43.60	1108.80	60.37	81.74	40.83
Stddev		23.01	0.72	0.85	0.84	8.53	16.98	11.04	10.89	4.84	17.39	3.20	3.46	0.45
Skewness		-1.00	1.03	0.89	0.18	-4.27	-0.18	-0.62	-0.93	-0.96	-1.03	-3.58	-1.82	-0.24
Kurtosis		19.52	3.47	4.34	4.32	21.67	1.19	1.53	2.29	2.66	3.25	17.98	5.32	1.83
Min		1.00	52.70	39.90	42.70	1.00	1.00	11.20	17.30	30.70	1056.00	39.37	67.98	40.01
Max		55.90	56.10	43.00	46.10	54.00	42.00	41.00	48.40	49.20	1130.80	62.25	85.49	41.53
Max - Min		54.90	3.40	3.10	3.40	53.00	41.00	29.80	31.10	18.50	74.75	22.88	17.51	1.52
CVaR(95%)		-11.94	-0.59	-0.33	-1.00	-14.70	-4.79	-1.79	-1.94	-4.25	-1.54	-0.69	-1.63	-0.10
MaxDD		0.98	0.02	0.02	0.04	0.98	0.96	0.39	0.19	0.33	0.01	0.17	0.09	0.01
 #(DD(>=10%))		13	0	0	0	17	19	11	8	11	0	3	0	0

*Average of available bid and ask based on second-by-second data from 14:30:00 to 14:59:59 EST on May 6, 2010.


Table 3. Descriptive Statistics on Stock C as well as the SPX E-mini Futures on May 6, 2010.

Share	C	Emini								P&G	3M	Accen- ture		
		1	2	3	4	5	6	7	8				9	S&P*
Simulation		749	745	702	804	664	858	839	830	812	1794	1680	1346	656
Observations		33.65	40.70	41.79	39.68	44.85	36.66	32.72	35.00	35.96	1108.80	60.37	81.74	40.83
Mean		13.11	16.24	12.03	8.56	0.89	13.17	11.31	11.35	4.29	17.39	3.20	3.46	0.45
Stddev		-1.87	-1.72	-1.69	-2.70	0.81	0.23	-0.74	0.23	-0.88	-1.03	-3.58	-1.82	-0.24
Skewness		1.93	2.33	4.58	2.49	2.70	1.06	1.78	1.19	2.76	3.25	17.98	5.32	1.83
Kurtosis		1.00	1.00	1.00	1.00	43.80	24.50	12.90	23.90	25.20	1056.00	39.37	67.98	40.01
Min		41.90	51.10	50.00	45.00	47.30	53.30	42.90	50.50	41.10	1130.80	62.25	85.49	41.53
Max		40.90	50.10	49.00	44.00	3.50	28.80	30.00	26.60	15.90	74.75	22.88	17.51	1.52
Max - Min		-7.82	-20.49	-10.48	-22.44	-0.51	-14.68	-2.79	-11.69	-5.19	-1.54	-0.69	-1.63	-0.10
CVaR(95%)		0.97	0.98	0.98	0.98	0.03	0.51	0.27	0.49	0.33	0.01	0.17	0.09	0.01
MaxDD		11	23	13	30	0	24	18	25	32	0	3	0	0
 #(DD(>=10%))		*Average of available bid and ask based on second-by-second data from 14:30:00 to 14:59:59 EST on May 6, 2010.												

Table 4. Descriptive Statistics on Index as well as the SPX E-mini Futures on May 6, 2010.

Index	E-mini									P&G	3M	Accen- ture	
	1	2	3	4	5	6	7	8	9				S&P*
Simulation	684	622	637	711	714	749	746	744	752	1794	1680	1346	656
Observations	106.55	123.43	110.85	121.80	133.93	88.62	96.31	101.97	124.41	1108.80	60.37	81.74	40.83
Mean	40.31	29.40	30.59	16.69	25.91	45.59	33.12	31.76	11.77	17.39	3.20	3.46	0.45
Stddev	-0.86	-1.29	-0.70	-1.94	-1.90	0.28	-0.73	-0.31	-0.15	-1.03	-3.58	-1.82	-0.24
Kurtosis	61.30	2.03	9.87	7.05	6.47	1.40	1.90	1.60	1.74	3.25	17.98	5.32	1.83
Min	0.80	19.48	14.83	48.60	23.39	33.10	31.50	50.50	98.90	1056.00	39.37	67.98	40.01
Max	152.10	154.10	145.10	137.30	156.00	154.10	127.20	141.50	142.50	1130.80	62.25	85.49	41.53
Max - Min	151.30	134.62	130.27	88.70	132.61	121.00	95.70	91.00	43.60	74.75	22.88	17.51	1.52
CVaR(95%)	-41.74	-32.01	-28.45	-10.24	-35.88	-20.99	-17.25	-11.10	-4.49	-1.54	-0.69	-1.63	-0.10
MaxDD	0.99	0.82	0.86	0.44	0.80	0.40	0.43	0.20	0.15	0.01	0.17	0.09	0.01
 #(DD(>=10%))	52	60	58	12	43	28	23	12	2	0	3	0	0

*Average of available bid and ask based on second-by-second data from 14:30:00 to 14:59:59 EST on May 6, 2010.

Figure 17. Realized P&L in Simulation 1 for Different Agent Types.

Realized P&L in Simulation 1 for different agent types. (MM = Market Maker, ZI = Random Agents, TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller, SL = Stop-Loss Agents).


visually examine the degree of similarity between the return distribution in our base scenario of Simulation 2 and that from the SPX E-mini futures. Not surprisingly, their skewness (-1.29 for Simulation 2 vs. -1.03 for SPX E-mini) and kurtosis (2.03 for Simulation 2 vs. 3.25 for SPX E-mini) statistics are also quite close. This graph also shows how the base scenario evolves under the price limit triggers set at 40%, 30%, 20%, and 10%, with tighter and tighter fits against their corresponding normal distribution curves.

C. Agents P&Ls

We have plotted the realized and unrealized P&Ls for all agent types in Simulations 1 and 2 in Figures 17, 18, 19, and 20. From these base scenarios we make the following observations:

1. In the absence of market interventions, Market Makers almost always make profits by design of their trading algorithms.
2. Neither the ZI (or “random”) agents nor the trend follower TF agents are able to make consistent profits.

Figure 18. Unrealized P&L in Simulation 1 for Different Agent Types.




Unrealized P&L in Simulation 1 for different agent types. (MM= Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

3. As expected, Arbitrageurs may suffer heavy losses when the Index fails to converge to its fair values.
4. The Bear Market seller may or may not make any profits, depending on the market’s recovery path.
5. The Stop-Loss agents will almost always lose money in flash crash by selling at unusually low prices that consequently recover.

If trades are “busted” at a certain level, then the P&Ls of the Market Makers will become uncertain. Doing so is expected to have a highly negative impact on the Market Makers’ willingness to participate in the markets during flash crashes. Without their participation in such markets, the authors contend that (a) it will be even more likely for the market to break down faster when liquidity is withdrawn faster from the market and (b) it will be more difficult for the market to recover from the destabilizing effects of any “flash crash.”

In addition, the unrealized P&Ls for all agent types in Simulations 3, 4, 7, and 8 (Figures 21, 22, 23, and 24) show that:

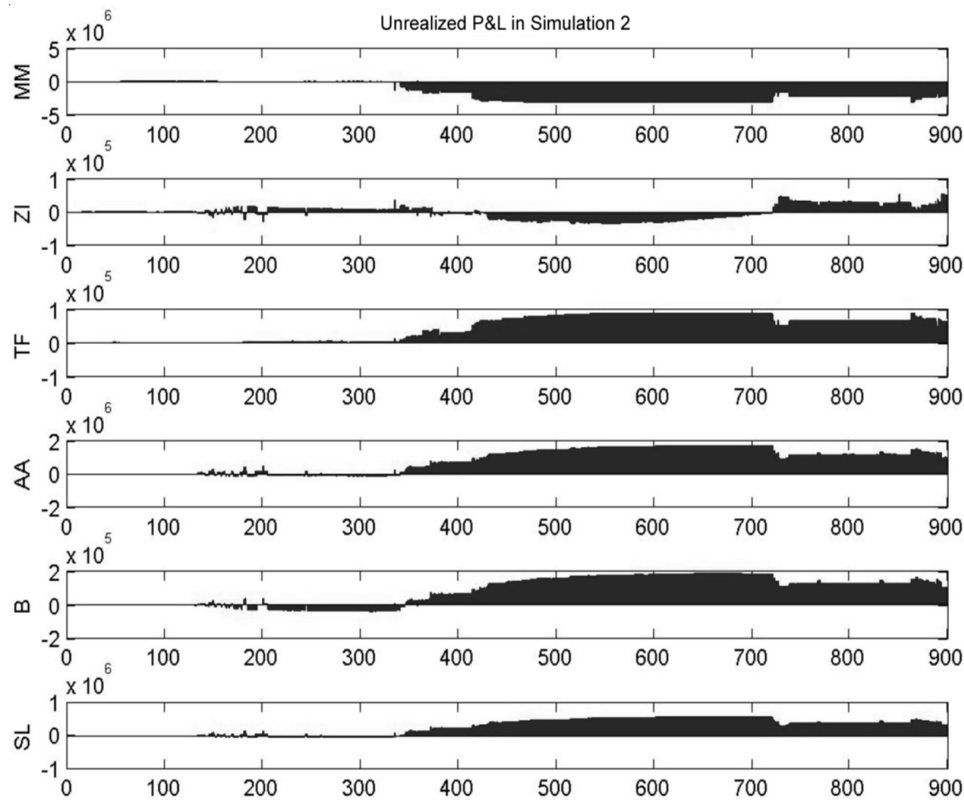
Figure 19. Realized P&L in Simulation 2 for Different Agent Types.

Realized P&L in Simulation 2 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

- Both imposing position limits by trader and changing the clearing mechanism from continuous time auction to discrete time auction may be ineffective in terms of eliminating “flash crash”-like symptoms, but these measures do not cause any unexpected changes to the P&L patterns among different types of market players.
- In Simulations 7 and 8 where price limits are imposed, it appears that certain professional traders are able to make profits at the expense of the Market Maker and to some extent the ZI (or “random”) agents.

Observation 2 is troubling, but not hugely surprising. When the market knows which direction a particular asset is going to trade because of regulatory intervention, professional traders can usually find ways to take advantage of the anticipated market movements. Market participants who are likely to be on the losing side of their trades will be the retail-like zero intelligence investors who typically deploy unsophisticated trading strategies assuming a fairly even distribution of market ups and downs, or market makers who are obligated to quote under the assumption that

Figure 20. Unrealized P&L in Simulation 2 for Different Agent Types.



Unrealized P&L in Simulation 2 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

bids and asks should be reasonably even and random. From a regulatory viewpoint, imposing price limits can be an effective policy to eliminate “flash crash”-like symptoms, but nonetheless one that may create unintended fairness issues for certain market participants.

1. “Busting” Trade

Finally, we used the base scenario of Simulation 2 to test the potential P&L impacts due to “busting trades” at or below 60% of the opening price of the asset traded:

1. If a long position is cancelled by the exchange after the trading session, then it is assumed that the agent has to “replace” the position at the asset’s closing price, resulting in a negative P&L impact.
2. If a short position is cancelled by the exchange after the trading session, then it is assumed that the agent has to “replace” the position at the asset’s closing price, resulting in a positive P&L impact.


Figure 21. Unrealized P&L in Simulation 3 for Different Agent Types.


Figure 21. Unrealized P&L in Simulation 3 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

The most interesting observation from Table 5 is that Market Makers and Zero-Intelligence end up bearing most of the impacts. These 2 agent types must quote or place trades based on the simple assumption that the bids and offers are evenly distributed. They are likely to suffer whenever there is a massive market adjustment in any one direction. Exchange officials should be aware of these unintended fairness issues before deploying the blunt tool to “bust” trades.

5. CONCLUSIONS AND RECOMMENDATIONS

The authors contend that the events of May 6, 2010 exhibit patterns consistent with the type of “flash crash” observed in their earlier study. While some commentators assigned blame on the May 6, 2010 “flash crash” to high-frequency trading, the authors suggest that the issue may be less about high-frequency trading per se, but rather the domination of market activities by trading strategies that are responding to the same set of market variables in similar ways, as well as various pre-existing schemes that modify the “rules of the game” in the middle of trading.

Figure 22. Unrealized P&L in Simulation 4 for Different Agent Types.




Unrealized P&L in Simulation 4 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

The consequent lack of market participants interested in the “other side” of their trades may result in a significant liquidity withdrawal during extreme market movements.

This paper describes an attempt to reconstruct the critical elements of the market events of May 6, 2010 based on the five hypotheses posed initially by the Joint CFTC-SEC Preliminary Report and the corresponding Final Report. The authors contend that the simulated asset price distributions have shown “reasonable resemblance” in descriptive statistics without over-fitting historical data.

Our specific recommendations are:

1. Any scheme to “slow down” trading does not address the fundamental demand and supply imbalance leading to flash crashes, and it may cause more problems than it solves.
2. In a “fragmented” market with parallel trading venues, the “action-reaction” nature of complex exchange rules to alter the speed of trading

Figure 23. Unrealized P&L in Simulation 7 for Different Agent Types.


Unrealized P&L in Simulation 7 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents)

may initiate a chain reaction that may drive liquidity further out of the aggregate market. Thus, it is important for parallel trading venues to coordinate their responses to avoid creating unintended domino effects.

3. The uneven slowing-down of trading at different trading venues often results in non-convergent fair values, because there is no or limited liquidity to complete one or more “legs” in an arbitrage trade. Arbitrageurs may suffer heavy losses in such markets, resulting in further withdrawal of liquidity due to their needs to “reverse out” from loss-making, incomplete arbitrage trades. Thus, it is important for parallel trading venues to coordinate the execution of their responses — in the event that going into a “slow mode” is the correct response, then its execution should be done in parallel by all relevant exchanges to avoid needlessly amplifying the uncertainties faced by market participants.

4. The problem appears to be less about the slowing-down of trading per

Figure 24. Unrealized P&L in Simulation 8 for Different Agent Types.



Unrealized P&L in Simulation 8 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents)

se. It is about the potential liquidity withdrawal due to the adjustments and chaos as a result of the initial slowing-down, as well as from the subsequent adjustments once the “normal” speed of trading is resumed.

5. “Busting trades” may discourage key participants such as Market Makers from trading in the markets as and when they are most needed. Unless there are clear technical errors involved, busting trades at arbitrary price levels is a blunt instrument that should be used sparingly and with extreme caution.

6. Price limits appear to be more effective than different implementation of positions limit in terms of stabilizing the market during the period of time when the market is finding its new equilibrium due to supply and demand imbalances.

7. Price limits do have limitations. When professional traders are

Table 5. Potential P&L Impacts of Different Agent Types.

Agent Type	Aggregated P&L without busted trades(\$)	Aggregated P&L with busted trades(\$)	Delta P&L(\$)
Market Maker (MM)	8,220,800	2,341.30	-8,218,458.70
Zero-Intelligence (ZI)	1,114,700	228,960.00	-885,740.00
Trend Follower (TF)	-5,930,600	184,590.00	6,115,190.00
Arbitrager (AA)	-132,040	-26,852.00	105,188.00
Bear Market (B)	-1,487,700	-148,520.00	1,339,180.00
Stop Loss (SL)	-1,581,800	-37,224.00	1,544,576.00

Potential P&L impacts of different agent types due to “busting trades” at 60% or below the opening price of each asset.

reasonably certain of potential market outcomes, they can normally find ways to make profits based on trading algorithms. That creates fairness issues for unsophisticated retail investors or market makers who are under obligations to quote. Therefore, the deployment of such blunt tools should be a regulatory policy of last resort.

References

- Cheng, S-F., 2007, Designing the Market Game for a Commodity Trading Simulation. *IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, San Jose, CA, November, 445-449.
- Easley, D., Lopez de Prado, M.M. and O'Hara, M., 2010, The Microstructure of the 'Flash Crash': Flow Toxicity, Liquidity Crashes and the Probability of Informed Trading. *The Journal of Portfolio Management*, 37(2), 118-128. Available at SSRN: <http://ssrn.com/abstract=1695041>.
- Findings Regarding the Market Events of May 6, 2010, 2010. *Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues*, September 30. Available from <http://www.cftc.gov/ucm/groups/public/@otherif/documents/ifdocs/staff-findings050610.pdf>.
- Kirilenko, A.A., Kyle, A.S., Samadi, M., and Tuzun, T., 2011, The Flash Crash: The Impact of High Frequency Trading on an Electronic Market. Available at SSRN: <http://ssrn.com/abstract=1686004>.

Lee, B., Cheng, S-F., and Koh, A., An Analysis of Extreme Price Shocks and Illiquidity among Systematic Trend Followers. *Review of Futures Markets*, 18, 385-419.

Preliminary Findings Regarding the Market Events of May 6, 2010, 2010. *Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues*, May 18. Available from <http://www.cftc.gov/ucm/groups/public/@otherif/documents/ifdocs/opa-jointreport-sec-051810.pdf>.

What Happened on May 6?, 2010. *CME Group Research Report*, May 10. Available from <http://www.cmegroup.com/files/May-6-2010-Incident.pdf>.

DIRECT MARKET ACCESS IN EXCHANGE-TRADED DERIVATIVES: EFFECTS OF ALGORITHMIC TRADING ON LIQUIDITY IN FUTURES MARKETS

Ahmet K. Karagozoglu*

Algorithmic trading (AT) and high frequency trading (HFT) afforded by direct market access (DMA) may have a greater impact on the exchange-traded derivatives markets than has been seen in the equity markets. This study breaks new ground to provide empirical evidence for the positive effects of AT on liquidity in the U.S. futures markets. To analyze the potential effects of electronic trading, this study provides an extensive review of the research in both equity and derivatives market microstructure. Using a unique dataset that directly and explicitly identifies algorithmic trading activity in exchange-traded derivatives, our research presents empirical evidence that AT decreases spreads (market width) and increases market depth in the Crude Oil, Euro FX, Eurodollar, S&P 500 E-mini, and 10-year U.S. Treasury Note futures contracts traded at the CME Group exchanges.

Electronic trading has been one of the most significant catalysts throughout the evolution of financial markets, especially for exchange-traded instruments. Emergence of electronic communication and/or crossing networks (ECNs) and their widespread use by various market participants resulted in a substantial change in the ownership and organizational structure of exchanges starting with the equity markets. Advances in technologies that directly impact trading in financial markets (e.g., telecommunication capacity, computational power) coupled with changes in the regulatory environment helped competitive market forces establish various trade execution venues. This increase in competition intensified the need to analyze and manage various components of trading costs and led to enhanced trading sophistication. As a result of these fundamental changes, techniques such as direct market access (DMA), smart order routing (SOR), algorithmic trading (AT), and high frequency trading (HFT) became the focus of attention for market participants,

*Ahmet K. Karagozoglu is an associate professor and Academic Director of the Martin B. Greenberg Trading Room, the Department of Finance, in the Zarb School of Business at Hofstra University. E-mail: finakk@hofstra.edu.

Acknowledgements: Author acknowledges the Institute for Financial Markets grant and summer research support of the Zarb School of Business at Hofstra University. Special thanks are due to John Labuszewski at the CME Group for providing the algorithmic trading and microstructure data as well as for his comments. Research assistance by Vanessa Singh and Xia Zhang is greatly appreciated.

Keywords: direct market access, algorithmic trading, liquidity, volatility, futures contracts
JEL Classification: G13, C13

exchanges, and regulators. Recently market and exchange characteristics of transparency, best execution, and latency have been the subject of research and analysis in addition to the more traditional factors of liquidity, volatility, and efficiency. Of course, given the recent turmoil in financial markets and high-profile losses, these factors have also attracted the attention of politicians and the public at large.

Extensive use of algorithmic trading (AT) activities emerged relatively more recently in the exchange-traded derivatives in comparison to the equity markets.¹ However, the impact of DMA, AT, and HFT on market quality and risk management may be more substantial for derivatives.² In order to analyze the potential effects of DMA, AT, and their resultant changes in exchange-traded derivatives markets, this study provides an extensive review of the research in both equity and derivatives market microstructure. Historically, exchanges in equity and derivatives markets had varying degrees of differences; however, the implementation of electronic trading has made these two markets more connected and trading practices are now more similar than ever before.

Based on a unique dataset that identifies algorithmic trading activity directly and explicitly, our research finds that AT decreases spreads and increases market depth in the Crude Oil, Euro FX, Eurodollar, S&P 500 E-mini, and 10-year U.S. Treasury Note futures contracts electronically traded at the CME Group exchanges. To the best of our knowledge, this study is the first to provide empirical evidence for effects of AT on liquidity in the U.S. futures markets. Similar to the findings for the U.S. equity markets by Hendershott, Jones, and Menkveld (2011) and for the German equity markets by Hendershott and Riordan (2009), we find that for the U.S. futures markets algorithmic trading has a positive effect on liquidity.

Section I presents an overview of concepts related to direct market access. Section II provides a review of the existing literature on equity and futures market microstructure; recent work on DMA, AT, and HFT; and draws conclusions for the exchange-traded derivatives markets. Section III describes the data used in this paper while section IV introduces the empirical methodology. Empirical results are discussed in section V and section VI offers conclusions.

I. OVERVIEW OF DIRECT MARKET ACCESS CONCEPTS

As with any major structural change and the emergence of new technology, the use of innovative trading technologies in financial markets had a profound impact on returns from short-term trading, long-term performance of investment portfolios, measurement and management of risk, as well as interconnectivity of various markets both domestically and globally. Market microstructure research (MMR) has focused

1. Electronic trading in CME's Globex platform started in 1992, and the Open Access Policy was implemented in 2000. The Open Access Policy allows customers to trade directly on CME Globex if their clearing firm provides a financial guarantee for their trading activity. This effectively means that CME provided DMA to investors starting in 2000. However, explicit identification of AT through "Tag 50" designation started more recently, in 2006.

2. The existence of multiple contract months and relatively more inter- and intra- market trading suggests that DMA, AT, and HFT may have a higher impact on the exchange-traded derivatives markets than on the equity markets.

on analyzing the effects of the changes in trading and execution rules, different trading venues, regulatory changes, impact of technological advances, and behavior of market participants in response to the developments in financial markets. MMR initially focused on equity markets primarily due to the availability of detailed transactions data and rapid changes in trading practices. Following the advent of electronic trading in derivatives markets, microstructure research focusing on exchange-traded derivatives, especially futures markets trading, increased significantly.

Similar to the developments in equity trading, participants in derivatives market are demanding more direct access to the markets (DMA) for reduced transaction costs, increased speed of executions, and decreased information leakage. As in the case of equities, electronic trading in futures enables the use of computers to execute trades, reducing errors as well as enabling more efficient post-trade reporting and analysis. Electronic trading in exchange-traded derivatives facilitates direct access to markets, which in turn allows algorithms to be used to generate quote updates and orders; eventually, increased sophistication and speed of trading systems — including exchanges' execution capabilities — leads to the high (and ultra-high) frequency trading.

DMA enables traders to connect directly to an exchange, using the exchange's native application programming interface (API) through its dedicated network.³ In its purest form, exchanges may provide DMA to market participants without explicit electronic order handling/authentication by intermediaries/brokers, called naked access. In other cases, intermediaries or brokerage houses facilitate DMA access. Different levels of DMA provided to various types of market participants have significant implications for transparency, fairness, and risk management.

Initially in equity markets, algorithmic trading (AT) referred to the use of computer programs to submit orders and execute trades in order to minimize the market impact costs. AT replicated the actions of human traders by determining the size and timing of purchases and sales of shares based on various mathematical models (algorithms).⁴ Contemporary AT encompasses almost all tasks that can be carried out by human market makers and traders. For example, posting of bid and/or ask quotes generated by computer models may be considered algorithmic market making and concurrent execution of several transactions across different assets/markets is algorithmic arbitrage. Additionally, electronic execution of trades to achieve various positions generated by financial models, both short- and longer-term investments in a range of assets, is also a form of algorithmic trading.

High frequency trading (HFT) occurs when the pace of transactions generated

3. Aitken, Harris, and Ji (2009) suggest that DMA is defined as electronic facilities that allow brokers to offer clients direct access to the exchange trading system through the broker's infrastructure without manual intervention by the broker.

4. Hendershott, Jones, and Menkveld (2011) provide a simple definition for algorithmic trading (AT) as "the use of computer algorithms to manage the trading process." They suggest that many observers view algorithms and AT from the standpoint of institutional buy-side investor and indicate correctly that "algorithms can also be used to formulate trading decisions and strategies as well as implement them."

by algorithms reaches a speed which human traders would not have been able to achieve.⁵ Increased competition and intensive use of AT and HFT necessitate that participants be physically closer to the order-matching engines of exchanges, creating the phenomenon called co-location.⁶ High frequency trading is a subset of algorithmic trading and AT is a subset of DMA activities. Direct market access includes “point-and-click” trading (e.g., by individual investors), automated trading activities that encompass low frequency trades, and the HFT with significantly large and fast submission of quotes and trades solely by computer programs.⁷

In an electronic trading environment in futures markets, DMA basically recreates the advantages of pit trading by allowing numerous market makers (locals) and traders to access and act on timely trade information. As a result, the efficiency of the pit environment is augmented with the use of technology in an electronic setting. DMA creates infinitely large electronic trading pits that can be interconnected in ways that were not possible in the physical pit-trading environment.

Another way to represent DMA from the point of view of an investor or a financial institution is that, rather than executing trades via a broker, trades are executed through a member of the exchange who has transaction privileges on the floor. In this case, co-locating could be analogous to such an individual or institution purchasing or renting the right to be physically present and trade at the floor of the exchange. The futures trading floor analogy for AT and HFT would be a local having beyond-human capabilities to analyze vast amounts of data, announce bids and asks with extreme rapidity, and confirm trades with others who could match his or her speed in announcing prices and quantities. In an electronic version of the above scenario, DMA, AT, HFT, and co-location enable access to prices and markets and offer the capabilities to transact that are not bound by location, distance, and human limitations. In this perspective, these new trading practices increase liquidity, decrease transaction costs, and improve the price discovery in exchange-traded derivatives markets.

The existence of multiple contract months and relatively more inter- and intra-market trading suggests that DMA as well as its by-products AT and HFT may have a higher impact on the exchange-traded derivatives markets than on the equity markets. Although there is a significant body of academic work in market microstructure research (MMR) covering both the equity and derivatives markets, empirical evidence on the effects of DMA, AT, and HFT in equity markets is new and limited. Even more, such research is very rare in exchange-traded derivatives markets.

Exchange-traded derivatives markets are in the process of experiencing the

5. Brogaard (2010) indicates that there are no clear and commonly accepted definitions for many of the terms in rapid trading and in computer controlled trading, and uses the definition HFT that Securities and Exchange Commission (SEC) uses, “professional traders acting in a proprietary capacity that engages in strategies that generate a large number of trades on a daily basis” (SEC, 2010, p. 3606).

6. SEC refers to co-location as “a service offered by trading centers that operate their own data centers and by third parties that host the matching engines of trading centers” (SEC, 2010, p. 3610).

7. We thank John Labuszewski at the CME Group for clarifying these subtle differences.

implementation of these innovative approaches at various levels. This research paper is intended to provide guidance to market participants, exchanges, and regulators by synthesizing the findings in equity MMR; the recent empirical work on the effects of DMA, AT, and HFT in stock markets; and microstructure research in derivatives markets. It presents empirical evidence on early stages of DMA and AT in futures markets and discusses the implications of these developments for exchange-traded derivatives markets.

II. REVIEW OF LITERATURE

Literature on direct market access and algorithmic trading in equity markets is limited and in exchange-traded derivatives markets, almost nonexistent. However, previous research focusing on various aspects of equity and derivatives market microstructure provides insights about how DMA, AT, and HFT impact derivatives trading.

A. Equity Market Microstructure

Considering the importance of price discovery and contributions of various market participants to this process, analyzing the relative informational advantages of these agents is important because DMA, AT, and HFT may cause changes in different agents' participation in trading while possibly altering the balance of asymmetric information.

It has been shown that electronic access to equity markets increases liquidity, reduces trade size, alters volatility, reduces returns to market making/specialist systems, and increases transparency. However, DMA may eventually lead to alternative trading venues and fragmentation of liquidity. Based on these findings, is there a chance that DMA, AT, and HFT will also result in the fragmented liquidity and creation of alternative execution venues observed in equity markets? If so, what might be the results of these changes in futures markets? Exchanges and regulators need to examine implications of such potential developments in exchange-traded derivatives markets.

Conrad, Johnson, and Wahal (2003) investigate the execution costs of trades sent to traditional and alternative trading systems in equity markets and conclude that orders sent to traditional brokers have higher execution costs than those executed by alternative trading systems such as electronic communication networks (ECNs). Barclay, Hendershott, and McCormick (2003) examine the competition among different trading venues in the United States and show that ECNs attract more informed orders than NASDAQ market makers.

Anand and Subrahmanyam (2008) compare the informational advantages of intermediaries with those of other investors using confidential transactions data from the Toronto Stock Exchange (TSX). They find that intermediaries account for greater price discovery than other institutional and individual investors, in spite of

8. They also note that TSX is a completely electronic and highly transparent environment, and in the context of individual stocks, the potential for informational effects is known to be stronger than in basket securities, derivatives, and futures indexes.

initiating fewer trades and volume.⁸ Their empirical results indicate that intermediaries contribute more to price discovery and hence tend to be more informed, even in a transparent electronic market where such an advantage is not driven by a privileged view of the market on a trading floor.

Saar (2001) shows that market intermediaries possess important order flow information that gives them an informational advantage. However, there is a possibility that the higher information share of market intermediaries may be a result of front running or stepping ahead by brokers. But Anand and Subrahmanyam (2008) investigate these activities and find no evidence of such trading by intermediaries on the TSX.

These findings suggest that with the increased use of DMA, AT, and HFT in derivatives markets, the informational role of intermediaries and entities with co-location privileges needs to be closely monitored for potential information asymmetry generation. The potential impact of DMA in terms of fragmenting liquidity in exchange-traded derivatives needs to be investigated. The nature of the intermediation provided by futures commission merchants (FCMs) may change, and, in turn, could equalize access to markets.

The influence of market transparency on market quality is investigated in several papers. Hendershott and Jones (2005) find that more transparency is associated with better market quality, which has been a crucial competitive advantage for ECNs in the United States. Bessembinder, Maxwell, and Venkataraman (2006) focus on the impact of transaction reporting on execution costs for corporate bonds and find a significant reduction of execution costs following the introduction of transaction reporting. Avgouleas and Degiannakis (2005) examine the impact of pre-trade transparency on market volume by using trading volume data before and after the introduction of a central order book at the London Stock Exchange (LSE). They conclude that when trading shifts from the quote-driven to the order-driven market structure, transparency increases significantly.

Bloomfield and O'Hara (2000) suggest that the demand for sunshine trading and order splitting reduces the competitive advantage of low-transparency markets; they question the long-term viability of transparent markets particularly in large, well-monitored markets with low information asymmetries where such regulated transparency may be of less value. Tuttle (2003) finds that NASDAQ traders tend to use hidden orders more in stocks with high idiosyncratic risk and high volatility, and he concludes that this is consistent with the idea that hidden orders reduce the adverse selection risk for liquidity providers. Tuttle's findings provide a competing hypothesis to Bloomfield and O'Hara that anonymity becomes more appealing when adverse selection risk and volatility are low, as this lowers the free option value of limit orders. Theissen (2002) also finds that, while the adverse selection component is larger in the anonymous electronic trading system in the German market for stocks of all sizes, small stocks also exhibit larger realized spreads when traded anonymously.

The implication of these results for the exchange-traded derivatives is that the level of transparency of the limit order book has a significant impact on the trading

costs for market participants with differential liquidity-related trading orientation. Given that there are multiple contract months and relatively more inter- and intra-market trading in derivatives markets, higher levels of limit order book transparency may be more desirable.

Anonymity plays a key role in market participants' trading strategies as part of their efforts to obtain best execution. In recent years, the SEC has been requiring higher standards of intermediary accountability in order execution practices, while exchanges are attempting to respond to market's demand for greater anonymity. Barclay et al. (2003) find that informed traders prefer using anonymous ECNs compared to transacting non-anonymously with NASDAQ dealers. Anecdotal evidence also indicates that institutional direct market access participants usually conduct their algorithmic trades anonymously. Furthermore, Frino, Johnstone, and Zheng (2010) examine whether the identity of a broker involved in transactions contains information. Using a sample of transactions from the Australian Stock Exchange — where broker identity is transparent — they provide evidence that consecutive buyer- and/or seller-initiated transactions by the same broker have a relatively high permanent price impact. Their findings imply that broker identity conveys information to market participants, and that markets in which broker identity is disclosed are likely to be more efficient.

Grammig, Schiereck, and Theissen (2001) find that for the German stock market the probability of informed trading is higher in the anonymous electronic trading system compared to the non-anonymous trading floor, while Reiss and Werner (2005) find that in London informed traders tend to go to the non-anonymous direct interdealer market. They conclude that adverse selection is less prevalent in anonymous brokered markets.

De Winne and D'hondt (2007) investigate why traders hide their orders and how other traders respond to hidden depth. Their empirical findings suggest that traders use hidden orders to manage both exposure risk and picking off risk. They show that hidden depth increases order aggressiveness, and when hidden depth is discovered, order submissions are adjusted to seize the opportunity for depth improvement, suggesting that either this hidden depth is not associated with informed trading or the risk of trading with an informed trader is offset by the improvement in depth. However, Anand and Weaver (2004) report that hidden quantity can be used to reduce price impact if the probability of non-execution is small. Pardo and Pascual (2007) show that the execution of hidden volume increases during periods of intense trading when aggressive orders are clustered. To minimize the non-execution risk, hidden order traders can wait for a higher trading aggressiveness on the opposite side of the market, reduce implicit trading costs, and find faster trading executions.

Comerton-Forde and Tang (2009) characterize the impact of anonymous orders in a limit order market where identity disclosure is voluntary. They find that anonymously initiated trades tend to be more informative than non-anonymous ones, with cumulative excess returns positively related to trade size and security activity levels. Their empirical results indicate that anonymous orders are traded at lower

spreads than non-anonymous orders only for the most actively traded stocks; market orders that are anonymous result in higher price impact (pointing to high adverse selection cost) and in lower realized spreads (suggesting lower order processing and inventory management costs) than non-anonymous market orders. They conclude that anonymous trading is dependent on the order aggressiveness and the type of order originator.

Increased use of the DMA to submit quote-revisions and orders generated by algorithms in exchange-traded derivatives is likely to increase the merits of allowing voluntary disclosure rules for specific futures markets and contract months. Given that many expiration (contract) months are traded in futures markets, DMA and AT increase the spread trading as well as pricing efficiency of deferred-month contracts. However, any adverse selection cost impact of anonymous orders in longer-dated contracts is likely to be transmitted to more liquid front-month contracts. Therefore, the optimal level of anonymity in algorithmic and high frequency trading in exchange-traded derivatives needs to be investigated.

Aitken et al. (2009) investigate trade-based manipulation, as proxied by the daily incidence of ramping alerts, in 34 security markets worldwide during the 2000–2005 period. They suggest that closing call auctions, direct market access, specific regulations, and real-time surveillance (RTS) procedures and enforcement assure better market integrity and enhance market efficiency.⁹ They conclude that reduction in liquidity caused by higher volatility affects the order submission of liquidity suppliers who submit orders less aggressively. Specifically, their findings indicate that direct market access (DMA) reduces ramping manipulation, which Aitkin et al. interpret as “DMA facilitates algorithmic countertrading strategies that can circumvent the pump and dump tactics of a ramping manipulator.” Cumming and Johan (2008) examine trading regulations with corresponding surveillance technology to monitor alerts and find that comprehensive rules prohibiting trade-based manipulation generate higher turnover and larger market caps.

These findings point to the importance of both pre- and post-trade real-time risk analysis. One possible solution is to co-locate the risk control algorithms of clearing houses and financial intermediaries with the exchanges’ trade-matching engines where the servers of market participants engaging in AT and HFT activities are co-locating. Also, a regulator or self-regulator algorithm trader might co-locate at that physical location in order to facilitate detection and rapid response to improper trading activity that might be taking place at extreme speeds.

B. Microstructure of Exchange-Traded Derivatives

A significant amount of research in exchange-traded derivatives markets focuses on the effects of the move from floor-based trading to electronic trading. Various authors study the effects of such a move on the liquidity, bid-ask spreads, trading

9. Cumming and Johan (2008) suggest that trading activity increases if exchanges adopt surveillance procedures and regulations that assure market integrity (similar to findings of Eleswarapu and Venkataraman 2006). Pagano and Schwartz (2003) and Comerton-Forde and Rydge (2006) investigate implementation of closing call auctions to improve market quality.

volume, and behavior of market participants in both U.S. and global exchanges. More recent articles focus on the changes in market structures and market quality using higher frequency trading and quote data in futures markets.

Liquidity costs are considerably lower in the electronic market than in the open outcry market (Shah and Brorsen 2010). Huang (2004) analyzes the determinants of bid-ask spreads for the Taiwan Futures Exchange (TAIFEX) and Singapore Exchange-Derivatives Trading (SGX-DT) futures and finds that volatility and the information asymmetry are the major factors affecting the spreads and that the information asymmetry component is significantly lower in the electronically traded TAIFEX contract than in the open-outcry SGX-DT futures.

Ates and Wang (2005), focusing on the electronic and floor-traded contracts based on S&P 500 and NASDAQ 100 indexes, investigate the relative efficiency in terms of contributions to price discovery and find that contribution of electronically traded contracts is higher. Tse and Zobotina (2001) examine the FTSE 100 index futures trading following the transition to electronic trading and find a decrease in bid-ask spreads; however, they also find that the open-outcry trading has higher market quality and higher information content.

Frino, Lepone, and Wearin (2008) study the intraday pattern of quoted depth in interest rate futures contracts traded at the Sidney Futures Exchange (SFE), which is a competitive dealer market, and find that depth is lowest at the open, considerably higher during the final hours of trading, and highest at the close, which is a pattern at odds with the ones observed in specialist markets. Their results show that an increase in quoted depth is due to a narrowing in bid-ask spreads, and they conclude that this observation at the close of trading is driven by dealers' rebalancing inventories.

Chung and Chiang (2006) examine the price clustering in the DJIA, S&P 500, and NASDAQ-100 index futures by comparing the electronically and floor-traded contracts and find that prices are significantly more clustered in open-outcry trading; they attribute this to higher levels of human participation in trading on the floor.

Frino et al. (2008) investigate the influence of large trades executed by outside customers on futures prices at the CME and find that the permanent price impact (information effect) of large buyer-initiated trades is greater than that of large seller-initiated trades, while the temporary price impact (liquidity effects) of seller-initiated trades is greater.

Chakravarty and Li (2003) find that dual traders in futures markets are informed and act as liquidity suppliers. Anand and Chakravarty (2007) analyze price discovery across trade sizes in options markets and find that small- and medium-size trades are responsible for the majority of price discovery.

Wagener and Riordan (2009) study the lead-lag effect between the Deutscher Aktien Index (DAX) spot index and DAX index futures under asymmetric latency in the exchange infrastructure by focusing on the introduction of the exchange electronic trading platform Xetra Release 8.0, which significantly reduced the trading latency. Their empirical results suggest that a decrease in relative latency between the Deutsche Börse systems Xetra and Eurex leads to a higher degree of market integration, and they conclude that "a significant improvement in the cash market

infrastructure cutting network latency reduces the execution risk.”¹⁰

Webb, Muthuswamy, and Segara (2007) investigate the frequency of market clearing and the changes in trading hours for stock index futures contracts at the TAIFEX and SGX to measure the effect of increases in clearing on the volatility of futures prices. They find that simultaneous opening times for the TAIFEX, which batches orders at the open, and the SGX, which does not, is associated with a significant reduction in the volatility in SGX.

Bortoli et al. (2006) investigate the effects of an increase in pre-trade

transparency on trading behavior in the Share Price Index (SPI) futures traded at the SFE. Their research covers the time period in 2001 when the exchange increased the limit order book disclosure from depth at the best bid-ask prices to depth at the three best bid-ask prices. They find a decline in depth at the best quotes and an increase in the proportion of market orders exceeding depth at the best quotes. Their conclusion is that when pre-trade transparency increases, “limit order traders charge market order traders a higher premium for execution certainty by withdrawing depth from the best quotes, but not by increasing bid-ask spreads.”

Tse, Xiang, and Fung (2006), investigating the Euro FX and Yen FX futures traded at the CME, show that electronic futures trading contributes more to price discovery than both online spot and floor futures trading while online spot trading dominates electronic futures. Cabrera, Wang, and Yang (2009) find that the Electronic Broking Services (EBS) electronic interdealer broker dominates both electronic and floor traded currency futures. Beckett (2010), using high frequency data on

Deutsche Börse. They find that AT affects liquidity almost equally in supply (when liquidity is expensive) and demand (when it is cheap), and they also show that algo trades and quotes are more informative than those generated by humans. They suggest that this is achieved by AT “placing more efficient quotes and demanding liquidity to move the prices towards the efficient price.” Chaboud et al. (2009) investigate the effects of AT in the spot foreign exchange markets and find that AT activity and volatility are not correlated, and that the order flow generated by AT does not affect the return variance.

Hendershott et al. (2011) investigate the impact of algorithmic trading on market liquidity by using the electronic message traffic as a proxy for algorithmic trading activity in the NYSE stocks and find that AT and liquidity are positively related. By considering the implementation of auto-quoting on the NYSE as an exogenous event, the authors show that algorithms result in more message traffic, and as quoted and effective spreads narrow adverse selection declines. They interpret this as an “indication that algorithmic trading does causally improve liquidity.”

Brogaard (2010) investigates the impact of high frequency traders on equities markets by considering how the strategies utilized are related to liquidity, price efficiency, and volatility. The study shows that contribution to price discovery of trades and quotes of HFT is greater than others and their activity reduces volatility. Empirical results indicate that high frequency traders demand liquidity at smaller order sizes and that trades surrounding a demanded HFT execute faster. These results suggest that high frequency trading does not increase volatility. Brogaard interprets these findings to suggest that “HFT plays a very important role in price efficiency and the price discovery process and high frequency trading provides more useful information to the price generation process.” Castura, Litzenberger, and Gorelick (2010), focusing on Russell 1000 and Russell 2000 stocks, investigate the impact of HFT on equity market quality. They find that while the ratio of HFT to total market activity is growing, equity markets appear to become more efficient with tighter spreads, greater liquidity at the inside, and less mean reversion of mid-market quotes; they correlate this with the growth in automation and speed on equity exchanges.

Hasbrouck and Saar (2009) find that, in electronic markets with the increase in AT, limit orders are cancelled very quickly, and they often correspond to modifications resulting in a new limit order at an updated price or in a market order. Hendershott et al. (2011) point out that the Regulation National Market System (Reg NMS) is designed to increase competition among liquidity suppliers, and their findings suggest that algorithmic liquidity suppliers play an important role in the supply of liquidity.

Chordia, Roll, and Subrahmanyam (2008) suggest that recent increases in trading volume and the reduction in the average trade size can be attributed to AT.¹² Garvey and Wu (2010) investigate the execution quality of electronic trading with

12. Brownlees, Cipollini, and Gallo (2010) develop a dynamic model for intraday volume which incorporates the existence of algorithmic trading.

geographically dispersed locations and trading speeds and find that “speed differences are costly to traders and that speed-advantaged traders engage in strategies that are more conducive to speed.”

Gerig and Michayluk (2010) develop a theoretical model that explains the increase in the high frequency automated trading volume. Their model shows that automated liquidity providers are able to price securities more precisely than traditional market makers so that they are able to transact the majority of order flow and cause prices to be more efficient. Model predictions also include that the informed investors’ profits decrease, uninformed investors lose less money, and trading activity of uninformed traders increases as a result of lower transaction costs.

Overall, empirical evidence to date suggests that the increased use of algorithmic and high frequency trading, facilitated by direct market access, has a positive effect on market liquidity in equity markets both domestically and globally. When this result is coupled with the lack of empirical evidence pointing to an increased price volatility attributed to AT and HFT, it is not too optimistic to expect that their impact is likely to be positive in exchange-traded derivatives markets as well.

III. DATA AND DESCRIPTIVE STATISTICS

A. Algorithmic Trading and Liquidity Measures

This study uses a unique dataset obtained from the CME Group for five futures contracts (Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note) traded at the CME Group exchanges. It includes several microstructure variables: percentage of volume attributed to automated trading systems in the specific market that day (*ATS*); percent of message traffic attributed to automated trading systems (*MSG*); the average bid-ask spread for a given size order during a trading day (*Width*); and the number of contracts displayed at the “top-of-the-book,” showing average size-in terms of contracts-of the best bid and best ask quotes in the limit order book (*Depth*).¹³

Among the many surveillance measures the CME Group’s market regulation division uses are the “Tag 50 ID” numbers to analyze the effect of algorithmic trading activities on the liquidity and quality of futures and options contracts traded on its exchanges (CME, CBOT, NYMEX, and COMEX). Identification of algorithmic trading activity “is facilitated by CME Globex policy that requires automated trading systems (ATs) to declare themselves as such” where ATS is referred to as “a system that automates the generation and routing of orders to Globex.”¹⁴

Market participants trading at the CME Group exchanges are required by the

13. CME Group, Algorithmic Trading and Market Dynamics, July 15, 2010. CME refers to the Depth variable as market resilience, which is the average width of the bid–offer spread for a specified size order. Depth is defined as the number of contracts on average at the “top of the book” or best bid or offer.

14. CME Group, Algorithmic Trading and Market Dynamics, July 15, 2010.

CME Group Rule 576 to include an operator ID, also referred to as the “Tag 50 ID” or “User ID” with each order they enter into the CME Globex electronic trading system.¹⁵ Although CME required its members who use algorithmic trading systems (ATS) to identify themselves with the “Tag 50 ID” starting in 2006, full implementation by all trading systems was not immediate. Therefore, microstructure data on *ATS* and *MSG* variables appear to be more reliable after mid 2008. As a result, this study covers the time period May 1, 2008, to May 27, 2010.¹⁶

The uniqueness of the dataset used in this study is due to the explicit identification of algorithmic trading (AT) volume, which is the proportion of executed orders originated from an ATS compared to the total electronic orders executed (*variable ATS*). CME Group data also provides the proportional volume of electronic message traffic attributed to ATS (*variable MSG*). Identification of the amount of electronic messages generated by AT, in addition to the actual AT trades, is necessary because the literature and anecdotal evidence indicate that ATs generate a large amount of bid and ask quotes which they cancel/lift over a short horizon. We believe that our study is the first to use such detailed identifiers of AT in exchange-traded U.S. derivatives markets.

B. Price and Trading Data on Futures Contracts

Daily open, high, low, and settlement prices, the daily total trading volume (*TrdVolu*), and open interest (*OpInt*) for the five contracts under investigation are obtained from the Reuters/CRB database. The Reuters/CRB database also contains the implied volatility (*ImpVola*) for each of the contracts based on the near-the-money futures options and the 200-day rolling historical volatility measure (*HisVola*).

C. Market Control Variables

In order to control for changes in the market conditions, various other variables are extracted from the Reuters/CRB database: AAA-corporate bond yield (*CorpAAA*); BAA corporate bond yield (*CorpBAA*); corporate credit spread (*CorpSprd* = *CorpBAA* – *CorpAAA*); yield on 3-month Treasury Bill (*Tbill3mo*); difference between the AAA-corporate bond yield and the yield on 10-year Treasury Note (*DefSprd*); difference between the yields on 10-year Treasury Note and the 3-month Treasury Bill (*TermSprd*); daily stock index levels for Dow Jones Industrial Average (*DOW*), NASDAQ composite (*NASDAQ*), New York Stock Exchange Composite (*NYSE*), Russell 1000 (*Russell1000*), and S&P 500 (*SP500*); daily values of Goldman Sachs Commodity Index (*GSCI*), U.S. Dollar Index (*DollarInd*),

15. See CME Group, Market Regulation Advisory Notice RA0915-5, “Operator ID (‘Tag 50 ID’) Required on All CME Globex Orders.” These IDs are “unique to the party who entered the order. For orders entered manually, the Tag 50 ID must be unique to the individual entering the order into CME Globex. For orders entered by an automated trading system (‘ATS’), the Tag 50 ID must be unique to the person, or the identified team of persons on the same shift, who are responsible for the operation of the ATS. All Tag 50 IDs must be unique at the level of the clearing member firm” (p. 1).

16. The data for the *ATS*, *MSG*, *Width*, and *Depth* variables are from the regular trading hours.

spot Gold price (*GOLD*), Reuters/CRB Commodity Index (ReutersCRBind), and the CBOE's Volatility Index (*VIX*).¹⁷

Table 1 presents the descriptive statistics on the futures microstructure variables. Percentage of trading volume from algorithmic trading systems appears to be highest in Euro FX (72.17%) and lowest in Crude Oil (32.43%) while for other contracts *ATS* ranges from 40% to 50%. A possible explanation for this observation is the existence of a highly liquid, electronic market for FX forwards that facilitates high frequency cross-market and cross-currency trades.¹⁸ Figure 1 displays the relative *ATS* and its time variation for the five contracts. Results for the percentage of electronic message (*MSG*) traffic emanating from AT indicate that the Euro FX contract has the highest proportion (88.33%) while the Eurodollar contract attains the lowest (55.87%). This suggests that almost half of the electronic message traffic in Eurodollar futures is generated by non-algorithmic activity. Figure 2 shows the *MSG* and its time-variation. Figure 3 graphs the *ATS* and Figure 4 graphs the *MSG*.¹⁹

Observations for the *Width* (bid-ask spread) and market *Depth* indicate that Eurodollar futures has the smallest width and largest depth among the five contracts, suggesting that the high liquidity of this contract attracts more "human" electronic orders/quotes, which tend to be revised more frequently than the ones from algorithms. We observe that the Crude Oil contract has the widest spread and least depth. Crude Oil futures did not start trading on an electronic system as early as other financial futures such as Euro FX and E-mini S&P 500. Spread trading is more prevalent in a physical commodity market such as crude oil, and spreads move more slowly compared to the outright futures prices. These market-specific characteristics may explain the relatively low algorithmic trading activity in the Crude Oil contract, and as a result its low liquidity can be attributed to limited electronic cross-market and cross-commodity trading. There are relatively more liquid and electronic cross-market and cross-asset trading possibilities for both E-mini and Treasury note futures. Figures 5 and 6 display the *Width* and *Depth* across five contracts and their time variation. These two graphs show the relative increases in spreads and decreases in market depth during the third quarter of 2008 as a result of the recent financial crisis.

Descriptive statistics for the trading volume, open interest, and volatility variables are provided in Table 2. In order to understand variation in the market variables prior to the start of our microstructure data period, comparison of these statistics for two time periods is presented: the "before" period is April 10, 2006, to April 30, 2008; the "after" period is May 1, 2008, to May 27, 2010.²⁰ Figures 7 and 8 graph

17. These control variables chosen to take into account the changes in the commodity, corporate debt, credit, currency, energy, equity fixed-income markets as well as the changes in volatility.

18. Findings of Tse, Xiang, and Fung (2006) and Cabrera, Wang, and Yang (2009) may point to this interpretation.

19. Figure 3 graphs the *ATS* and Figure 4 graphs the *MSG* approximately one month before and after May 6, 2010, the day referred to as the "Flash Crash." A casual inspection of these figures does not suggest an extraordinary change in *ATS* and *MSG* on that day.

20. Mean and median of market variables (using both parametric and non-parametric tests) are found to be different during the 2-year period before and after May 1, 2008 (except for mean of GSCI).


Table 1. Descriptive Statistics on Futures Microstructure Variables: ATS, MSG, Width and Depth, May 1, 2008, to May 27, 2010.

	<u>CrudeOil</u>	<u>EuroFX</u>	<u>EuroDollar</u>	<u>EminiSP</u>	<u>Tnote10yr</u>
Mean	32.43%	72.17%	44.10%	48.09%	47.48%
Median	31.50%	72.89%	42.98%	47.91%	48.03%
Max	43.50%	80.97%	56.42%	59.22%	58.08%
Min	23.95%	55.24%	31.71%	36.56%	26.69%
Std. Dev.	3.98%	4.32%	5.80%	4.02%	5.35%
Skewness	0.3550	-0.9213	0.1340	0.1039	-1.0271
Kurtosis	2.2348	3.9365	1.8630	2.6756	4.5370
<i>MSG</i>					
	<u>CrudeOil</u>	<u>EuroFX</u>	<u>EuroDollar</u>	<u>EminiSP</u>	<u>Tnote10yr</u>
Mean	70.67%	88.33%	55.87%	71.48%	65.89%
Median	68.96%	89.01%	55.12%	71.51%	66.57%
Max	85.65%	95.07%	85.65%	81.44%	84.18%
Min	57.74%	75.07%	21.53%	59.47%	48.20%
Std. Dev.	6.12%	3.83%	7.36%	3.78%	4.88%
Skewness	0.5050	-0.7179	0.2774	-0.0478	-0.1736
Kurtosis	2.1064	3.0393	4.9103	2.7085	3.4623
<i>Width</i>					
	<u>CrudeOil</u>	<u>EuroFX</u>	<u>EuroDollar</u>	<u>EminiSP</u>	<u>Tnote10yr</u>
Mean	48.08349	22.80801	18.34247	21.60916	28.57527
Median	41.35478	18.7203	13.7488	20.74676	25.93447
Max	107.8332	75.27579	58.69449	62.13548	95.25045
Min	13.53045	13.0642	12.59267	12.50082	15.63671
Std. Dev.	18.97225	9.878861	9.992397	9.022524	13.55803
Skewness	0.6779	1.2075	2.0948	1.3547	1.2368
Kurtosis	2.4031	4.5995	6.3964	5.2716	4.5477
<i>Depth</i>					
	<u>CrudeOil</u>	<u>EuroFX</u>	<u>EuroDollar</u>	<u>EminiSP</u>	<u>Tnote10yr</u>
Mean	6.10853	21.53783	1279.785	397.1073	409.0063
Median	6.051945	21.44758	717.1916	348.3574	343.4008
Max	11.13911	48.83141	10062.65	1244.024	1350.825
Min	3.20584	6.041679	93.03325	68.40597	75.09946
Std. Dev.	1.936459	9.209597	1723.291	213.2097	264.3048
Skewness	0.4160	0.2744	3.0227	1.1495	1.1995
Kurtosis	2.1246	2.2577	12.5923	4.5287	4.1469

Note: *ATS* is the percentage of volume attributed to automated trading systems in the specific market that day; *MSG* is the percent of message traffic attributed to automated trading systems; *Width* is the average bid-ask spread for a given size order during a trading day; *Depth* is the number of contracts displayed at the “top-of-the-book” (i.e., average size-in terms of contracts-of the best bid and best ask quotes in the limit order book). The data for the *ATS*, *MSG*, *Width*, and *Depth* variables are from regular trading hours.


Data source: CME Group

Figure 1. Proportion of Electronically Executed Orders Originated by Algorithmic Trading.




Proportion of Electronically Executed Orders Originated by Algorithmic Trading, ATS by Contract: May 1, 2008, to May 27, 2010 for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts. Data source: CME Group; the data for the *ATS*, *MSG*, *Width*, and *Depth* variables are from the regular trading hours.

Figure 2. Proportion of Electronic Messages (orders/quotes) Emanating from Algorithmic Trading (MSG).



Proportion of Electronic Messages (orders/quotes) Emanating from Algorithmic Trading (MSG) by Contract: May 1, 2008, to May 27, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts. Data source: CME Group; the data for the *ATS*, *MSG*, *Width* and *Depth* variables are from the regular trading hours.


Figure 3. Proportion of Electronically Executed Orders Originated by Algorithmic Trading.



Proportion of Electronically Executed Orders Originated by Algorithmic Trading for the Period April 1, 2010, to May 27, 2010, showing "Flash Crash" of May 6, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts.

Data source: CME Group; the data for the *ATS*, *MSG*, *Width* and *Depth* variables are from the regular trading hours.


Figure 4. Proportion of Electronic Messages (orders/quotes) Emanating from Algorithmic Trading.



Proportion of Electronic Messages (orders/quotes) Emanating from Algorithmic Trading for the Period April 1, 2010, to May 27, 2018, showing "Flash Crash" of May 6, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts.

Data source: CME Group; the data for the *ATS*, *MSG*, *Width* and *Depth* variables are from the regular trading hours.


Figure 5. Market Width for the Period May 1, 2008, to May 27, 2010.



Market Width for the Period May 1, 2008, to May 27, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts.

Data source: CME Group; the data for the *ATS*, *MSG*, *Width* and *Depth* variables are from the regular trading hours.

Figure 6. Market Depth for the Period May 1, 2008, to May 27, 2010.



Market Depth for the Period May 1, 2008, to May 27, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts.

Data source: CME Group; the data for the *ATS*, *MSG*, *Width* and *Depth* variables are from the regular trading hours.

Table 2. Descriptive Statistics on Five Futures Contracts: Daily Volume, Open Interest, and Volatility.

(1) Time	CrudeOil	Trading Volume & Open Interest			Volatility Implied & Historical			Intraday Volatility		
		<u>Ln(TrdVolu)</u>	<u>Ln(OpInt)</u>	-	<u>ImpVola</u>	<u>HisVola</u>	-	<u>GarKla</u>	<u>Pakinson</u>	<u>Range</u>
Before	Mean	12.87936	14.07090	0.31037	0.26492	0.01386	0.00032	0.02782	0.00032	0.00032
	Median	12.95356	14.10261	0.29780	0.26180	0.01291	0.00024	0.02585	0.00025	0.00025
	Std. Dev.	0.39796	0.12323	0.04809	0.06619	0.00517	0.00028	0.01052	0.00037	0.00037
	Skewness	-1.02338	-0.57196	1.04733	0.68003	1.45966	3.47890	1.47393	9.75343	9.75343
	Kurtosis	5.43091	2.29499	4.45689	3.46680	6.88324	23.56477	6.81036	148.58290	148.58290
After	Mean	13.21007	14.02149	0.52866	0.45146	0.02365	0.00105	0.04692	0.00103	0.00103
	Median	13.21756	14.00876	0.45900	0.36910	0.01934	0.00055	0.03900	0.00056	0.00056
	Std. Dev.	0.29514	0.07379	0.20988	0.21182	0.01410	0.00138	0.02675	0.00137	0.00137
	Skewness	-0.65126	0.44812	1.07943	0.93223	1.82701	3.97193	1.54600	3.78359	3.78359
	Kurtosis	6.31197	2.72576	3.16691	2.64784	7.95801	30.94680	6.35204	24.69317	24.69317

Table 2, continued. Descriptive Statistics on Five Futures Contracts: Daily Volume, Open Interest, and Volatility.

(2) EuroFX Time	Trading Volume & Open Interest			Volatility Implied & Historical			Intraday Volatility		
	<u>Ln(TrdVolu)</u>	<u>Ln(OnInt)</u>	<u>ImpVola</u>	<u>HsVola</u>	<u>GarKla</u>	<u>Pakinson</u>	<u>Range</u>	<u>RSY94</u>	
Before									
Mean	12.01282	12.16782	0.07649	0.06626	0.00359	0.00002	0.00720	0.00002	
Median	12.03239	12.18965	0.07420	0.06250	0.00334	0.00002	0.00669	0.00002	
Std. Dev.	0.33997	0.14822	0.01750	0.01862	0.00153	0.00002	0.00308	0.00002	
Skewness	-0.87207	-0.29675	0.51258	0.40865	1.09825	2.16552	1.09303	4.38765	
Kurtosis	7.31627	2.40108	2.62918	2.03892	4.18738	8.60837	4.16081	36.96949	
After									
Mean	12.32943	11.99436	0.13971	0.12240	0.00648	0.00008	0.01304	0.00008	
Median	12.33109	12.01053	0.11620	0.10530	0.00565	0.00005	0.01133	0.00005	
Std. Dev.	0.40165	0.23707	0.04953	0.04336	0.00323	0.00009	0.00654	0.00010	
Skewness	-1.22556	0.26549	1.25194	0.80376	1.54021	2.99151	1.54917	3.25996	
Kurtosis	10.71546	2.76698	4.00326	2.50596	5.83391	13.93990	5.86484	17.20603	

Table 2, continued. Descriptive Statistics on Five Futures Contracts: Daily Volume, Open Interest, and Volatility.

Time	Statistic	Trading Volume & Open Interest		Volatility				Range	RSY94
		<u>Ln(TrdVolu)</u>	<u>Ln(OpInt)</u>	<u>ImpVola</u>	<u>HsVola</u>	<u>GarKla</u>	<u>Pakinson</u>		
(3) EuroDollar	Mean	14.61597	16.16021	0.12799	0.13504	0.00027	0.00000	0.00113	0.00000
	Median	14.63891	16.14970	0.07510	0.07575	0.00018	0.00000	0.00100	0.00000
	Std. Dev.	0.40724	0.07360	0.10976	0.17608	0.00030	0.00000	0.00063	0.00000
	Skewness	-0.88732	0.15756	1.24030	2.27956	3.13282	3.47453	1.52332	4.57961
	Kurtosis	5.78757	2.02780	3.79660	7.64453	19.34330	19.57783	6.19571	34.65420
After	Mean	14.40346	15.82571	0.72905	0.58973	0.00034	0.00000	0.00144	0.00000
	Median	14.40958	15.80338	0.78230	0.54810	0.00023	0.00000	0.00117	0.00000
	Std. Dev.	0.42409	0.14628	0.30475	0.32711	0.00040	0.00000	0.00107	0.00001
	Skewness	-1.91023	0.48073	-0.24035	0.87563	3.01185	19.06164	6.34548	21.91379
	Kurtosis	14.69838	2.22836	2.66144	3.66717	15.17884	404.52580	78.41297	493.27890

Table 2, continued. Descriptive Statistics on Five Futures Contracts: Daily Volume, Open Interest, and Volatility.

Time	Statistic	Trading Volume & Open Interest				Volatility				Intraday Volatility			
		<u>Ln(TrdVolu)</u>	<u>Ln(OpInt)</u>	<u>ImpVola</u>	<u>HisVola</u>	<u>GarKla</u>	<u>Pakinson</u>	<u>Range</u>	<u>RSY94</u>	<u>GarKla</u>	<u>Pakinson</u>	<u>Range</u>	<u>RSY94</u>
(4) Emini S&P 500	Mean	14.13279	14.42281	0.15564	0.13727	0.00678	0.00010	0.01370	0.00010	0.00010	0.01370	0.00010	
	Median	14.08976	14.44146	0.13480	0.13130	0.00553	0.00004	0.01110	0.00005	0.00004	0.01110	0.00005	
	Std. Dev.	0.49210	0.20364	0.05956	0.06304	0.00431	0.00015	0.00876	0.00020	0.00015	0.00876	0.00020	
	Skewness	-0.11723	-0.03448	0.71801	0.56820	1.89127	4.32681	1.88043	7.95242	4.32681	1.88043	7.95242	
	Kurtosis	3.19595	2.36717	2.36358	2.35308	8.09170	29.54153	7.89987	86.37391	8.09170	29.54153	7.89987	
After	Mean	14.59885	14.77799	0.28511	0.27249	0.01326	0.00041	0.02688	0.00039	0.00041	0.02688	0.00039	
	Median	14.59838	14.75500	0.22930	0.20040	0.01017	0.00015	0.02044	0.00015	0.00015	0.02044	0.00015	
	Std. Dev.	0.41200	0.12662	0.13799	0.19326	0.00991	0.00077	0.02051	0.00077	0.00077	0.02051	0.00077	
	Skewness	-1.81441	0.71461	1.49343	1.75631	2.08392	3.92632	2.16004	4.62776	3.92632	2.16004	4.62776	
	Kurtosis	12.97335	3.67853	5.21569	5.49366	7.88584	20.86218	8.34482	29.37146	7.88584	20.86218	8.34482	

Table 2, continued. Descriptive Statistics on Five Futures Contracts: Daily Volume, Open Interest, and Volatility.

Time	Statistic	Trading Volume & Open Interest		Volatility			Intraday Volatility		
		<u>Ln(TrdVolu)</u>	<u>Ln(OpInt)</u>	<u>ImpVola</u>	<u>HisVola</u>	<u>GarKla</u>	<u>Pakinson</u>	<u>Range</u>	<u>RSY94</u>
(5) Treasury Note (10-year)	Mean	13.93702	14.69171	0.05746	0.05007	0.00258	0.00001	0.00530	0.00001
	Median	13.95781	14.67468	0.05070	0.04340	0.00221	0.00001	0.00462	0.00001
	Std. Dev.	0.45410	0.10604	0.01885	0.02062	0.00138	0.00002	0.00285	0.00002
	Skewness	-0.49608	0.41215	0.81686	1.21966	1.53552	3.51397	1.57170	5.06524
	Kurtosis	4.75827	2.49628	2.57332	3.72802	6.24795	20.79764	6.35896	45.21718
Before	Mean	13.56462	14.09699	0.08919	0.08430	0.00408	0.00003	0.00837	0.00003
	Median	13.60130	14.03630	0.08385	0.07840	0.00363	0.00002	0.00756	0.00002
	Std. Dev.	0.52420	0.23324	0.02445	0.03132	0.00210	0.00005	0.00429	0.00004
	Skewness	-1.42213	0.56219	0.48520	1.04812	2.20825	7.72542	2.38856	4.71498
	Kurtosis	9.65794	2.08146	2.21400	3.39042	11.86176	92.58063	15.02140	33.75686
After	Mean	13.56462	14.09699	0.08919	0.08430	0.00408	0.00003	0.00837	0.00003
	Median	13.60130	14.03630	0.08385	0.07840	0.00363	0.00002	0.00756	0.00002
	Std. Dev.	0.52420	0.23324	0.02445	0.03132	0.00210	0.00005	0.00429	0.00004
	Skewness	-1.42213	0.56219	0.48520	1.04812	2.20825	7.72542	2.38856	4.71498
	Kurtosis	9.65794	2.08146	2.21400	3.39042	11.86176	92.58063	15.02140	33.75686

Note: Before Time: April 10, 2006, to April 30, 2008; After Time: May 01, 2008, to May 27, 2010. Daily total trading volume (*TrdVolu*), daily total open interest (*OpInt*), and implied volatility (*ImpVola*) based on the near-the-money options traded on futures, 200-day rolling historical volatility measure (*HisVola*).

Data source: CME Group; ATS and MSG data is available from May 01, 2008, to May 27, 2010.

the trading volume and open interest for the five contracts during the two years before and after the start of our AT data. While Figures 7 and 8 show no obvious trend, the ratio of trading volume to open interest presented in Figure 9 suggests a positive time trend across all contracts with differing magnitudes. Figures 10 and 11 display the estimates of the implied and the intraday volatility of futures prices. Although the main focus of the paper is not to statistically analyze these factors individually, these graphs help visualize the market conditions specific to the futures contracts under investigation.

We also include in our analysis various variables to control for conditions in the overall financial markets. Table 3 contains the statistics for the market control variables and provides before and after comparisons. Figure 12 graphs select market control variables (VIX, CorpSprd, GSCI, Gold, and S&P 500) over the four years (April 2006 to May 2010).

Using parametric and non-parametric tests for the mean and median of contract specific variables, we investigate potential changes in trading volume, open interest, implied and historical volatility, and four different measures of intraday volatility (Garman-Klass, Parkinson, Range, and RSY94). For all five contracts, we observe an increase in all volatility measures before (April 10, 2006, to April 30, 2008) and after (May 1, 2008, to May 27, 2010) availability of ATS data in our study. Except for the E-mini S&P 500 contract, open interest appears to decrease in the after period.

These descriptive statistics are casual graphical observations and simple univariate comparisons of means and medians. Our intention is not to model the before and after effects based on ATS data availability but rather to use these variables in a microstructure model to control for changes in markets specific to each contract in addition to the overall economy.

IV. EMPIRICAL METHODS


In this section we describe the empirical methods used in estimating the intraday price volatility and the models used in investigating the effects of DMA and algorithmic trading on futures market liquidity. Liquidity measures used are the daily average width and depth provided by the CME and calculated using the intraday quotes and transaction prices.

A. Estimating Intraday Volatility

In addition to the implied and historical volatility measures provided by the Reuters/CRB dataset, we estimate the intraday volatility (*IntVola*) of the futures prices using various methods, expecting that both short-term and long-term volatility affect market liquidity.

Finance literature, in particular futures markets research, contains numerous methods to estimate intraday volatility using the daily open (*OP*), high (*HP*), low (*LP*), and closing (*CP*) prices. The simplest estimator is the difference between the high and the low prices of the day:


Figure 7. Daily Trading Volume for the Period April 10, 2006, to May 27, 2010.



Daily Trading Volume for the Period April 10, 2006, to May 27, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts. Data Source: Reuters/CRB database.


Note: Casual observation of the graph suggests that there is no evident change in the trading volume after May 1, 2008 (start of the microstructure dataset used in this study).

Figure 8. Daily Total Open Interest for the Period April 10, 2006, to May 27, 2010.



Daily Total Open Interest for the Period April 10, 2006, to May 27, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts. Data Source: Reuters/CRB database.
 Note: Casual observation of the graph suggests that there is no evident change in the open interest after May 1, 2008 (start of the microstructure dataset used in this study).


Figure 9. Ratio of Trading Volume to Open Interest for the Period April 10, 2006, to May 27, 2010.



Ratio of Trading Volume to Open Interest for the Period April 10, 2006, to May 27, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts. Data source: Reuters/CRB database.


Note: Casual observation of the graph suggests that ratio of trading volume to open interest is increasing during the time frame under investigation, indicating an improvement in liquidity.

Figure 10. Implied Volatility for the Period April 10, 2006, to May 27, 2010.



Implied Volatility for the Period April 10, 2006, to May 27, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, and 10-year U.S. Treasury Note Contracts. Data source: Reuters/CRB database. Note: Implied volatility (*ImpVola*) for each of the contracts based on the near-the-money options traded on those futures. This graph suggests that there is a marked increase in volatility starting with the third quarter of 2008 (in line with recent turmoil in financial markets). However, this does not appear to be immediately after May 1, 2008 (start of the microstructure dataset used in this study).

Figure 11. Garman-Klass Estimate of Intraday Volatility for the Period April 10, 2006, to May 27, 2010.



Garman-Klass Estimate of Intraday Volatility for the Period April 10, 2006, to May 27, 2010, for Crude Oil, Euro FX, Eurodollar, E-mini S&P 500, 10-year Treasury Note. Intraday volatility is estimated using the Garman-Klass (*GarKla*) estimator. This graph suggests that there is a period of heightened intraday volatility for all contracts from the September 2008 until April 2009. Another observation is that the Garman-Klass estimate of the intraday volatility of Eurodollar contract is very small (order of 10⁻²) compared to the other four contracts. However, this high intraday volatility period does not appear to be around May 01, 2008 (start of the microstructure dataset used in this study). Data Source: Reuters/CRB database.

Table 3. Descriptive Statistics on Market Control Variables: April 10, 2006, to May 27, 2010.

Horizon	Statistic	<u>Corp AAA</u>	<u>Corp BAA</u>	<u>Corp Sprd</u>	<u>Tbill 3mo</u>	<u>Def Sprd</u>	<u>Term Sprd</u>
April 10, 2006 to April 30, 2008	Mean	5.5793	6.5536	0.9743	4.2186	1.0265	0.3373
	Median	5.55	6.56	0.91	4.853	0.803	0.152
	Std. Dev.	0.2108	0.2250	0.1732	1.1955	0.4110	0.7982
	Skewness	0.2511	0.0775	1.4610	-1.4156	1.2110	1.2355
	Kurtosis	2.0095	2.2059	3.9512	3.7450	3.1567	3.7752
	IQ Range	0.3600	0.3500	0.0900	1.2640	0.4940	0.9390
	CV	0.0378	0.0343	0.1778	0.2834	0.4004	2.3664
May 01, 2008 to May 27, 2010	Mean	5.4245	7.2348	1.8102	0.4505	1.9621	3.0086
	Median	5.365	7.075	1.48	0.155	1.7715	3.203
	Std. Dev.	0.3362	0.9504	0.8153	0.6267	0.4125	0.5584
	Skewness	0.9436	0.5829	0.7015	1.5112	0.6760	-0.4200
	Kurtosis	4.6777	2.2126	1.9179	3.5224	1.9874	1.7661
	IQ Range	0.3700	1.7250	1.6300	0.2120	0.7275	0.9730
	CV	0.0620	0.1314	0.4504	1.3912	0.2102	0.1856
April 10, 2006 to May 27, 2010	Mean	5.5021	6.8935	1.3915	2.3291	1.4947	1.6807
	Median	5.475	6.64	1.095	1.787	1.57	2.1105
	Std. Dev.	0.2908	0.7692	0.7220	2.1123	0.6232	1.5029
	Skewness	0.4506	1.5055	1.5710	0.2067	0.2739	-0.1190
	Kurtosis	4.1149	4.5344	4.1597	1.2876	2.0547	1.4215
	IQ Range	0.3700	0.7200	0.5700	4.6980	1.0160	3.0525
	CV	0.0529	0.1116	0.5189	0.9069	0.4169	0.8942

Table 3, continued. Descriptive Statistics on Market Control Variables: April 10, 2006, to May 27, 2010.

Horizon	Statistic	<u>DOW</u>	<u>NASDAQ</u>	<u>NYSE</u>	<u>Russell1000</u>	<u>SP500</u>
April 10, 2006 to April 30, 2008	Mean	4,217.73	2,425.01	9,128.36	762.99	1,401.80
	Median	4,225.97	2,430.86	9,139.57	766.33	1,408.21
	Std. Dev.	259.89	190.07	641.76	48.58	88.34
	Skewness	-0.1904	-0.0457	-0.1427	-0.1047	-0.1018
	Kurtosis	1.9050	2.3005	2.0043	1.9074	1.9160
	IQ Range	421.72	268.17	1,029.14	80.27	148.15
	CV	0.0616	0.0784	0.0703	0.0637	0.0630
May 01, 2008 to May 27, 2010	Mean	3,400.42	2,012.76	6,789.97	574.95	1,051.88
	Median	3,379.32	2,123.93	6,899.68	584.91	1,066.19
	Std. Dev.	533.86	346.05	1,234.00	96.84	173.12
	Skewness	0.2752	-0.3308	0.2761	0.1321	0.1614
	Kurtosis	2.3995	1.7879	2.4248	2.1651	2.2283
	IQ Range	746.30	608.04	1,690.72	153.96	265.63
	CV	0.1570	0.1719	0.1817	0.1684	0.1646
April 10, 2006 to May 27, 2010	Mean	3,807.11	2,217.90	7,953.55	668.52	1,226.00
	Median	3,898.49	2,300.05	8,320.19	696.61	1,277.58
	Std. Dev.	586.33	347.26	1,528.84	121.36	222.63
	Skewness	-0.5618	-0.7831	-0.4731	-0.5187	-0.4874
	Kurtosis	2.2776	2.8530	2.1000	2.1420	2.0934
	IQ Range	911.08	385.46	2,319.66	183.40	344.65
	CV	0.1540	0.1566	0.1922	0.1815	0.1816


Table 3, continued. Descriptive Statistics on Market Control Variables: April 10, 2006, to May 27, 2010.

Horizon	Statistic	<u>Gold</u>	<u>DollarInd</u>	<u>GSCI</u>	<u>ReutersCRB</u>	<u>VIX</u>
April 10, 2006 to April 30, 2008	Mean	706.62	81.41	512.43	424.35	17.14249
	Median	663.53	82.53	483.44	404.98	15.235
	Std. Dev.	110.51	4.4409	83.43	52.04	5.7511
	Skewness	1.0993	-0.6508	1.0789	1.2762	0.6121
	Kurtosis	2.9892	2.3270	3.2554	3.6471	2.1157
	IQ					
	Range	138.14	7.4610	121.55	61.41	10.2900
	CV	0.1564	0.0546	0.1628	0.1226	0.3355
May 01, 2008 to May 27, 2010	Mean	964.55	79.87	510.85	442.37	31.23839
	Median	938.36	79.85	486.96	450.95	25.45
	Std. Dev.	124.10	4.4334	144.36	67.00	13.6837
	Skewness	0.2608	0.0943	0.9555	0.3701	1.2994
	Kurtosis	2.2023	2.0955	3.1350	2.3668	4.0689
	IQ					
	Range	206.22	7.2350	130.40	98.05	18.1800
	CV	0.1287	0.0555	0.2826	0.1515	0.4380
April 10, 2006 to May 27, 2010	Mean	835.71	80.63	511.63	433.40	24.22429
	Median	848.48	80.87	484.11	413.41	21.68
	Std. Dev.	174.48	4.5017	117.99	60.67	12.6548
	Skewness	0.3015	-0.2644	1.0575	0.7580	1.6549
	Kurtosis	2.0330	2.0103	3.9091	2.7847	5.9705
	IQ					
	Range	281.60	7.8910	104.27	86.22	11.9300
	CV	0.2088	0.0558	0.2306	0.1400	0.5224

Notes: *CorpAAA*—AAA-corporate bond yield; *CorpBAA*—BAA corporate bond yield; *CorpSprd*—corporate credit spread (= *CorpBAA* – *CorpAAA*); *Tbill3mo*—yield on 3-month Treasury Bill; *DefSprd*—difference between the AAA-corporate bond yield and the yield on 10-year Treasury Note; *TermSprd*—difference between the yields on 10-year Treasury Note and the 3-month Treasury Bill; *DOW*—daily stock index levels for Dow Jones Industrial Average; *NASDAQ*—NASDAQ composite; *NYSE*—New York Stock Exchange Composite; *Russell1000*—Russell 1000; *SP500*—S&P 500; *GSCI*—daily values of Goldman Sachs Commodity Index; *DollarInd*—U.S. Dollar Index; *GOLD*—spot Gold price; *ReutersCRB*—Reuters/CRB Commodity Index; and *VIX*—the CBOE's Volatility Index.

Data sources: Reuters/CRB database; CME Group's ATS and MSG data is available from May 1, 2008, to May 27, 2010.

Figure 12. Market Control Variables, VIX, CorpSprd, GSCI, Gold, and SP500, April 10, 2006, to May 27, 2010.



Note: *VIX* is the CBOE's Volatility Index, *CorpSprd* is the corporate credit spread (calculated as the difference between the AAA rated and BAA rated corporate bond yields), *GOLD*, *GSCI*, and *SP500* are the daily spot levels of Gold, Goldman Sachs Commodity Index, and S&P 500 Index, respectively. Data source: Reuters/CRB database.

$$Range_t = Ln(HP_t) - Ln(LP_t) \quad (1)$$

Some researchers also used the simple difference of the two prices (Chan and Lien 2003). Parkinson (1980) proposes a revised version of the range estimator:

$$Parkinson_t = [Ln(HP_t) - Ln(LP_t)]^2 / [4Ln(2)] \quad (2)$$

Garman and Klass (1980) incorporate the opening and low prices of the day into the following estimate of intraday volatility:²¹

$$GarKla_t = \left\{ \frac{1}{2} [Ln(HP_t) - Ln(LP_t)]^2 \right\} - \left\{ [2\ln(2) - 1] [Ln(CP_t) - Ln(OP_t)]^2 \right\} \quad (3)$$

A version of the Garman-Klass estimator independent of the drift is proposed by Rogers, Satchell, and Yoon (1994):²²

$$RSY94_t = \left\{ [Ln(HP_t) - Ln(OP_t)] [Ln(HP_t) - Ln(CP_t)] \right\} - \left\{ [Ln(LP_t) - Ln(OP_t)] [Ln(LP_t) - Ln(CP_t)] \right\} \quad (4)$$

All four of these intraday volatility estimators rely on the daily range based analysis with varying levels of efficiency. Based on the futures markets research, we use the Garman-Klass estimates of intraday volatility in our empirical analysis. We also repeat empirical tests using other estimators and find that our results do not materially change.

B. Modeling Liquidity and AT

In order to investigate the effects of DMA and AT on the liquidity of futures contracts traded at the CME, we use a model similar to the one used by Hendershott, Jones, and Menkveld (2011). They model the relationship between the liquidity and their proxy of algorithmic trading as:

$$Liq_{i,t} = \alpha_i + \beta AT_{i,t} + \delta' X_{i,t} + \varepsilon_{i,t} \quad (5)$$

where $Liq_{i,t}$ is a measure of liquidity for stock i on day t , $AT_{i,t}$ is their proxy for the algorithmic trading, and $X_{i,t}$ is a vector of control variables (which they choose to be share turnover, volatility, the inverse of share price, and log market cap).²³ They

21. Chen, Daigler, and Parhizgari (2006) and Shu and Zhang (2006) illustrate that volatility estimates using the Garman-Klass method and the high frequency realized volatility measures provide equivalent results.

22. Yang and Zhang (2000) discuss modifications to the RSY94 estimator.

23. Hendershott et al. (2011) include both fixed effects and time dummies in their model.

estimate the panel regressions in equation (5) using standard errors that are robust to general cross-section and time-series heteroskedasticity and within-group autocorrelation (Arellano and Bond 1991).

Our empirical tests use two different direct measures of algorithmic trading provided by the CME: *ATS*, percentage of trading volume identified as originating from algorithms, and *MSG*, percentage of message traffic identified as originating from algorithms. Our empirical tests do not suffer as much from the measurement error as Hendershott, Jones, and Menkveld's (2011) proxy for AT, normalized measure of electronic message traffic.²⁴ We also use two measures of liquidity, average market width and depth, for each contract. Our control variables include those specific to the contracts *GSCI*, gold price, and CBOE's volatility index *VIX*: estimates of intraday and implied volatility, trading volume and open interest, as well as market-related factors.

We estimate the following general model using various cross-sectional time series (CSTS) techniques:

$$Liq_{i,t} = \alpha_i + \beta_i \text{Algo}_{i,t} + \delta_i' \mathbf{X}_{i,t} + \varphi_i' \mathbf{Z}_{i,t} + \varepsilon_{i,t} \quad (6)$$

where $Liq_{i,t}$ is either of our liquidity measures *ATS* or *MSG*; $\text{Algo}_{i,t}$ is either of our direct measure of algorithmic trading, $\mathbf{X}_{i,t}$ is a vector of control variables on each futures contract (*IntVola*, intraday volatility; *ImpVola*, implied volatility; *OpInt*, open interest; *TrdVola*, trading volume) and $\mathbf{Z}_{i,t}$ is a vector of market controls (*GSCI*, Goldman Sachs Commodity Index; *Gold*, price of gold; *VIX*, CBOE's volatility index). Explicitly, we first estimate models without market controls:

$$Liq_{i,t} = \alpha_i + \beta_i A_{i,t} + \delta_{1,i} \text{IntVola}_{i,t} + \delta_{2,i} \text{ImpVola}_{i,t} + \delta_{3,i} \text{OpInt}_{i,t} + \delta_{4,i} \text{TrdVola}_{i,t} + \varepsilon_{i,t} \quad (7)$$

where

$$Liq_{i,t} = \begin{cases} \text{Width}_{i,t} \\ \text{Depth}_{i,t} \end{cases}, \text{ and } A_{i,t} = \begin{cases} \text{ATS}_{i,t} \\ \text{MSG}_{i,t} \end{cases}. \quad (8)$$

In order to provide robust estimation results, we use the following alternative panel estimation methods: (a) random-effects GLS regressions with autoregressive errors AR(1); (b) standard fixed-effects panel regression using the between-regression estimator (when we exclude market controls from the independent variables). When we include the vector of market controls in our analysis, we estimate the following models using (c) standard fixed-effects panel regression with using the between regression estimator and (d) fixed-effects cross-sectional time-series regression with first-order autoregressive disturbances:

24. Hendershott et al. (2011) state that they "cannot directly observe whether a particular order is generated by a computer algorithm," which is due to the nature of the NYSE data they use in their analysis. They indicate that "the rate of electronic message traffic may be a useful proxy for the amount of algorithmic trading taking place," which they normalize by dividing number of electronic messages by trading volume of each stock on a given day.

$$Liq_{i,t} = \alpha_i + \beta_i A_{i,t} + \delta_{1,i} \text{IntVola}_{i,t} + \delta_{2,i} \text{ImpVola}_{i,t} + \delta_{3,i} \text{OpInt}_{i,t} + \delta_{4,i} \text{TrdVol}_{i,t} + \varphi_1 \text{GSCI}_t + \varphi_2 \text{Gold}_t + \varphi_3 \text{VIX}_t + \varepsilon_{i,t} \quad (9)$$

$$\text{where } Liq_{i,t} = \begin{cases} \text{Width}_{i,t} \\ \text{Depth}_{i,t} \end{cases}, \text{ and } A_{i,t} = \begin{cases} \text{ATS}_{i,t} \\ \text{MSG}_{i,t} \end{cases}. \quad (10)$$

We estimate equation (6) with various market control variables and find that the results do not materially change; therefore, we report our findings using the vector of market controls that include the GSCI, Gold, and the VIX.

V. EMPIRICAL RESULTS

Table 4 presents the empirical results for the effects of algorithmic trading on liquidity using only the contract specific factors as control variables (specifically equations 7 and 8). The results using both the random-effects GLS regressions with AR(1) and the fixed-effects models are consistent. After controlling for intraday and implied volatilities, trading volume and open interest, we find that an increase in the proportion of trading associated with algorithmic trading systems (*ATS*) decreases the width (spreads) and increases the market depth. When an AT's proportion of electronic message traffic (*MSG*) is used as a measure of algorithmic trading, we observe the same results. Our models explain relatively large portions of within and between variation in the cross-sectional time series data, and coefficient estimates of *ATS* and *MSG* are all significant at 1%.

Estimated coefficients of volatility, volume, and open interest are consistent with the findings in futures MMR. (See, e.g., Wang, Yau, and Baptiste 1997; Wang and Yao 2000; Girma and Mougoue 2002; Bryant and Haigh 2004; and Frank and Garcia 2009.) *Width* (spreads) increases with both measures of volatility and decreases with trading volume and open interest; their effect on *Depth* is reversed. Our results for the volatility are robust to the measurement of short-term (intraday) volatility and longer-term (implied) volatility.

The changes we observe by considering only the futures contract-specific factors may in fact be influenced by other dynamics of overall financial markets. Table 5 presents findings when we include both futures contract and market control variables in our cross-sectional time series regressions (specifically equations 9 and 10). Results based on cross-sectional time series estimation using both the fixed-effects and fixed-effects with AR(1) disturbances are consistent and confirm the findings presented in Table 4.

We again observe that trading volume of *ATS* (as well as their proportion of electronic message traffic, *MSG*) decreases the *Width* while increasing the market *Depth*, after controlling for both futures contract-specific and market-wide factors. While the coefficient estimates of futures contract-specific control factors retain their signs and significance, the inclusion of market-wide factors increases the

Table 4. Effects of Algorithmic Trading on Liquidity, May 1, 2008, to May 27, 2010: Cross-Sectional Time Series (Panel) Data Analysis for $Liq_{i,t} = \alpha_i + \beta_i Algo_{i,t} + \delta_i X_{i,t} + \varepsilon_{i,t}$

	Random-effects GLS Regression with AR(1)		CSTC with Fixed-effects	
	Width	Depth	Width	Depth
<i>ATS</i>	-63.07 (-15.22)*	3346.60 (14.23)*	-105.99 (-20.57)*	7793.50 (22.99)*
<i>MSG</i>	-23.7574 (-5.29)*	1601.95 (7.44)*	-42.93 (-9.71)*	3631.73 (12.13)*
<i>IntVola(GarKla)</i>	437.20 (13.21)*	-737.35 (-0.51)	671.62 (17.74)*	-6535.28 (-2.44)**
<i>ImpVola</i>	12.88 (8.96)*	383.50 (4.69)*	14.1302 (9.4)*	637.61 (5.74)*
<i>OpInt</i>	-9.84E-07 (-0.82)	2.02E-04 (13.19)*	-2.6874 (-1.73)**	1.73E-04 (5.89)*
<i>TrdVolu</i>	-3.20E-06 (-6.94)*	-6.39E-05 (-3.33)*	-8.9189 (-11.84)*	1.68E-05 (0.49)
<i>Constant</i>	57.00 (22.75)*	-1783.76 (-12.11)*	227.94 (11.46)*	-4003.83 (-20.87)*

Table 4, continued. Effects of Algorithmic Trading on Liquidity, May 1, 2008, to May 27, 2010: Cross-Sectional Time Series (Panel)
 Data Analysis for $Liq_{i,t} = \alpha_i + \beta_i \text{Algo}_{i,t} + \delta_i X_{i,t} + \varepsilon_{i,t}$

	Random-effects GLS Regression with AR(1)		CSTC with Fixed-effects	
	Width	Depth	Width	Depth
R2 within	0.3577	0.2667	0.4276	0.2223
R2 between	0.7803	0.5991	0.9013	0.1308
R2 overall	0.4734	0.4292	0.5617	0.1203
# of obs.	2220	2220	2210	2267
Wald Chi ² (6)	805.12	510.62	328.75	129.05

*, **, and *** denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Note: Two direct measures of algorithmic trading ($\text{Algo}_{i,t}$) are $\text{ATS}_{i,t}$ – percentage of volume attributed to automated trading systems and $\text{MSG}_{i,t}$ – percent of message traffic attributed to automated trading systems. Two measures of the liquidity ($\text{Liq}_{i,t}$) are $\text{Width}_{i,t}$ – average bid–ask spread for a given size order during a trading day, and $\text{Depth}_{i,t}$ – number of contracts displayed at the “top-of-the-book” (i.e., average contract size of the best bid and best ask quotes). $X_{i,t}$ is a vector of control variables on each futures contract; TrdVolu is daily total trading volume, OpInt is daily total open interest, $\text{ImpVolu}(\text{GarKla})$ is the Garman-Klass estimate of intraday volatility, and ImpVolu is implied volatility for each of the contracts based on the near-the-money options traded on those futures. The data for the ATS , MSG , Width and Depth variables are from regular trading hours.

Table 5. Effects of Algorithmic Trading on Liquidity, Controlling for Market Factors, May 1, 2008, to May 27, 2010: Cross-Sectional Time Series (CSTS) Data Analysis for $Liq_{i,t} = \alpha_i + \beta_1 Algo_{i,t} + \delta_i X_{i,t} + \varphi_i Z_{i,t} + \varepsilon_{i,t}$

	CSTS with Fixed-effects		CSTC with Fixed-effect & AR(1)	
	Width	Depth	Width	Depth
<i>ATS</i>	-56.33 (-13.01)*	5570.94 (15.61)*	-32.3193 (-8.57)*	1616.71 (7.66)*
<i>MSG</i>		2826.60 (10.35)*	0.0549 (0.02)	326.01 (1.86)**
<i>IntVola(GarKla)</i>	13.248 (3.99)*	18219.66 (6.65)*	82.08 (2.87)*	4492.32 (3.2)*
<i>ImpVola</i>	13.10 (10.4)*	651.74 (6.23)*	12.28 (9.05)*	563.24 (6.8)*
<i>OpInt</i>	-11.27 (-8.41)*	211.34 (1.91)**	-6.12E-07 (-2.28)**	1.55E-04 (10.41)*
<i>TrdVola</i>	-4.64 (-7.71)*	-168.45 (-3.68)*	-2.82E-06 (-7.18)*	-6.25E-05 (9.92)*
<i>VIX</i>	0.6435 (25.99)*	-14.30 (-7.08)*	0.7785 (37.65)*	-23.86 (-18.2)*
<i>GSCI</i>	0.0362 (19.06)*	0.1630 (1.05)	0.0387 (22.35)*	-0.9603 (-8.04)*
<i>GOLD</i>	-0.0123 (-5.63)*	1.4233 (7.86)*	-0.0083 (-4.71)*	0.4533 (4.11)*
<i>Constant</i>	246.23 (14.63)*	-4379.05 (-3.15)*	7.5402 (14.94)*	-169.85 (-6.29)*

Table 5, continued. Effects of Algorithmic Trading on Liquidity, Controlling for Market Factors, May 1, 2008, to May 27, 2010: Cross-Sectional Time Series (CSTS) Data Analysis for $Liq_{i,t} = \alpha_i + \beta_i \text{Algo}_{i,t} + \delta_i X_{i,t} + \varphi_i Z_{i,t} + \varepsilon_{i,t}$

	CSTS with Fixed-effects		CSTC with Fixed-effect & AR(1)	
	Width	Depth	Width	Depth
R2 within	0.6562	0.6376	0.589	0.2187
R2 between	0.4978	0.145	0.7664	0.8302
R2 overall	0.5539	0.3468	0.643	0.5475
# of obs.	2210	2210	2212	2259
F(8,2197)(8,2244)	524.2	483.17	393.99	78.57

*, **, and *** denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Note: Two direct measures of algorithmic trading ($\text{Algo}_{i,t}$) are $\text{ATS}_{i,t}$ – percentage of volume attributed to automated trading systems and $\text{MSG}_{i,t}$ – percentage of message traffic attributed to automated trading systems. Two measures of the liquidity ($\text{Liq}_{i,t}$) are $\text{Width}_{i,t}$ – average bid–ask spread for a given size order during a trading day, and $\text{Depth}_{i,t}$ – number of contracts displayed at the “top-of-the-book” (i.e., average size-in terms of contracts-of the best bid and best ask quotes). $X_{i,t}$ is a vector of control variables on each futures contract where TrdVolu is daily total trading volume, OpInt is daily total open interest, $\text{IntVol}(\text{GarKla})$ is the Garman-Klass estimate of intraday volatility, and ImpVol is implied volatility for each of the contracts based on the near-the-money options traded on those futures. $Z_{i,t}$ is a vector of market controls where GSCI is the Goldman Sachs Commodity Index, GOLD is spot Gold price, and VIX is the CBOE’s Volatility Index. The data for the ATS , MSG , Width and Depth variables are from regular trading hours.

within and between R-squared values of our models.²⁵

Our empirical results for the effects of AT on the liquidity in futures markets using direct measures that identify algorithm-generated trades and quote revisions confirm the findings for the U.S. equity markets by Hendershott, Jones, and Menkveld (2011) and the findings for the German equity markets by Hendershott and Riordan (2009). While we employ a very similar model to the one used by Hendershott, Jones, and Menkveld, our measures of AT activity do not suffer from their measurement errors. Results presented in our Tables 4 and 5 are based on four different cross-sectional time series modeling techniques and two separate direct measures of AT activity; after controlling volatility, trading volume, open interest and other market-wide factors, the findings indicate that algorithmic trading has a significant positive impact on market liquidity. This is evidenced by a decrease in spreads and an increase in depth. The nature of our dataset obtained from the CME Group precludes us from analyzing the informativeness of individual AT generated trades and quotes.

VI. CONCLUSIONS

Although the extensive use of algorithmic trading (AT) activities emerged relatively more recently in the exchange-traded derivatives in comparison to the equity markets, their impact on market quality and risk management may be more substantial. In order to analyze the potential effects of DMA, AT, and their accompanied changes in exchange-traded derivatives markets, this study provides an extensive review of the research in both equity and derivatives market microstructure.

After synthesizing the very recent and limited empirical evidence for the effects of algorithmic trading in equity markets, our research presents empirical results based on a unique dataset of algorithmic trading activity in five futures contracts electronically traded at the CME Group exchanges. To the best of our knowledge, this study is the first to provide such empirical evidence for the U.S. futures markets.

The uniqueness of the dataset used in this study is due to the explicit identification (direct measurement) of algorithmic trading volume — the proportion of executed orders originated from ATS to the total electronic orders executed (*variable ATS*). CME Group data also include the proportional volume of electronic message traffic attributed to ATS (*variable MSG*). Our empirical results are based on the Crude Oil, Euro FX, Eurodollar, S&P 500 E-mini, and 10-year U.S. Treasury Note futures, for the time period between May 1, 2008, and May 27, 2010.

After controlling for short- and longer-term volatility, trading volume, and open interest, as well as other market-wide factors, we find that an increase in the proportion of trading associated with algorithmic trading systems (*ATS*) decreases the width (spreads) and increases the market depth in futures trading. When an AT's proportion of electronic message traffic (*MSG*) is used as a measure of

25. We estimate equations (9) and (10) using various combinations of market control variables and find no material change in our overall results for the impact of AT on liquidity.

algorithmic trading activity, we observe similar statistically significant results. Our models explain relatively large portions of within and between variations in the cross-sectional time series data, and our coefficient estimates for the volatility, volume, and open interest all have the expected signs and significance. Similar to recent research in equity markets, our results for the U.S. futures markets conclude that algorithmic trading has a positive impact on market liquidity.

It is our intent that this paper will provide guidance to market participants, exchanges, and regulators because it presents empirical evidence on early stages of DMA and AT in futures markets and discusses the implications of these developments for exchange-traded derivatives markets.

References

- Aitken, M. J., Harris, F. H. deB., and Ji, S., 2009, Trade-Based Manipulation and Market Efficiency: A Cross-Market Comparison. *Working paper, 22nd Australasian Finance and Banking Conference 2009*, November 18.
- Anand, A. and Chakravarty, S., 2007, Stealth Trading in Options Markets. *Journal of Financial and Quantitative Analysis*, 42(1), 167-187.
- Anand, A. and Subrahmanyam, A., 2008, Information and the Intermediary: Are Market Intermediaries Informed Traders in Electronic Markets? *Journal of Financial and Quantitative Analysis*, 43(1), 1-28.
- Anand, A. and Weaver, D., 2004, Can Order Exposure Be Mandated? *Journal of Financial Markets*, 7, 405-426.
- Arellano, M. and Bond, S., 1991, Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. *The Review of Economic Studies*, 58, 277-297.
- Ates, A. and Wang, G.H.K., 2005, Information Transmission in Electronic versus Open-Outcry Trading Systems: An Analysis of U.S. Equity Index Futures Markets. *Journal of Futures Markets*, 25, 679-715.
- Avgouleas, E. and Degiannakis, S., 2005, The Impact of the EC Financial Instruments Markets Directive on the Trading Volume of EU Equity Markets. *Athens University of Economics and Business, Statistics Technical Report No. 219*. Available at SSRN: <http://ssrn.com/abstract=828804>
- Barclay, M., Hendershott, T. and McCormick, D.T., 2003, Competition Among Trading Venues: Information and Trading on Electronic Communications Networks. *Journal of Finance*, 58(6), 2637-2666.
- Bessembinder, H., Maxwell, W., and Venkataraman, K., 2006, Market Transparency, Liquidity Externalities, and Institutional Trading Costs in Corporate Bonds. *Journal of Financial Economics*, 82, 251-288.
- Bloomfield, R. and O'Hara, M., 2000, Can Transparent Markets Survive? *Journal of Financial Economics*, 55(3), 425-459.
- Bortoli, L., Frino, A., Jarnecic, E., and Johnstone, D., 2006, Limit Order Book Transparency, Execution Risk, and Market Liquidity: Evidence from the Sydney Futures Exchange. *Journal of Futures Markets*, 26(12), 1147-1167.

- Brogaard, J. A., 2010, High Frequency Trading and Its Impact on Market Quality. *Working Paper, Northwestern University*.
- Brownlees, C.T., Cipollini, F., and Gallo, G.M., 2010, Intra-daily Volume Modeling and Prediction for Algorithmic Trading. *Journal of Financial Econometrics*, forthcoming.
- Bryant, H. and Haigh, M., 2004. Bid-Ask Spreads in Commodity Futures Markets. *Applied Financial Economics*, 14, 923-936.
- Cabrera, J., Wang, T., and Yang, J., 2009, Do Futures Lead Price Discovery in Electronic Foreign Exchange Markets? *Journal of Futures Markets*, 29(2), 137-156.
- Castura, J., Litzenberger, R., and Gorelick, R., 2010, Market Efficiency and Microstructure Evolution in U.S. Equity Markets: A High-Frequency Perspective. *Technological Trading in the Markets*, U.S. Commodity Futures Trading Commission, July 14.
- Chaboud, A., Chiquoine, B., Hjalmarsson, E., and Vega, C., 2009, Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. *Working paper, No. 980, Federal Reserve Board*.
- Chakravarty, S. and Li, K., 2003, An Examination of Own Account Trading by Dual Traders in Futures Markets. *Journal of Financial Economics*, 69(2), 375-397.
- Chan, L. and Lien, D., 2003, Using High, Low, Open, and Closing Prices to Estimate the Effects of Cash Settlement on Futures Prices. *International Review of Financial Analysis*, 12, 35-47.
- Chen, Z., Daigler, R.T., and Parhizgari, A.M., 2006, Persistence of Volatility in Futures Markets. *Journal of Futures Markets*, 26(6), 571-594.
- Chordia, T., Roll, R., and Subrahmanyam, A., 2008. Why Has Trading Volume Increased? *Technical report, University of California at Los Angeles*.
- Chung, H. and Chiang, S., 2006, Price Clustering in E-Mini and Floor-Traded Index Futures. *Journal of Futures Markets*, 26(3), 269-295.
- Comerton-Forde, C. and Rydge, J., 2006, The Current State of Asia-Pacific Stock Exchanges: A Critical Review of Market Design. *Pacific-Basin Finance Journal*, 14, 1-32.
- Comerton-Forde, C. and Tang, K. M., 2009, Anonymity, Liquidity and Fragmentation. *Journal of Financial Markets*, 12(3), 337-367.
- Conrad, J., Johnson, K.M., and Wahal, S., 2003, Institutional Trading and Alternative Trading Systems. *Journal of Financial Economics*, 70(1), 99-134.
- Cumming, D. and Johan, S., 2008, Global Market Surveillance. *American Law and Economics Review*, 10(2), 454-506.
- De Winne, R. and D'hondt, C., 2007, Hide-and-Seek in the Market: Placing and Detecting Hidden Orders. *Review of Finance*, 11(4), 663-692.
- Easley, D., Hendershott, T., and Ramadorai, T., 2009, Leveling the Trading Field. *Working paper*, accessed at <http://ssrn.com/abstract=961041>.
- Eleswarapu, V. and Venkataraman, K., 2006, The Impact of Political and Legal Institutions on Equity Trading Costs: A Cross-Country Analysis. *Review of Financial Studies*, 19, 1081-1111.

- Frank, J. and Garcia, P., 2009, Time-Varying Risk Premium? Further Evidence in Agricultural Futures Markets. *Applied Economics*, 41(6), 715-725.
- Frino, A., Bjursell, J., Wang, G.H.K., and Lepone, A., 2008, Large Trades and Intraday Futures Price Behavior. *Journal of Futures Markets*, 28(12), 1147-1181.
- Frino, A., Johnstone, D., and Zheng, H., 2010, Anonymity, Stealth Trading, and the Information Content of Broker Identity. *Financial Review*, 45(3), 501-522.
- Frino, A., Lepone, A., and Wearin, G., 2008, Intraday Behavior of Market Depth in a Competitive Dealer Market: A Note. *Journal of Futures Markets*, 28(3), 294-307.
- Garman, M.B. and Klass, M.J., 1980, On the Estimation of Security Price Volatilities from Historical Data. *Journal of Business*, 53(1), 67-78.
- Garvey, R. and Wu, F., 2010, Speed, Distance, and Electronic Trading: New Evidence on Why Location Matters. *Journal of Financial Markets*, 13, 367-396.
- Gerig, A. and Michayluk, D., 2010, Automated Liquidity Provision and the Demise of Traditional Market Making. *Working paper, University of Technology, Sydney*.
- Girma, P. B. and Mougoue, M., 2002, An Empirical Examination of the Relation between Futures Volatility, Volume, and Open Interest. *Journal of Futures Markets*, 22(11), 1083-1102.
- Grammig, J., Schiereck, D., and Theissen, E., 2001, Knowing Me, Knowing You: Trader Anonymity and Informed Trading in Parallel Markets. *Journal of Financial Markets*, 4, 385-412.
- Hasbrouck, J. and Saar, G., 2009, Technology and Liquidity Provision: The Blurring of Traditional Definitions. *Journal of Financial Markets*, 12(2), 143-172.
- Hendershott, T. and Jones, C., 2005, Island Goes Dark: Transparency, Fragmentation, and Regulation. *Review of Financial Studies*, 18(3), 743-793.
- Hendershott, T., Jones, C.M., and Menkveld, A.J., 2011, Does Algorithmic Trading Improve Liquidity? *Journal of Finance*, 66(1), 1-33.
- Hendershott, T. and Riordan, R., 2009, Algorithmic Trading and Information. *Working paper, University of California at Berkeley and Karlsruhe Institute of Technology*.
- Huang, Y. C., 2004, The Components of Bid-Ask Spread And Their Determinants: TAIFEX Versus SGX-DT. *Journal of Futures Markets*, 24(9), 835-860.
- Pagano, M.S. and Schwartz, R.A., 2003, A Closing Call's Impact on Market Quality at Euronext Paris. *Journal of Financial Economics*, 68(3), 439-484.
- Pardo, A. and Pascual, R., 2007, On the Hidden Side of Liquidity, *Working Paper, University of Valencia and Universidad de las Islas Baleares*, accessed at <http://ssrn.com/abstract=459000>.
- Parkinson, M., 1980, The Extreme Value Method for Estimating the Variance of the Rate of Return. *Journal of Business*, 53(1), 61-65.
- Poskitt, R., 2010, Price Discovery in Electronic Foreign Exchange Markets: The Sterling/Dollar Market. *Journal of Futures Markets*, 30(6), 590-606.
- Reiss, P. C. and Werner, I. M., 2005, Anonymity, Adverse Selection, and the Sorting of Interdealer Trades. *Review of Financial Studies*, 18(2), 599-636.

- Riordan, R. and Storckenmaier, A., 2009, Exchange System Innovation and Algorithmic Liquidity Supply. *Working paper, Karlsruhe Institute of Technology.*
- Rogers, L.C.G., Satchell, S. E., and Yoon, Y. 1994, Estimating the Volatility of Stock Prices: A Comparison of Methods that Use High and Low Prices. *Applied Financial Economics*, 4(3), 241-247.
- Saar, G., 2001, Investor Uncertainty and Order Flow Information. *Working paper, New York University.*
- Securities and Exchange Commission (SEC) Release No. 34 61358, 2010, Concept Release on Equity Market Structure; Proposed Rule. *Federal Register*, 75(13), 3594-3614.
- Shah, S. and Brorsen, B.W., 2010, Transition to Electronic Trading of Kansas City Board of Trade Wheat Futures. *Working paper, Southern Agricultural Economics Association 2010 Annual Meeting.*
- Shu, J. and Zhang, J.E., 2006, Testing Range Estimators of Historical Volatility. *Journal of Futures Markets*, 26(3), 297-313.
- Theissen, E., 2002, Price Discovery in Floor and Screen Trading Systems. *Journal of Empirical Finance*, 9(4), 455-474.
- Tse, Y., Xiang, J., and Fung, J.K.W., 2006, Price Discovery in the Foreign Exchange Futures Market. *Journal of Futures Markets*, 26(11), 1131-1143.
- Tse, Y. and Zobotina, T. V., 2001, Transaction Costs and Market Quality: Open Outcry Versus Electronic Trading. *Journal of Futures Markets*, 21(8), 713-735.
- Tuttle, L. A., 2003, Hidden Orders, Trading Costs and Information. University of Kansas Working Paper. SSRN: <http://ssrn.com/abstract=676019>
- Wagener, M. and Riordan, R., 2009, System Latency in Linked Spot and Futures Markets. Proceedings of the Fifteenth Americas Conference on Information Systems, pp. 231-235.
- Wang, G.H.K., Yau, J., and Baptiste, T., 1997, Trading Volume and Transaction Costs in Futures Markets. *Journal of Futures Markets*, 17(7), 757-780.
- Wang, G.H.K. and Yao, J., 2000, Trading Volume, Bid-Ask Spread, and Price Volatility in Futures Markets. *Journal of Futures Markets*, 20(10), 943-970.
- Webb, R. I., Muthuswamy, J., and Segara, R., 2007, Market Microstructure Effects on Volatility at the TAIFEX. *Journal of Futures Markets*, 27(12), 1219-1243.
- Yang, D. and Zhang, Q., 2000, Drift-Independent Volatility Estimation Based on High, Low, Open and Close Prices. *Journal of Business*, 73, 477-491.