

FUTURES MARKET VOLATILITY: WHAT HAS CHANGED?

NICOLAS P.B. BOLLEN and ROBERT E. WHALEY*

The evolution of trading practices in futures markets, including growth of high-frequency trading, has raised concerns about market quality. This study investigates whether excess futures return volatility, as an encompassing gauge of market quality, has changed over time. Daily measures of realized volatility are computed using 5-minute returns of 15 electronically traded futures contracts. Two benchmarks are used to control for changes in the rate of information flow: option implied volatility and long horizon volatility estimates. Relative to the benchmarks, realized volatility has not changed, indicating that changes in trading practices have not led to a deterioration of market quality. © 2014 Wiley Periodicals, Inc. *Jrl Fut Mark* 35:426–454, 2015

1. INTRODUCTION

Sparked by increased competition and advances in technology, futures markets have undertaken a number of structural changes over the past decade. One prominent example is the rapid growth of high-frequency trading, which accounted for 61% of futures market volume in 2013Q1 compared to 47% in 2008.¹ Some market observers have voiced concern regarding the impact these changes have had on market quality, and claim that futures market volatility has increased as a result.² This paper studies whether the statistical properties of realized return volatility of 15 electronically traded futures contracts, measured primarily using intraday 5-minute returns, have in fact changed in recent years.

Modeling the impact of changes in market microstructure on the volatility of futures returns is no straightforward task. Observed or realized volatility can certainly be affected by market microstructure, including bid/ask spreads, electronic versus pit trading, and the rise in algorithmic trading. However, futures prices also respond rapidly to new information; hence, changes in the rate of information flow, such as the increase that occurred during the financial crisis, also have a direct effect on volatility. As a consequence, it is important to disentangle the microstructural component of realized volatility from changes in fundamental volatility so that statements can be made about market quality.

Nicolas P.B. Bollen is the E. Bronson Ingram Professor of Finance, Owen Graduate School of Management, Vanderbilt University, Nashville, Tennessee. Robert E. Whaley is the Valere Blair Potter Professor of Finance and Director, Financial Markets Research Center, Owen Graduate School of Management, Vanderbilt University, Nashville, Tennessee.

*Correspondence author, Valere Blair Potter Professor of Finance and Director, Financial Markets Research Center, Owen Graduate School of Management, Vanderbilt University, 401 21st Avenue South, Nashville, Tennessee 37203. Tel: +615-343-7747, Fax: 615-343-7177, e-mail: whaley@vanderbilt.edu

Received September 2013; Accepted February 2014

¹See “High-Speed Traders Exploit Loophole,” *Wall Street Journal*, 1 May 2013.

²The CFTC solicited public comments in 2013 on a Concept Release regarding the need to expand testing and supervision of high-speed trading strategies, including limits on the speed of order submissions.

In this study, two benchmarks for fundamental volatility are used to control for changes in the rate of information flow. The benchmarks allow an assessment of the contribution of changes in market microstructure to changes in realized volatility. The first benchmark is the implied volatility in equity index options markets. The level of the CBOE Market Volatility Index or “VIX,” for example, is the market’s best assessment of the expected return volatility of the S&P 500 index over the next 30 days. An option’s implied volatility generally differs from the underlying asset’s realized volatility due to microstructural effects in both markets.³ If changes in microstructure have increased realized volatility, then the difference between implied and realized volatility should change. However, there is no evidence that the difference between the realized volatility for the CME Group’s E-mini S&P 500 futures contract and the VIX has changed over time. The same is true for the relation between the FTSE 100 Volatility Index or “VFTSE,” the implied return volatility of the FTSE 100 index, and the realized volatility of the corresponding NYSE Liffe FTSE 100 index futures contract. Realized volatility for Eurex’s DAX futures contract, in contrast, has risen relative to the corresponding DAX Volatility Index or “VDAX.” This change has occurred only in the last year of the sample, however, and is likely due to recent increases in macroeconomic uncertainty related to the Euro zone crisis.

The second benchmark is return volatility measured over an extended horizon. To understand how this approach works, assume that futures prices are noisy due to microstructural effects such as bid/ask price bounce, price discreteness, and price impact. As argued by Bandi and Russell (2006, 2008), the amount of noise in the futures price is independent of whether you measure returns over 5 minutes or 10 days. Consequently, the “signal-to-noise ratio” (i.e., amount of true information about price change that you are extracting from the data relative to the amount of microstructural noise) is much greater for longer distancing intervals than for shorter ones. Consistent with Bandi and Russell (2006, 2008), we find that realized volatility for shorter holding periods is higher than realized volatility computed over longer periods, thereby confirming the presence of microstructural effects. But, more importantly, the relative magnitudes have not changed meaningfully through time.

Taken together, these two results indicate that, after controlling for changes in the rate of information flow, there is no evidence to suggest that the realized volatility of returns in electronically traded futures markets has changed through time. This conclusion may appear to conflict with discrete disruptions in markets that have been linked to high-frequency trading, including the flash crash of May 2010, as well as several recent academic studies including Lee (2015) and Jarrow and Protter (2012). Note, however, that our measure of realized volatility is based on squared 5-minute returns; hence, we estimate volatility at the daily and monthly frequencies, which are the frequencies that matter for longer-term investors.

2. LITERATURE REVIEW

According to the Futures Industry Association, the number of futures and options contracts traded worldwide tripled between 2003 and 2011.⁴ In addition, as we show in Section 5, volatility for some futures contracts has doubled in recent years. There are two competing explanations. The first is that an increase in the rate of information flow leads to increases in both volatility and trading volume. The second is that changes in trading practices have led to abnormal increases in trading volume, which have added noise to prices and hence volatility. Our data-driven analysis later in the paper determines which of the two explanations has more compelling empirical support. In this section, we provide a review of underlying theory.

³See, for example, Figlewski (1989), Bollen and Whaley (2004), and Garleanu, Pedersen, and Poteshman (2009).

⁴FIA Annual Volume survey (2013).

As described in Li and Wu (2006), classic market microstructure theory in Glosten and Milgrom (1985), Admati and Pfleiderer (1988), and Foster and Vishwanathan (1990), among others, links trading volume and price changes to the arrival of information. The mixture of distribution hypothesis of Clark (1973) is commonly used to study this issue, positing a joint dependence of price changes and volume on latent information flow, with mixed results. Clark (1973), Epps and Epps (1976), and Tauchen and Pitts (1983) find support for the model. Later studies, including Heimstra and Jones (1994), Lamoureux and Lastrapes (1994), and Richardson and Smith (1994), question the earlier findings. Andersen (1996) reconciles the conflicting results of prior studies by specifying an alternative process of information arrival.

Li and Wu (2006) find that the positive volume-volatility relation is driven by informed trading and the rate of information flow. After controlling for information flow the relation turns negative, consistent with information-free “liquidity trading” dampening volatility. Theoretically, the negative relation between liquidity trading and price volatility exists in any Bayesian learning model, in which market makers infer information from sequence of trades.⁵ As argued by Kyle (1985) an increase in liquidity trading increases the depth of the market, that is, the willingness of liquidity suppliers to commit to trading, which leads to greater price stability. Market makers and other liquidity suppliers can accommodate random variations in order flow without suffering any price impact. Chordia, Roll, and Subrahmanyam (2001), for example, find that higher liquidity is generally associated with lower volatility in the absence of information. These results suggest that if the rise of algorithmic trading has generated additional information-free trading volume, then volatility may in fact be reduced.

An alternative view of recent changes in trading practices, typically presented in the popular press as opposed to academic research, makes several claims that together suggest that the quality of financial markets has deteriorated in recent years due to high-frequency trading. High-frequency traders, for example, are described as being selective in their provision of liquidity, so that the theoretical results described above do not apply. In addition, order submission and cancellation strategies are employed at the millisecond level in order to learn about trading demands of other market participants, enabling front-running strategies that exacerbate the short-term price impact of trades. Consequently, the microstructure component of volatility may be expected to increase.

We turn next to a description of the historical evolution of volatility as well as different ways of measuring volatility as we begin to differentiate between these two explanations.


3. VOLATILITY, MACROECONOMIC EVENTS, AND REALIZED VOLATILITY MEASUREMENT

Volatility changes through time depending on the rate that new information arrives in the marketplace. To illustrate, Figure 1 shows the behavior of U.S. stock market volatility from January 1986 through June 2012 as measured by the CBOE’s market volatility index or VIX.⁶ Two features are salient.

First, there are periodic spikes in volatility, 20 of which are associated with macroeconomic events in Figure 1. Unexpected changes in interest rates, the probability of recession, and the rate of bank failures are some events that increase the level of anxiety in the marketplace. Two large spikes are far larger than the others. The first, labeled event number four, corresponds to the stock market crash of October 19, 1987, a day on which the Dow Jones

⁵See, for example, Madhavan (1995).

⁶The series plotted in Figure 1 is actually the VXO, the original form of the VIX when it was released in 1993. For an explanation of the differences between the indexes, see Whaley (2009).



Legen	Date	VIX	Event
1	3/21/1986	26.91	OPEC agrees to drop production resulting in a sharp increase in crude oil prices
2	9/12/1986	27.69	Inflation fears and portfolio insurance programs are blamed for a one-day 4.61% drop in the DJIA
3	4/27/1987	31.46	Dollar falls to 39-year low against the yen and inflation hits 5%
4	10/19/1987	150.19	DJIA drops over 22% on the day called "Black Monday"
5	8/23/1990	38.07	Saddam Hussein appears on state television with Western hostages following the August 2 Iraqi invasion of Kuwait
6	1/15/1991	36.93	Iraq ignores U.N. deadline for withdrawal from Kuwait prompting the beginning of Operation Desert Storm
7	10/7/1992	21.12	Pessimistic economic statistics fuel recession fears
8	4/4/1994	22.50	Stocks drop as long-term interest rates rise unexpectedly
9	3/8/1996	24.37	DJIA drops 3% in contrarian fashion following job growth, lowering likelihood of Fed stimulus
10	7/23/1996	24.43	Jagged trading triggers the NYSE uptick rule for the seventh consecutive trading day
11	10/27/1997	39.96	Stock markets plummet worldwide due to Asian economic crisis
12	8/31/1998	48.33	DJIA drops 19% in August in the weeks following the Russian Default
13	4/14/2000	39.33	Nasdaq drops 25% in one week ushering in the post-bubble period
14	3/22/2001	39.70	CPI rises more than expected, dampening hopes of Fed rate cut
15	9/20/2001	49.04	Markets re-open following September 11 terrorist attack
16	10/9/2002	49.48	Stocks reach 2002 lows culminating an 18-month drop from dot-com era peak
17	9/15/2008	31.70	Lehman Brothers files for Chapter 11 bankruptcy protection
18	11/20/2008	87.24	S&P 500 drops to an 11 1/2 year low following continued signs of economic contraction
19	5/20/2010	43.63	U.S. stock indices fell into correction following continued evidence of a slow economic recovery
20	8/8/2011	50.13	First trading day following S&P downgrade of U.S. credit rating; fears of European debt crisis mount

FIGURE 1

Volatility from January 1986 to June 2012.

Shown is the daily closing level of the CBOE Volatility Index (VIX) from January 1986 to June 2012. Spikes corresponding to 20 important events are indicated. Note that the VIX closed at 150.19 on the stock market crash of September 19, 1987, although the scale of the chart is capped at 90.

Industrial Average dropped 22%. The VIX actually closed over 150 on that day, well beyond the maximum on the figure's scale. The second extreme spike, labeled number 18, occurred at the height of the current global financial crisis on November 20, 2008, when the VIX closed at about 87. These spikes are quite transient, indicating that implied volatility, as a measure of perceived uncertainty and risk, can change as quickly as the rate of information flow.

Second, volatility appears to go through long periods of relatively elevated or depressed levels. The 5-year period of 1992 through 1996, for example, was relatively benign, and featured an average VIX level of just 14.3. In contrast, the VIX almost doubled to 26.7 over the following 6-year period of 1997 through 2002, which included a number of important events, including the Asian crisis, the September 11, 2001 attacks, and massive drops in stock prices, especially those of technology companies listed on NASDAQ. Similarly, VIX averaged just 14.1 from July 2003 through June 2007, but then almost doubled to 26.2 over the last 5 years of the sample, July 2007 through June 2012, corresponding to the global financial crisis.

Figure 2 provides a more granular view of the VIX in the time surrounding the financial crisis. Macroeconomic phenomena including the European debt crisis, bankruptcy filings, and regulatory uncertainty all contributed to prolonged levels of high volatility reaching into 2011. Markets were rocked by events that no one could have foreseen, including a downgrade of the United States. Treasury's credit rating and a tsunami that spawned a nuclear disaster in Japan. This long list of fundamental sources of risk and uncertainty complicates any study of market microstructure and its potential impact on volatility. As a consequence, in Section 6, we develop two measures that are relatively free of microstructural effects in order to construct appropriate benchmarks for futures return volatility. But first we describe the variety of approaches we use to measure realized volatility.

3.1. Constructing Daily Realized Volatility

Daily measurements of realized volatility are constructed by first dividing each 24-hour day on which trading occurs into 288 5-minute periods, denoted by t where $t = 1, \dots, 288$. For each 5-minute period within which a trade occurred, the last trade price is recorded, denoted by p_t . Starting with the second 5-minute period, a return is computed as:

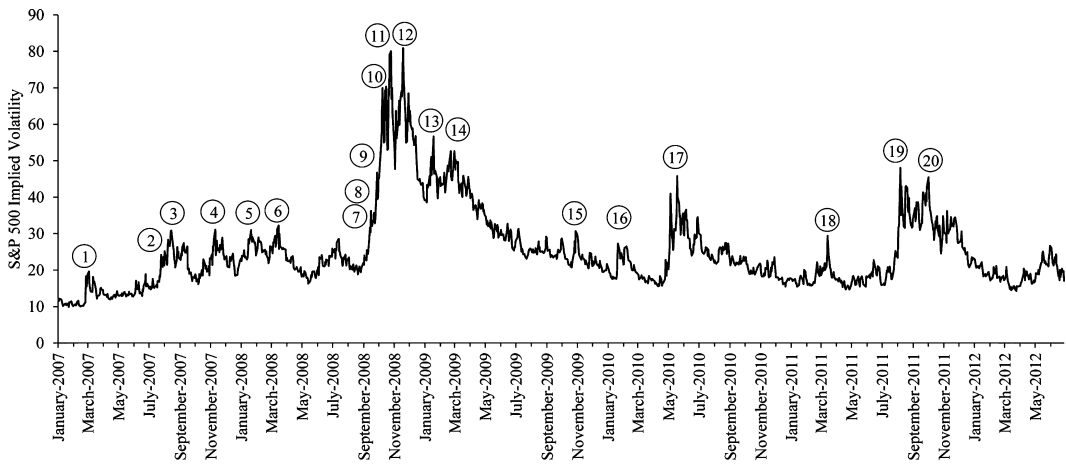
$$r_t = \ln\left(\frac{p_t}{p_{t-1}}\right) \quad (1)$$

if both period t and the previous contained a trade, otherwise the period does not contribute to the day's variance measurement. Realized return variance is then computed as the sum of squared returns, scaled as follows:

$$\frac{288}{n} \sum_{t=2}^{288} r_t^2 \quad (2)$$

where n is the number of 5-minute returns recorded during the period. Scaling by 288 standardizes the measure to allow comparison across days and across contracts with different trading hours. We define "daily realized volatility" as the square root of Equation (2).

Empirical analyses of realized volatility constructed from 5-minute returns include studies by Andersen, Bollerslev, Diebold, and Labys (2001) and Andersen, Bollerslev, Diebold, and Ebens (2001) of exchange rates and stock returns, respectively. In both papers, the focus is on the distribution of realized volatility as well as its serial correlation. Realized volatility features positive skewness and substantial excess kurtosis, whereas log realized volatility appears close to Gaussian, which is exploited in subsequent statistical analysis. Thomakos and Wang (2003) find similar results using 5-minute returns of Treasury Bond,



Legend	Date	VIX	Event
1	2/27/2007	18.30	Federal Home Loan Mortgage Corp announces it will no longer buy riskiest subprime securities
2	7/31/2007	23.52	Bear Stearns liquidates two hedge funds that invested in MBS
3	8/16/2007	30.83	Fitch downgrades Countrywide Financial Corp to BBB+
4	11/12/2007	31.09	Bank of America, Citigroup, JPMorgan agree to establish a \$75 billion fund to buy troubled assets
5	1/22/2008	31.01	FOMC votes to reduce Federal Funds rate by 75 basis points to 3.5%
6	3/14/2008	31.16	Federal Reserve approves JPMorgan bail out of Bear Stearns
7	9/15/2008	31.70	Lehman Brothers files for Chapter 11 bankruptcy protection
8	9/25/2008	32.82	JPMorgan wins bid to acquire Washington Mutual in FDIC orchestrated auction
9	9/29/2008	46.72	U.S. House of Representatives rejects legislation to authorize the U.S. Treasury to purchase troubled assets
10	10/17/2008	70.33	Disappointing economic statistics lead to dramatic daily changes in equity index levels
11	10/27/2008	80.06	U.S. Treasury injects \$125 billion into nine major U.S. banks
12	11/20/2008	80.86	S&P 500 drops to an 11 1/2 year low following continued signs of economic contraction
13	1/20/2009	51.00	U.K. banking crisis intensifies; Barack Obama inauguration
14	3/5/2009	50.17	U.S. equity markets reach new lows dragged down by financials, including Citigroup, which trades at less than \$1 per share
15	10/30/2009	30.69	VIX increases by 38% in one week reflecting fears of slowing recovery
16	1/22/2010	27.31	U.S. stocks drop by 2% over concerns of President Obama's banking reform plans
17	5/20/2010	45.79	U.S. stock indices fell into correction following continued evidence of a slowing economic recovery
18	3/16/2011	29.40	Fukushima Nuclear Power Plant situation worsens following Japanese tsunami
19	8/8/2011	48.00	First trading day following S&P downgrade of U.S. credit rating; fears of European debt crisis mount
20	10/3/2011	45.45	Greece misses a deficit target despite austerity measures increasing probability of bankruptcy

FIGURE 2
Crisis timeline.

Shown is the daily closing level of the CBOE Volatility Index (VIX) from January 2007 to June 2012. Spikes corresponding to 20 of the important events of the global financial crisis are indicated.

S&P 500, and Eurodollar futures contracts. Thomakos and Wang mitigate the impact of bid-ask bounce by filtering the data using an MA(1) model and using the resulting residuals as a proxy for true returns. Daily log volatility is persistent with autocorrelations significant for over 50 days in all three studies, consistent with the voluminous GARCH literature that relies on estimates of latent volatility inferred from daily returns. Realized volatility is also used in recent work on high-frequency trading, as in Brogaard (2012).

3.2. Range-Based Estimators of Realized Volatility

For robustness, we compute several range-based estimators of realized volatility, which use only the open, high, low, and closing prices observed in a trading day. Range-based estimators are used in several studies of high-frequency trading, including Hendershott, Jones, and Meukveld (2011), who construct daily volatility estimates. Similarly, Hasbrouck and Saar (2011) use the spread between high and low midpoint quotes over 10-minute windows as a measure of volatility over a much shorter horizon.

Let O, H, L , and C denote these prices and define percentage changes from the open as $u \equiv \ln(H/O)$, $d \equiv \ln(L/O)$, and $c \equiv \ln(C/O)$. Parkinson (1980) developed the first range-based estimator using only high and low prices. The intuition is that the higher the volatility, the larger the observed range of prices observed over the course of a trading day. Using the above definitions, and the range observed over n trading days, the Parkinson estimator can be computed as:

$$V_P = \frac{1}{4 \ln(2)} \times \frac{1}{n} \sum_{i=1}^n (u_i - d_i)^2 \quad (3)$$

The Parkinson estimator is valid only for processes with zero-drift. Rogers and Satchell (1991) derive an alternative estimator that accommodates a non-zero drift and also has substantially lower sampling error. The Rogers and Satchel estimator, again using prices from n trading days, can be computed as:

$$V_{RS} = \frac{1}{n} \sum_{i=1}^n (u_i - c_i)u_i + (d_i - c_i)d_i \quad (4)$$

We compute both estimates and use non-overlapping daily, weekly, and monthly intervals in the calculations.⁷

3.3. GARCH Processes for Time Variation in Volatility

As discussed above, one of the most salient features of volatility is its time variation. Consequently, it will be useful to implement processes that explicitly accommodate changes in volatility. We use the GARCH(1,1) volatility model, which, for daily returns y_t , can be expressed as:

$$\begin{aligned} y_t &= \mu + \varepsilon_t, \\ \varepsilon_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &= \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \end{aligned} \quad (5)$$

⁷Yang and Zhang (2000) modify the Rogers and Satchel estimator to reflect the volatility of overnight returns, which they describe as “opening jumps” between the prior day’s close and the current day’s open. Since we focus on intraday volatility, we use the Rogers and Satchel estimator.

Two aspects of the GARCH estimates will be especially useful. First, the coefficient β on lagged variance provides a measure of the speed with which volatility reverts to average levels. Short-lived spikes in volatility will result in faster convergence to the long-run mean and lower values for β . Second, the long-run variance implied by the GARCH(1,1) model is given by $\omega/(1 - \alpha - \beta)$, which can differ from sample estimates given the ability of the GARCH(1,1) model to accommodate short-run changes in volatility.

4. DATA

Our analysis is based on electronic trade-by-trade data provided by the Intercontinental Exchange (ICE), Eurex, NYSE Euronext (NYSE Liffe), and The CME Group (CME).⁸ While the trade-by-trade data, in some cases, contain spread trades and block trades, these were eliminated from the subsequent analyses to focus in an unfettered way on trading activity. The 15 specific contracts are listed in Table I, together with their ticker symbols and time-series start and end dates. For the remainder of the paper, contracts are referred to by their ticker symbols. Seven are interest rate futures contracts, five are stock index futures contracts, two are crude oil futures contracts, and one is an agricultural futures contract. Three of the exchanges also provided end-of-day data, including the daily open, high, low, and closing prices of the futures contracts, which are used to construct range-based estimates of volatility. End-of-day data were obtained from Price-Data.com as a check on data integrity. Daily data for three popular stock market volatility indexes—the VIX, VDAX, and VFTSE—were downloaded from Datastream.

Before examining return volatility in the 15 contracts, several market microstructure variables are estimated to gain insight regarding the effect of changes in trading practices. In particular, we measure changes in average trade size in contracts and in price clustering. Trade size indicates whether trading practices have changed over the sample period considered. Prices of many asset types tend to cluster on certain multiples of a given market's minimum price increment or tick size. As noted by Brown, Laux, and Schacter (1991) and Harris (1991), smaller tick sizes can reduce the average speed of execution by increasing the number of possible prices at which to trade, thereby complicating negotiation. Bollen and Christie (2009) show that markets can endogenously create optimal tick sizes by clustering trade prices appropriately. Ball, Torous, and Tschoegl (1985) and Harris (1991) find that stocks with higher volatility and lower trade frequency feature more clustering, suggesting that the degree of clustering increases in the difficulty in obtaining information about an asset. These results motivate us to test whether clustering has changed in futures markets reflecting any change in market quality. To illustrate, Figure 3 shows the percentage of trade prices at each decimal increment observed for the CL contract. If the 100 possible increments were used equally we would expect each to occur about 1% of the time. Though all are used, those ending on multiples of "nickels" are much frequent.

Table II shows changes in the two microstructure variables. For the contracts traded on ICE (B, SB, and TF) we report results for the years 2008 and 2011 ("Pre" and "Post") since SB and TF do not begin trading until 2008. For all other contracts we report results for 2004 and 2011 to gain a longer-term perspective on changes in the markets.

⁸The length of the different time series varies from exchange to exchange, and the time-series generally end in May 2012.

TABLE I
Fifteen Futures Contract Time-Series Provided by Futures Exchanges

<i>Exchange</i>	<i>Contract</i>	<i>Ticker</i>	<i>Begins</i>	<i>Ends</i>
Intercontinental Exchange (ICE)	Brent Crude	B	11/4/2004	7/9/2012
	Russell 2000 Index	TF	1/2/2008	7/9/2012
	Sugar #11	SB	1/2/2008	7/9/2012
Eurex	DAX	FDAX	5/2/2002	5/15/2012
	Euro-Stoxx 50 Index	FESX	5/2/2002	5/15/2012
	Euro-Bund	FGBL	5/2/2002	5/15/2012
	Euro-Bobl	FGBM	5/2/2002	5/15/2012
NYSE Liffe	FTSE 100 Index	Z	1/4/2000	5/31/2012
	3-Month Euro (Euribor)	I	1/4/2000	5/31/2012
	3-Month Sterling (Short Sterling)	L	1/2/2001	5/31/2012
	Long Gilt	R	1/4/2000	5/31/2012
CME Group (CME)	Eurodollar	ED	8/6/1992	5/31/2012
	E-mini S&P 500 Index	ES	1/1/2000	5/31/2012
	Light Sweet (WTI) Crude Oil	CL	11/30/1999	5/31/2012
	10-Year U.S. Treasury Note	TY	1/1/2004	5/31/2012

Note. Listed are futures contracts time-series that serve as the basis of our analysis. The time and sales data provided by the exchanges contain time-stamped trade information. The end-of-day (EOD) summary data contain daily open, high, low, and close prices as well as number of contracts traded and open interest, and were also provided by the exchanges. The Price-Data data are also daily summary data and were purchased from Price-Data.com.




FIGURE 3

Clustering in the CL contract.

Shown is the percentage of trade prices in the CME's crude oil futures contract (CL) at each possible decimal increment. The 615,191 prices used are recorded from the last trade in each 5-minute period from December 1, 1999 through May 15, 2012.

TABLE II
Market Microstructure

Ticker	A. Clustering						B. Trade size		
	Pre		Post		Change		Pre	Post	P-value
	Freq (%)	P-value	Freq (%)	P-value	Δfreq (%)	P-value			
CL	1.63	0.0000	1.16	0.0000	-0.47	0.0000	7.68	2.13	0.0000
ED	5.22	0.0593	4.76	0.0148	-0.45	0.0024	67.43	45.37	0.0000
ES	25.83	0.0000	25.66	0.0001	-0.17	0.4819	12.70	12.80	0.7474
TY	24.62	0.0314	25.13	0.4314	0.51	0.0342	31.03	17.17	0.0000
FDAX	5.34	0.0043	6.27	0.0000	0.93	0.0000	6.41	4.56	0.0000
FESX	0.77	0.0000	1.91	0.0000	1.14	0.0000	51.88	46.24	0.0014
FGBL	0.84	0.0038	1.30	0.0000	0.45	0.0000	97.84	35.03	0.0000
FGBM	1.22	0.0000	1.17	0.0005	-0.05	0.4962	117.16	49.43	0.0000
B	1.88	0.0000	1.45	0.0000	-0.43	0.0000	1.53	1.51	0.1796
SB	0.95	0.2875	0.90	0.0710	-0.05	0.5220	4.01	1.96	0.0000
TF	1.27	0.0000	1.07	0.0723	-0.20	0.0083	1.54	1.35	0.0000
I	5.27	0.0336	4.84	0.1508	-0.44	0.0112	80.99	41.08	0.0000
L	1.35	0.0000					88.52	54.35	0.0000
R	0.95	0.4306	0.88	0.0329	-0.08	0.3118	10.02	6.48	0.0000
Z	5.23	0.0765	6.16	0.0000	0.93	0.0000	4.07	2.82	0.0000

Note. Listed for each contract is the frequency with which the decimal component of trade prices are observed with last two digits equal to 50 for a “Pre” period (2008 for B, SB, and TF; 2004 for all others) and a “Post” period (2011 for all contracts). The *P*-value tests whether the frequency is statistically significantly different from expected under the null hypothesis of a uniform distribution given the contract’s tick size. For ES and TY the null predicts a 25% probability. For ED, FDAX, I, and Z the null predicts a 5% probability. For all other contracts the null predicts a 1% probability. Also listed is the change in the frequency from 2008 to 2011. The *P*-value tests whether the frequencies are statistically significantly different from each other.

The clustering measure is the percentage of trade prices with the last two digits of the decimal increment equal to 50. Given the contracts’ tick sizes, the null hypothesis predicts 25% for the ES and TY contracts, 5% for the ED, FDAX, I, and Z contracts, and 1% for all others. Ten of the 15 contracts feature significant clustering in the Pre-period, though the economic magnitudes are modest as expected in the highly liquid futures markets. In the Post-period, the L contract is not considered given the extremely low short-term interest rates in 2011, which make clustering irrelevant. Of the remaining 14 contracts, nine feature significant clustering. More importantly, only four contracts feature significant increases in clustering across the two periods, the FDAX, FESX, FGBL, and Z contracts. Note that these are all European contracts, and the increase in clustering may be explained by the increase in fundamental uncertainty in the financial crisis period.

In contrast, the average trade size in contracts has fallen substantially for all but two of the contracts. The average trade size for the CL contract fell from 7.68 contracts to 2.13, for example, likely reflecting the increased presence of computer-driven trading. We turn next to an assessment of whether these microstructure changes are associated with changes in volatility.

5. REALIZED VOLATILITY ESTIMATES

This section reports estimates of realized volatility using 5-minute returns of the 15 futures contracts in our sample.

5.1. Summary Statistics

Tables III and IV contain summary statistics of realized volatility. Realized volatility is computed daily, as described in Section 3, and is annualized. Only a single contract on an underlying is used in a given day. Usually, it is the contract with the highest trading volume. The days on which we roll from the nearby to the second nearby contracts range from 1 to 45 days before expiration. For some contracts, like ED, only quarterly expirations were used since the trading volumes of non-quarterly expiration contracts are meager.

Table III contains summary statistics of realized volatility estimated over each contract's entire sample period. The average volatilities across asset classes are in line with expectations. For the CME contracts, for example, CL volatility is highest at 31.7%. ED volatility, at the other end of the risk spectrum, is at 1.1%.⁹ The ES and the TY contracts had average volatilities of 20.0% and 6.6%, respectively. Average levels of volatility across exchanges generally correspond along asset classes. For equities, the Z contract traded on the NYSE Liffe features average volatility of 24.7%, slightly higher than the ES volatility traded on the CME. The volatility of the TF contract, traded on the ICE, is higher still, at 34.4%, reflecting the inverse relation between firm size and volatility. The FDAX and FESX contracts, traded on Eurex, have average volatilities of 30.4% and 35.3%, respectively, reflecting the uncertainty surrounding the debt crises in a number of European countries as well as uncertainty about the future of the Euro.

For short-term fixed income securities, the I and L contracts, traded on the NYSE Liffe, have average volatilities of 1.0% and 1.7%, respectively, similar to the ED volatility traded on the CME. For longer-term fixed income securities, the FGBL and FGBM contracts, traded on the Eurex, feature average volatilities of 8.7% and 5.6%, respectively, comparable to the TY contract volatility traded on the CME. Volatility of the R contract, traded on the NYSE Liffe, is at 8.3%.

Commodities tend to have the highest volatilities, with the B and SB contracts traded on ICE featuring volatilities of 37.9% and 56.8%, respectively, the former similar in magnitude to the CL volatility traded on the CME.

For all contracts, the Jarque–Bera test easily rejects the hypothesis that volatility is normally distributed, consistent with prior research such as Thomakos and Wang (2003).

Table IV shows summary statistics for the 12 contracts with data extending back to January 2004 over two sub-periods: a “Pre-crisis” period from January 2004 through June 2007, and a “Crisis” period from July 2007 through the first half of May 2012. Dramatic increases in the average level and volatility of volatility are observed in a number of contracts. For the four equity index contracts, ES, FDAX, FESX, and Z, for example, the average level of volatility doubles in each case. Substantial increases are also present in CL, FGBL, and FGBM. As described in Section 3, these increases can be attributed to the uncertainty created by the global financial crisis.

5.2. GARCH Processes for Time Variation in Volatility

As discussed in Section 3, Figures 1 and 2 illustrate how volatility has undergone cycles of high and low levels over the sample period, and has featured numerous spikes attributable to

⁹Note that for purposes of comparison, Eurodollar volatility is being expressed in terms of percent change in price. In practice, however, Eurodollar volatility is most often quoted in terms of percent change in the Eurodollar interest rate.

TABLE III
Summary Statistics

<i>CME</i>				
	<i>CL</i>	<i>ED</i>	<i>ES</i>	<i>TY</i>
# Obs.	3,114	2,148	3,118	2,353
First	19991201	20040105	20000103	20040102
Last	20120515	20120515	20120515	20120515
Avg.	31.66	1.07	20.01	6.64
Std. dev	16.44	0.50	13.48	2.91
Skewness	2.43	3.34	6.11	2.69
Kurtosis	11.59	19.22	98.08	19.27
J-B	12,632.75	27,550.55	1,193,850.81	28,794.25
<i>P</i> -value	0.00	0.00	0.00	0.00
<i>Eurex</i>				
	<i>FDAX</i>	<i>FESX</i>	<i>FGBL</i>	<i>FGBM</i>
# Obs.	2,555	2,558	2,559	2,559
First	20020502	20020502	20020502	20020502
Last	20120515	20120515	20120515	20120515
Avg.	30.37	35.28	8.70	5.61
Std. dev	20.65	26.12	6.18	3.49
Skewness	2.39	2.67	3.00	2.93
Kurtosis	12.52	14.63	14.23	15.25
J-B	12,075.21	17,457.76	17,274.91	19,661.81
<i>P</i> -value	0.00	0.00	0.00	0.00
<i>ICE</i>				
	<i>B</i>	<i>SB</i>	<i>TF</i>	
# Obs.	1,877	1,101	1,065	
First	20050214	20080102	20080319	
Last	20120515	20120515	20120515	
Avg.	37.85	56.84	34.40	
Std. dev	18.08	17.42	20.36	
Skewness	2.37	1.29	2.34	
Kurtosis	10.05	6.65	11.75	
J-B	5,650.15	914.37	4,371.38	
<i>P</i> -value	0.00	0.00	0.00	
<i>NYSE Liffe</i>				
	<i>I</i>	<i>L</i>	<i>R</i>	<i>Z</i>
# Obs.	3,002	2,093	3,036	2,935
First	20000104	20010102	20000104	20000104

continued

TABLE III
(Continued)

	<i>NYSE Liffe</i>			
	<i>I</i>	<i>L</i>	<i>R</i>	<i>Z</i>
Last	20120515	20120515	20120515	20120515
Avg.	1.00	1.69	8.27	24.72
Std. dev	0.37	0.58	3.14	15.23
Skewness	4.37	5.51	2.00	2.69
Kurtosis	36.76	72.80	11.82	15.06
J-B	152,094.52	435,519.13	11,873.67	21,336.56
P-value	0.00	0.00	0.00	0.00

Note. Listed are summary statistics of daily observations of realized volatility, computed as the sum of squared 5-minute returns using the last trade price within each 5-minute interval. Each day's realized volatility is scaled to reflect trading in all 288 intervals. Realized volatility is converted to annualize volatility assuming a 252-trading day year.

macroeconomic events. To provide some insight regarding the nature of the time variation in volatility, we plot in Figure 4 daily returns of the three equity index contracts, ES, FDAX, and Z, for which implied volatility indexes are available.

For all three equity index contracts, the daily returns feature classic volatility clustering throughout the sample. Large swings in returns indicate periods corresponding to rapid changes in the macroeconomic environment. Note also that the correspondence between daily variation in returns and daily levels of the volatility indexes is tight, suggesting a strong link between relatively high-frequency observations of volatility and the 30 calendar day measures of implied volatility. For this reason, we use implied volatility as one of our volatility benchmarks to provide a test for a change in the role of microstructural effects on realized volatility controlling for changes in the rate of information flow.

To determine whether volatility processes themselves have changed over the pre-crisis and crisis periods, we estimate parameters of a GARCH(1,1) for each contract. The results are reported in Table V. For four of the contracts (FGBL, B, L, R), we fail to reject a constant volatility model in favor of the GARCH(1,1) during the pre-crisis period, and for these the constant volatility estimate is listed. Perhaps the most important result here is that the long-run volatility implied by the GARCH parameters features substantial increases for some of the contracts, though not nearly as large as the raw averages. The FDAX long-run volatility, for example, increases from 13.4% to 24.1% using the GARCH parameters, whereas the average volatility listed in Table IV increases from 17.0% to 35.0%. The reason for this is that the GARCH model accounts for the transience of spikes in volatility.

We provide one additional analysis of the time-series behavior of realized volatility by computing the autocorrelation function at the daily frequency with 100 lags. Figures 5 through 7 display the autocorrelation functions of the three equity index contracts featured above, the ES, FDAX, and Z, respectively. The high level of serial correlation at short lags, and the slow decay, reflects the volatility clustering shown in Figure 4. More importantly, the differences in the pre-crisis and crisis periods in the bottom panels is stark—in all three contracts serial correlation has increased substantially in recent years, again likely the reflection of an increase in underlying fundamental volatility.

TABLE IV
Summary Statistics Over Pre-Crisis and Crisis Subsets

	<i>Panel A. Pre-crisis</i>				<i>Panel B. Crisis</i>			
	<i>CME</i>				<i>CME</i>			
	<i>CL</i>	<i>ED</i>	<i>ES</i>	<i>TY</i>	<i>CL</i>	<i>ED</i>	<i>ES</i>	<i>TY</i>
# Obs.	903	878	881	879	1,510	1,270	1,232	1,474
Avg.	23.80	0.95	11.54	5.58	34.97	1.15	23.25	7.26
Std. dev	7.13	0.32	3.34	2.38	19.35	0.58	15.42	3.01
Skewness	1.88	5.60	1.26	4.77	2.00	2.69	3.15	2.24
Kurtosis	10.49	56.73	4.94	39.12	7.99	13.24	19.13	17.20
J-B	2,645.00	110,216.63	371.88	51,100.76	2,570.83	7,075.48	15,392.93	13,623.46
P-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

	<i>Panel A. Pre-crisis</i>				<i>Panel B. Crisis</i>			
	<i>Eurex</i>				<i>Eurex</i>			
	<i>FDAX</i>	<i>FESX</i>	<i>FGBL</i>	<i>FGBM</i>	<i>FDAX</i>	<i>FESX</i>	<i>FGBL</i>	<i>FGBM</i>
# Obs.	892	894	894	894	1,242	1,242	1,242	1,242
Avg.	17.04	18.12	5.34	3.74	35.04	44.85	11.50	7.05
Std. dev	5.85	5.72	1.59	1.69	22.61	30.31	7.68	4.18
Skewness	1.34	1.65	2.92	9.17	2.58	2.44	2.16	2.20
Kurtosis	6.22	7.50	19.21	133.97	13.49	12.19	8.26	9.49
J-B	653.03	1,157.51	11,060.26	651,460.95	7,078.67	5,600.01	2,401.08	3,179.84
P-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

	<i>Panel A. Pre-crisis</i>				<i>Panel B. Crisis</i>			
	<i>NYSE Liffe</i>				<i>NYSE Liffe</i>			
	<i>I</i>	<i>L</i>	<i>R</i>	<i>Z</i>	<i>I</i>	<i>L</i>	<i>R</i>	<i>Z</i>
# Obs.	864	468	861	860	1,227	1,127	1,203	1,209
Avg.	0.81	1.58	6.04	13.95	1.07	1.66	10.05	29.26
Std. dev	0.15	0.41	1.76	4.84	0.45	0.64	3.37	17.60
Skewness	1.68	2.12	3.55	2.66	4.18	6.60	2.29	2.69
Kurtosis	11.31	15.60	29.67	16.22	30.59	86.22	13.43	13.43
J-B	2,892.37	3,447.01	27,329.87	7,271.99	42,481.45	333,365.01	6,499.79	6,936.41
P-value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Note. Listed are summary statistics of daily observations of realized volatility, computed as the sum of squared 5-minute returns using the last trade price within each 5-minute interval. Each day's realized volatility is scaled to reflect trading in all 288 intervals. Realized volatility is converted to annualized volatility assuming a 252-trading day year. "Pre-crisis" uses data from January 2, 2004 through June 29, 2007. "Crisis" uses data from July 2, 2007 through May 25, 2012.




FIGURE 4

Daily returns and volatility indexes.

Shown are daily open-to-close returns of the ES, FDAX, and Z contracts along with the closing levels of the VIX, VDAX, and VFTSE volatility indexes. The three data series begin on January 3, 2000, January 4, 2000, and October 24, 2006, respectively. All series run through May 15, 2012.

TABLE V
GARCH Parameter Estimates

Ticker	Panel A. Pre-crisis				Panel B. Crisis			
	ω	α	β	LR	ω	α	β	LR
CL	0.0977	0.0236	0.9408	26.30	0.0445	0.0457	0.9419	30.11
	0.0966	0.0202	0.0000		0.0558	0.0000	0.0000	
ED	0.0001	0.3100	0.5754	0.55	0.0000	0.1442	0.8458	0.69
	0.2834	0.0801	0.0244		0.0806	0.0002	0.0000	
ES	0.0395	0.0537	0.8527	10.32	0.0373	0.1135	0.8667	21.76
	0.0151	0.0091	0.0000		0.0055	0.0000	0.0000	
TY	0.0017	0.0202	0.9564	4.29	0.0052	0.0476	0.9200	6.38
	0.3423	0.0426	0.0000		0.9033	0.7459	0.0235	
FDAX	0.0284	0.0306	0.9299	13.45	0.0740	0.1006	0.8674	24.12
	0.0924	0.0894	0.0000		0.0225	0.0000	0.0000	
FESX	0.0196	0.0246	0.9462	12.98	0.0808	0.1065	0.8692	28.98
	0.0972	0.1213	0.0000		0.0027	0.0000	0.0000	
FGBL	0.0544			3.70	0.0035	0.0546	0.9237	6.42
	0.0000				0.0706	0.0001	0.0000	
FGBM	0.0002	0.0113	0.9787	2.30	0.0009	0.0412	0.9454	4.10
	0.8703	0.5111	0.0000		0.0558	0.0033	0.0000	
B	2.3716			24.45	0.0336	0.0517	0.9364	26.63
	0.0000				0.1218	0.0009	0.0000	
I	0.0000	0.2302	0.7519	0.40	0.0001	0.2234	0.7471	0.81
	0.0016	0.0000	0.0000		0.1534	0.0657	0.0000	
L	0.0006			0.40	0.0000	0.1590	0.8310	0.70
	0.0000				0.0504	0.0001	0.0000	
R	0.0594			3.87	0.0060	0.0391	0.9255	6.55
	0.0000				0.7620	0.3915	0.0000	
Z	0.0341	0.1044	0.7941	9.21	0.0544	0.1332	0.8459	25.56
	0.0105	0.0005	0.0000		0.0067	0.0000	0.0000	

Note. Listed are GARCH(1,1) parameter estimates based on de-meaned daily open-to-close returns. Below each coefficient estimate is the associated *P*-value testing for statistical significance. Also listed is the long-run annualized volatility (LR) implied by the parameter estimates. “Pre-crisis” uses data from January 2, 2004 through June 29, 2007. “Crisis” uses data from July 2, 2007 through May 15, 2012. For four of the contracts (FGBL, B, L, R) we failed to reject a constant volatility model in favor of the GARCH(1,1) during the pre-crisis period, and for these the constant volatility estimate is listed.

5.3. Range-Based Estimator as a Robustness Test

To test the robustness of our 5-minute return volatility measures, we compare them to the Rogers and Satchell (1991) range-based estimator that relies only on a daily record of the open, high, low, and closing prices for a contract. We compute both of the estimators over three observation windows: daily, non-overlapping five-trading-day periods, and non-overlapping 21-trading-day periods. The latter two correspond to weekly and monthly measurements. For each of the 15 futures contracts, we measure the linear correlation and the Spearman rank correlation between the two estimators over the full sample for each contract.

ACF for ES Realized Variance 1/3/2000 - 5/15/2012

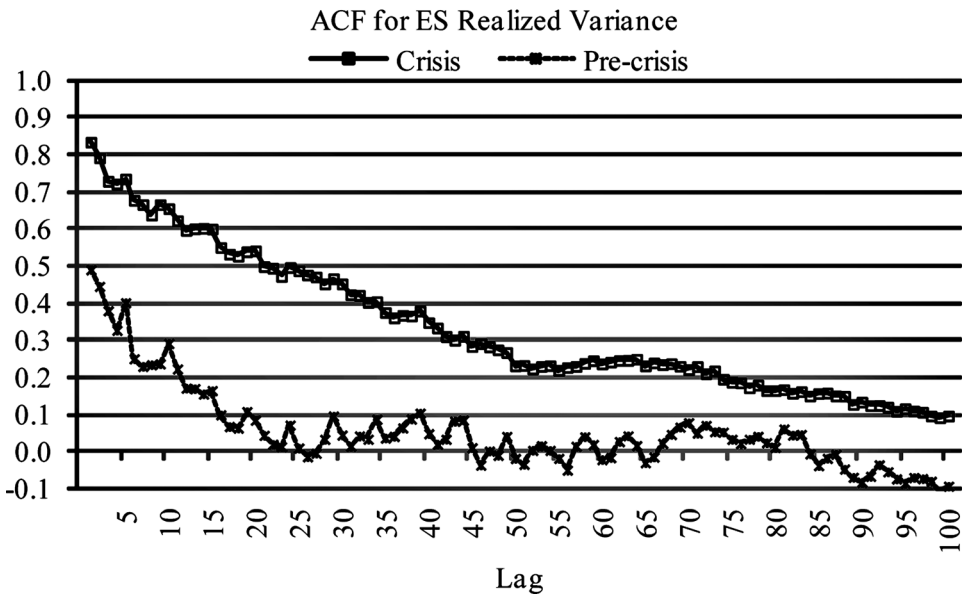



FIGURE 5

Autocorrelation functions of the E-mini contract.

Shown are autocorrelation functions of daily measures of realized variance based on 5-minute squared returns for the S&P 500 E-mini contract (ES) traded on the CME. The bottom figure shows the functions estimated over two subsets. "Crisis" uses data from July 2, 2007 through May 25, 2012. "Pre-crisis" uses data from January 2, 2004 through June 29, 2007.

The correlation results, reported in Table VI, are noteworthy in a number of respects. First, the correlations tend to increase with the length of the measurement window. For the ES contract, for example, the linear correlation between the two volatility estimators is 0.40 at the daily frequency compared to 0.73 at the weekly frequency. Second, the correlations are

ACF for FDAX Realized Variance 1/3/2000 - 5/15/2012

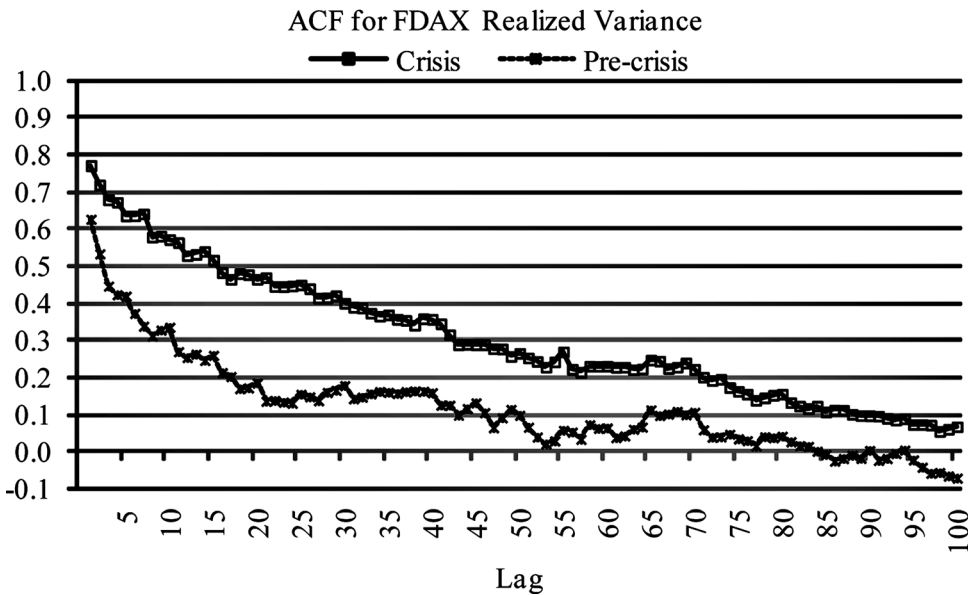



FIGURE 6

Autocorrelation functions of the DAX contract.

Shown are autocorrelation functions of daily measures of realized variance based on 5-minute squared returns for the DAX contract (FDAX) traded on the Eurex. The bottom figure shows the functions estimated over two subsets. “Crisis” uses data from July 2, 2007 through May 25, 2012. “Pre-crisis” uses data from January 2, 2004 through June 29, 2007.

generally quite high. At the monthly frequency, 10 of the 15 contracts feature a linear correlation above 0.85. Third, the Spearman rank correlations generally correspond quite closely to the linear correlations, suggesting that the linear correlations are not spuriously high due to outliers or non-normalities in volatility.

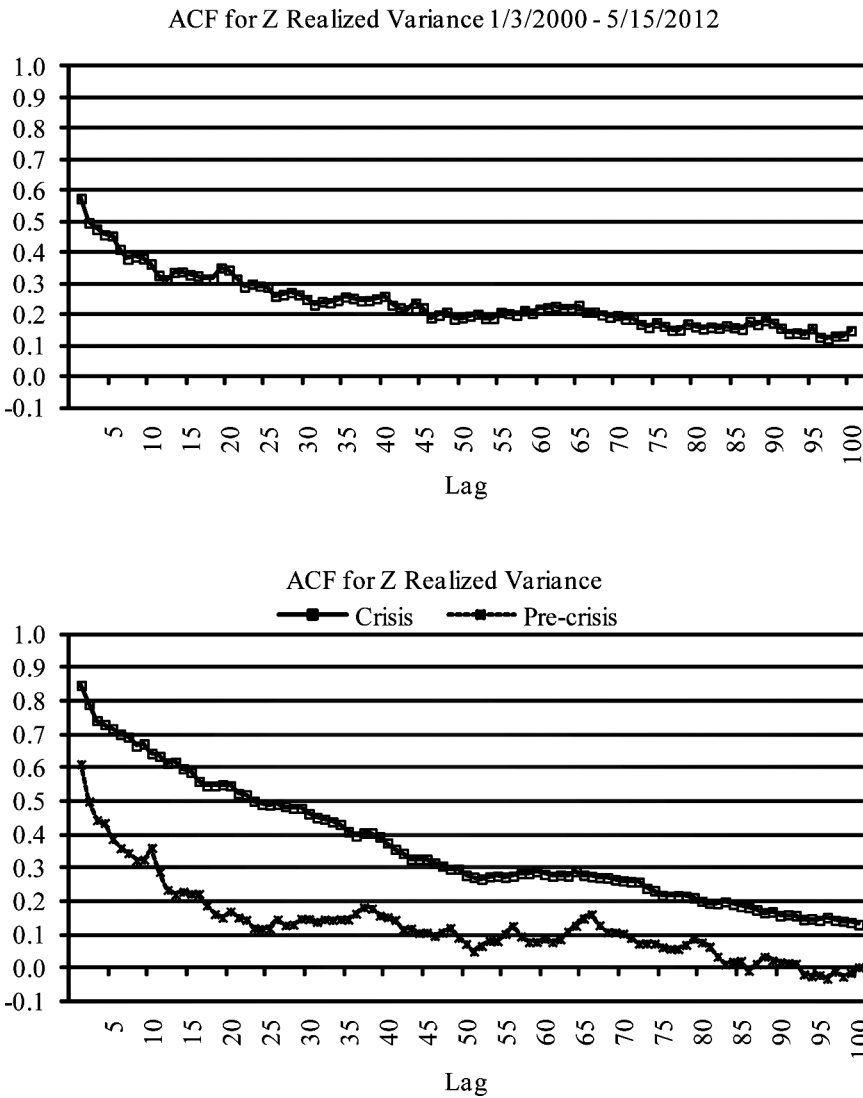


FIGURE 7

Autocorrelation functions of the FTSE contract.

Shown are autocorrelation functions of daily measures of realized variance based on 5-minute squared returns for the FTSE contract (Z) traded on NYSE Liffe. The bottom figure shows the functions estimated over two subsets. “Crisis” uses data from July 2, 2007 through May 25, 2012. “Pre-crisis” uses data from January 2, 2004 through June 29, 2007.

6. BENCHMARKING REALIZED VOLATILITY MOVEMENTS

As noted earlier, realized volatility changes when there is a change in the rate of information disseminating into the marketplace. We have documented significant changes over time in the level of volatility for all 15 futures contracts, reflecting the dramatic events of the financial crisis and the resulting uncertainty in markets around the world. Realized volatility is also affected by changes in market microstructure. Changes in market structure, including the increased dominance of electronic platforms and the rise in algorithmic trading, have

TABLE VI
Correlations Across Volatility Measures

<i>Ticker</i>	<i>Daily obs.</i>	<i>Correlation</i>	<i>1 day</i>	<i>5 days</i>	<i>21 days</i>
CL	3,114	Linear	0.7544	0.9199	0.9555
		Spearman	0.5121	0.6190	0.7332
ED	2,148	Linear	0.8522	0.9617	0.9867
		Spearman	0.5998	0.7836	0.8243
ES	3,118	Linear	0.3993	0.7346	0.8587
		Spearman	0.6522	0.8207	0.8548
TY	2,353	Linear	0.7883	0.9111	0.9700
		Spearman	0.6307	0.7639	0.8452
FDAX	2,555	Linear	0.7299	0.8835	0.8957
		Spearman	0.6451	0.6978	0.6719
FESX	2,558	Linear	0.6008	0.6842	0.6813
		Spearman	0.6151	0.6562	0.6187
FGBL	2,559	Linear	0.5544	0.7101	0.7311
		Spearman	0.5407	0.6649	0.6850
FGBM	2,559	Linear	0.5537	0.6659	0.6348
		Spearman	0.5384	0.6405	0.6810
B	1,877	Linear	0.7971	0.9515	0.9849
		Spearman	0.5423	0.7195	0.7349
SB	1,101	Linear	0.7405	0.9116	0.9533
		Spearman	0.5494	0.7283	0.8130
TF	1,065	Linear	0.8250	0.9618	0.9649
		Spearman	0.6056	0.6778	0.7355
I	3,002	Linear	0.1867	0.3800	0.5562
		Spearman	0.5480	0.7412	0.8103
L	2,093	Linear	0.1619	0.3137	0.4688
		Spearman	0.3931	0.5739	0.6176
R	3,036	Linear	0.7458	0.9029	0.9581
		Spearman	0.5740	0.7664	0.8238
Z	2,935	Linear	0.8415	0.9581	0.9853
		Spearman	0.6595	0.8105	0.8738

Note. Listed are correlations between realized volatility based on the sum of squared 5-minute returns and the Rogers and Satchell (1991) range-based estimate based on daily open, high, low, and closing prices. Volatility measures are computed over single-day periods as well as non-overlapping 5-day and 21-day periods. Linear correlations and Spearman rank correlations are listed. All Spearman correlations are statistically significant at the 1% level.

occurred contemporaneously with the increase in fundamental volatility in recent years, making inference difficult regarding the impact of changes in market structure on market quality.

To make this more concrete, note that observed trade prices are noisy due to market microstructure issues such as bid-ask price bounce, price discreteness (minimum tick size),¹⁰

¹⁰The bid-ask price bounce, for example, acknowledges that trade prices are likely to have occurred at the bid or the ask, depending on the motive for the trade. Indeed, Roll (1984) shows that the serial covariance of the sequence of trade prices can be used to infer the size of the bid-ask spread in an informationally efficient market.

and price impact, and this noise inflates the level of realized price (return) volatility in the following fashion:

$$\sigma_{\text{realized}} = \sigma_{\text{true}} + \sigma_{\text{microstructure}} \quad (6)$$

In the analyses conducted thus far in this study, we do not explicitly address the distinction between realized volatility and true volatility. In this section, we do.

The first component is “true” volatility or “macro-level” volatility. The second component is not related to fundamental economics and is a product of market microstructure. One way to mitigate the effects of microstructure volatility on realized volatility is to use bid-ask price midpoints throughout the day rather than trade prices. Hasbrouck and Saar (2011), for example, use the spread between high and low midpoint quotes over 10-minute windows. Unfortunately, this approach is not feasible since intraday bid-ask quote data were unavailable for our study. A potential alternative method is the Smith and Whaley (1994) generalized method of moments procedure. Using the sequence of trade prices, Smith and Whaley show how true volatility and microstructure volatility can be estimated simultaneously. Indeed, this estimation procedure was developed specifically for time and sales data from the futures exchanges in an era when historical bid-ask price quotes were not recorded. Unfortunately, this approach, too, was set aside because bid/ask spreads in many markets have become so small that the estimator arrives a corner solution.

6.1. Using Implied Volatility as a Benchmark

One way to distinguish between realized volatility and true volatility is to use option prices. Since volatility is a parameter in the option valuation formula, and all of the remaining parameters are known, we can equate the formula to the observed option price to infer the level of expected future volatility in the underlying asset market. This serves as our proxy for true volatility,¹¹ unfettered by microstructural considerations. Since we do not have access to futures option prices, we rely on published volatility indexes, of which we identified three: (a) the CBOE’s Volatility Index VIX, which provides an estimate of volatility for the CME’s E-mini S&P 500 futures contract; (b) the VDAX, which provides an estimate of the volatility for Eurex’s DAX futures contract; and (c) the VFTSE, which provides an estimate of the volatility of the FTSE stock index futures contract.

Figure 8 focuses on the comparison of realized volatility of the ES contract and the VIX index over the period January 3, 2000 through May 15, 2012. Since the VIX is a measure of annualized volatility in percentage points, we scale our measure of daily realized volatility appropriately. Figure 8A shows the individual daily estimates of realized volatility. Three features are apparent. First, in general the two series track each other extremely closely, and in fact have a linear correlation of almost 80%. Considering that the VIX is a forward-looking estimate of the following 30 days of volatility derived from index option prices, whereas the realized volatility is a backward-looking estimate derived from futures prices, this correlation is somewhat surprising. One interpretation is that market participants weigh heavily the intraday volatility of the ES contract in their assessment of fundamental volatility. Second, though the two series are highly correlated, the VIX tends to exceed the realized volatility quite dependably. The average level of the VIX is 22.2% over this period, for example, compared to 20.0% for realized volatility. The difference can be interpreted as a volatility risk premium

¹¹While it is true that option prices are also subject to microstructural effects just like futures, the effects can be mitigated by using bid/ask price midpoints and multiple option contract prices. Indeed, the CBOE uses hundreds of out-of-the-money S&P 500 options in its determination of the level of VIX.

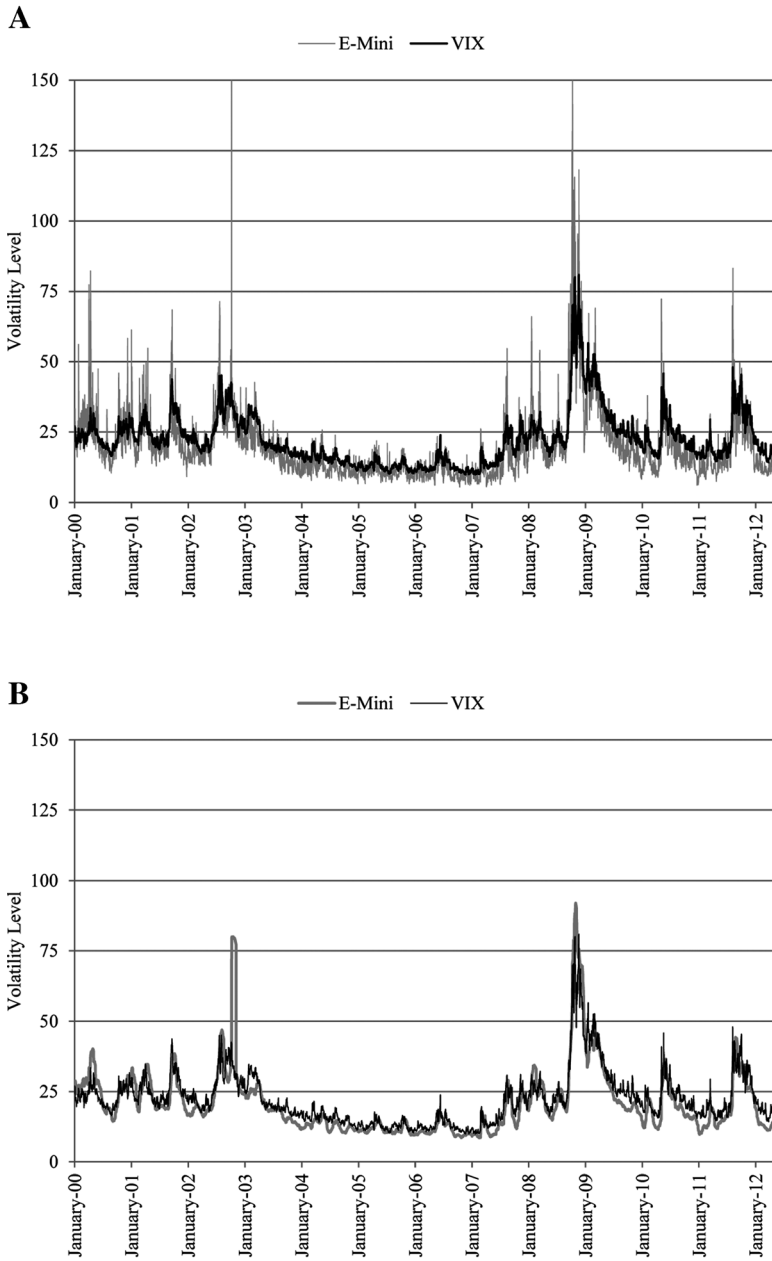


FIGURE 8

E-mini S&P 500 futures realized volatility versus VIX volatility index.

Shown are daily closing levels of “VIX,” the VIX volatility index, and “E-mini,” an annualized measure of realized volatility computed daily from 5-minute returns of the CME E-mini S&P 500 futures contract. Realized volatility is the square root of the sum of squared 5-minute returns using the last trade price within each 5-minute interval. Figures A and B show daily and monthly measures of realized volatility, respectively. The data run from January 3, 2000 through May 15, 2012.

incorporated into index option prices. Bollen and Whaley (2004) show that this premium is driven largely by the demand for stock portfolio insurance. An important attribute of the figure is, however, that the difference between implied and actual volatility appears to have increased in the latter part of the sample. In other words, realized volatility appears to have decreased relative to implied volatility even after the volatility risk premium is taken into account. Third, the spikes in the VIX are much smaller than the spikes in realized volatility. The reason is clear: as a forward-looking 30-day forecast of volatility, the VIX downplays the impact of volatility on any given day.

Figure 8B compares realized volatility to VIX by first averaging the current and past 20 observations of realized volatility before annualizing.¹² Here the spikes in realized volatility are generally equal in magnitude to the spikes in the VIX.¹³ The correlation between the two series is close to 90%. In our opinion, it is difficult to overstate the importance of this result. The VIX represents a benchmark for fundamental volatility that is free from microstructural effects in the underlying futures market. Figure 8A and B shows that our measure of realized volatility based on 10-minute returns tracks the VIX consistently from January 2000 through May 2012. If changes to market structure affected intraday volatility in a meaningful way, we would expect to see a divergence between realized volatility and the VIX after the changes were made. No such divergence is apparent.

In Figure 9, we compare the realized volatility of the FDAX contract to the VDAX volatility index. Figure 9A shows results using individual daily measures of realized volatility. The two series track each other quite closely, though not as closely as the E-mini realized volatility tracks the VIX. During the last 12 months of the sample, for example, the realized volatility consistently far exceeds the VDAX, and averages 51.8% versus 29.6% for the VDAX. One explanation for this phenomenon is the turmoil created by Germany's central role in maintaining financial order in the Euro zone. The FDAX market uncertainty was undoubtedly affected by events such as the credit downgrades in countries such as Ireland in April 2011 and Cyprus in September 2011 and the political upheaval arising from changes in governmental leadership in Ireland in February 2011, Portugal in June 2011, Spain in July 2011, Italy and Greece in November 2011, and France in May 2012. In Figure 9B, we see the same divergence between realized volatility in the FDAX and the level of volatility as measured by the VDAX. The correlation between the two series is close to 90%, just like the E-mini and the VIX, though again the divergence in the last 12 months is clear.

We compare the realized volatility of the Z contract to the VFTSE in Figure 10. Here the situation looks very similar to that of the ES contract and the VIX. A plausible macroeconomic explanation is that the United Kingdom financial market is less affected by trouble in the Euro zone than the Germany financial market.

6.2. Using Longer Horizon Volatility as a Benchmark

A second way to create a benchmark that abstracts from market microstructural effects is to compute volatility using returns of varying time horizons. Naturally, the measure of realized volatility from 5-minute returns can be significantly affected by microstructural effects including the bid-ask spread, high-frequency interactions between trading algorithms, and changes in liquidity. Bandi and Russell (2006, 2008) investigate decomposing variance estimated from high-frequency returns into two components: the variance due to efficient

¹²The 30-day horizon of VIX corresponds to roughly 21 trading days.

¹³The spike in realized variance in October 2002 can be traced to questionable prices late in the trading day on October 9.

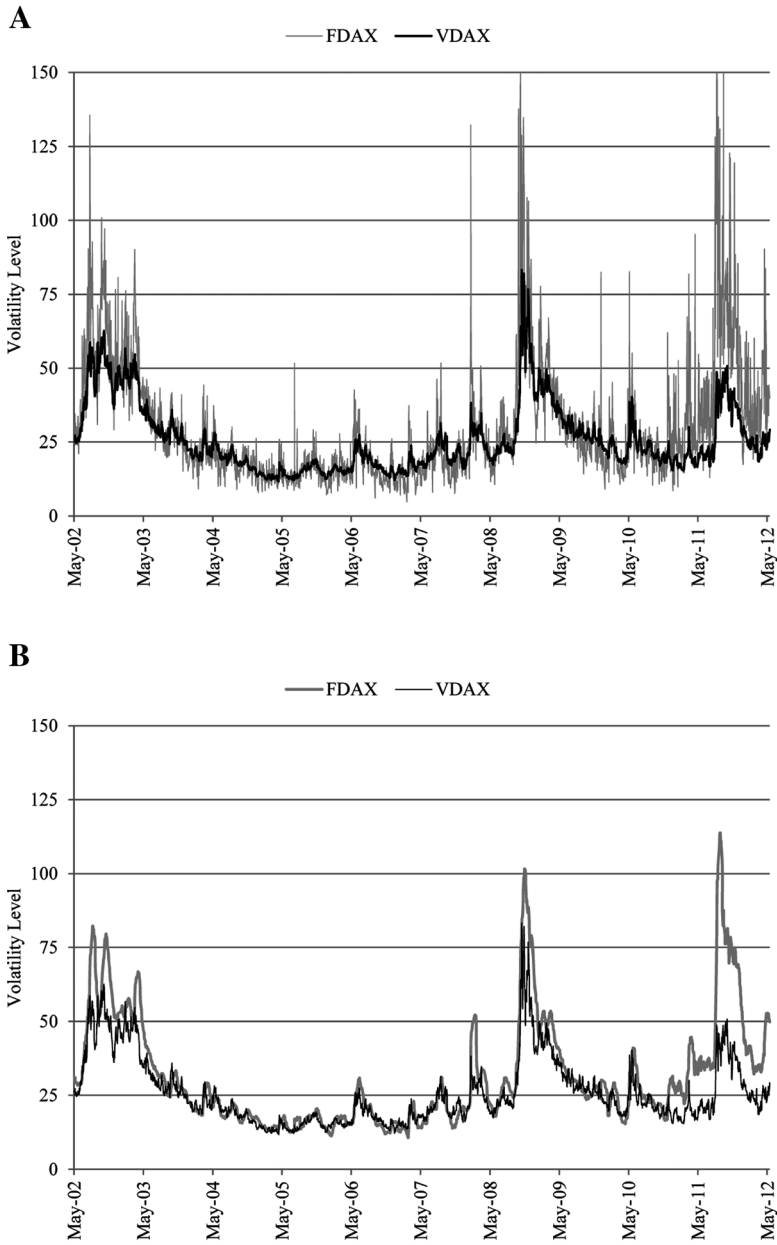


FIGURE 9

FDAX futures realized volatility versus VDAX volatility index.

Shown are daily closing levels of “VDAX,” the VDAX (new) volatility index, and “FDAX,” an annualized measure of realized volatility computed daily from 5-minute returns of the Eurex DAX futures contract. Realized volatility is the square root of the sum of squared 5-minute returns using the last trade price within each 5-minute interval. Figures A and B show daily and monthly measures of realized volatility, respectively. The data run from October 24, 2006 through May 15, 2012.




FIGURE 10

FTSE futures realized volatility versus VFTSE volatility index.

Shown are the daily closing levels of “VFTSE,” the VFTSE volatility index, and “FTSE,” an annualized measure of realized volatility computed daily from 5-minute returns of the NYSE Liffe FTSE 100 Index futures contract. Realized volatility is the square root of the sum of squared 5-minute returns using the last trade price within each 5-minute interval. Figures A and B show daily and monthly measures of realized volatility, respectively. The data run from January 4, 2000 through May 15, 2012.

returns that would prevail in a frictionless economy and the variance due to microstructure noise. The decomposition is achieved by sampling returns at different frequencies. Bandi and Russell show in a general specification in which the observed log price is the sum of the efficient log price and the log of a microstructure-induced noise term that as the observation window length decreases from 30- to 5-minute periods, the average squared observed return gradually increases to a consistent estimate of the second moment of microstructure noise. We therefore compare the annual volatility of each contract using 5-minute returns to the annual volatility using 10- and 30-minute returns to test for temporal changes in the impact of microstructure on market quality.

For each contract year, we construct an annual volatility by taking the square root of the sum of daily realized variance, in turn created from the sum of intraday squared 5-minute returns. These annual volatilities are standardized to a 252-trading day year. We then measure low-frequency volatility using returns from non-overlapping periods of 10 and 30 minutes within each day. Table VII shows the average percentage difference between the 5-minute volatility measure and the 10-minute (Panel A) or 30-minute returns (Panel B) for each contract. For B, SB, and TF the “Pre” year is 2008; for all other contracts it is 2004. For all contracts the “Post” year is 2011.

The volatility differences reported in Table VII are in all cases greater positive, which is consistent with Bandi and Russell (2006, 2008) since the realized 5-minute return volatility is more inflated by the impact of microstructure. Put differently, the “signal-to-noise ratio” (i.e., amount of true information about price change that you are extracting from the data relative to the amount of microstructural noise) is much greater for longer distancing intervals than for short ones. More importantly, in Panel A, only six of the contracts feature an increase from Pre to Post. Results in Panel B are qualitatively identical. For the four CME contracts, for example, CL and ES actually feature a significant reduction in noise, TY shows no change, and ED shows a significant increase. The other five contracts with an increase in noise are European (FDAX, FESX, FGBL, FGBM, and L). Increases in noise for these contracts are likely due to the continued liquidity problems caused by the Euro crisis rather than changes in trading practices.

7. CONCLUSIONS

The purpose of this study is to determine whether changes in trading practices have led to changes in futures market quality. We use excess futures contract return volatility as an encompassing measure of quality. While realized futures volatility is affected by changes in market microstructure, it is also affected by changes in the rate of information flow; hence, we must control for the dramatic events that have roiled financial markets in the past decade. Modeling true futures return volatility is no straightforward task. In general, it is very difficult to extract the microstructure component of realized volatility so that statements can be made about market quality.

In this study, we identify two benchmarks for fundamental volatility that permit direct tests for changes of the impact of market microstructure. The first is the use of implied volatility in equity index options markets. There is little evidence to suggest that the difference between implied and realized volatility has changed. In terms of extending this analysis, computing implied volatilities for markets that do not have published indexes, and investigating the differences between implied and realized volatilities would be worthwhile.

The second benchmark involves computing return volatility over different holding periods. When we compute volatility for different holding periods, we find that realized

TABLE VII
Volatility Ratios

Ticker	<i>Pre</i>		<i>Post</i>		<i>Change</i>	
	Average (%)	<i>P</i> -value	Average (%)	<i>P</i> -value	Δ Average (%)	<i>P</i> -value
Panel A. 5-minute/10-minute volatility – 1						
CL	9.74	0.0000	2.40	0.0000	–7.33	0.0000
ED	21.42	0.0000	28.63	0.0000	7.21	0.0000
ES	5.43	0.0000	2.58	0.0000	–2.85	0.0000
TY	4.72	0.0000	5.35	0.0000	0.63	0.4829
FDAX	3.30	0.0000	21.45	0.0000	18.15	0.0000
FESX	6.64	0.0000	34.64	0.0000	28.01	0.0000
FGBL	7.06	0.0000	37.63	0.0000	30.56	0.0000
FGBM	11.59	0.0000	36.75	0.0000	25.16	0.0000
B	1.44	0.0011	2.40	0.0000	0.96	0.1355
SB	5.46	0.0000	3.72	0.0000	–1.75	0.0345
TF	3.99	0.0000	2.83	0.0000	–1.16	0.1616
I	30.39	0.0000	23.13	0.0000	–7.26	0.0000
L	28.29	0.0000	33.27	0.0000	4.98	0.0527
R	4.30	0.0000	2.17	0.0001	–2.14	0.0532
Z	2.24	0.0000	2.65	0.0000	0.41	0.5409
Panel B. 5-minute/30-minute volatility – 1						
CL	23.79	0.0000	5.37	0.0000	–18.42	0.0000
ED	61.00	0.0000	86.05	0.0000	25.04	0.0000
ES	9.54	0.0000	6.33	0.0000	–3.21	0.0232
TY	12.84	0.0000	11.02	0.0000	–1.81	0.2909
FDAX	7.32	0.0000	57.99	0.0000	50.68	0.0000
FESX	13.12	0.0000	95.25	0.0000	82.13	0.0000
FGBL	15.39	0.0000	103.36	0.0000	87.97	0.0000
FGBM	23.51	0.0000	93.79	0.0000	70.28	0.0000
B	7.64	0.0000	5.70	0.0000	–1.94	0.1714
SB	12.31	0.0000	7.79	0.0000	–4.52	0.0127
TF	9.49	0.0000	5.18	0.0000	–4.31	0.0246
I	101.48	0.0000	60.14	0.0000	–41.33	0.0000
L	88.84	0.0000	101.90	0.0000	13.06	0.0433
R	9.66	0.0000	9.25	0.0000	–0.41	0.8333
Z	5.06	0.0000	7.30	0.0000	2.24	0.1575

Note. Listed for each contract is the average daily percentage difference between annualized volatility constructed from 5-minute squared returns and annualized volatility constructed from 10-minute (Panel A), and 30-minute (Panel B) returns for two separate years. For B, SB, and TF the “Pre” year is 2008; for all other contracts it is 2004. For all contracts the “Post” year is 2011. The *P*-values indicate whether the averages are statistically significantly different from zero. Also listed for each contract is the change across years. The *P*-values indicate whether the averages from the 2 years are statistically different from each other.

volatility for shorter periods is higher than for longer periods, thereby confirming the presence of microstructural effects. More importantly, the relative magnitudes have not increased meaningfully through time. In other words, there is scant evidence to suggest that the quality of electronically traded futures markets has changed through time.

REFERENCES

- Admati, A., & Pfleiderer, P. (1988). A theory of intraday patterns: Volume and price variability. *Review of Financial Studies*, 1, 3–40.
- Andersen, T. (1996). Return volatility and trading volume: An information flow interpretation of stochastic volatility. *Journal of Finance*, 51, 169–204.
- Andersen, T., Bollerslev, T., Diebold, F., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics*, 61, 43–76.
- Andersen, T., Bollerslev, T., Diebold, F., & Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, 96, 42–55.
- Ball, C., Torous, W., & Tschoegl, A. (1985). The degree of price resolution: The case of the gold market. *Journal of Futures Markets*, 5, 29–43.
- Bandi, F., & Russell, J. (2006). Separating microstructure noise from volatility. *Journal of Financial Economics*, 79, 655–692.
- Bandi, F., & Russell, J. (2008). Microstructure noise, realized variance, and optimal sampling. *Review of Economic Studies*, 75, 339–369.
- Bollen, N., & Christie, W. (2009). Market microstructure of the Pink Sheets. *Journal of Banking and Finance*, 33, 1326–1339.
- Bollen, N., & Whaley, R. E. (2004). Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance*, 59, 711–754.
- Broggard, J. (2012). High frequency trading and volatility. Working paper, Northwestern University.
- Brown, S., Laux, P., & Schacter, B. (1991). On the existence of optimal tick size. *Review of Futures Markets*, 10, 50–72.
- Chordia, T., Roll, R., & Subrahmanyam, A. (2001). Market liquidity and trading activity. *Journal of Finance*, 56, 501–530.
- Clark, P. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, 41, 135–155.
- Epps, T., & Epps, M. (1976). The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis. *Econometrica*, 44, 305–321.
- Figlewski, S. (1989). Options arbitrage in imperfect markets. *Journal of Finance*, 44, 1289–1311.
- Foster, F. D., & Vishwanathan, S. (1990). A theory of intraday variations in volumes, variances and trading costs in securities markets. *Review of Financial Studies*, 3, 593–624.
- Garleanu, N., Pedersen, L. H., & Poteshman, A. M. (2009). Demand-based option pricing. *Review of Financial Studies*, 22, 4259–4299.
- Glosten, L., & Milgrom, P. (1985). Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14, 71–100.
- Harris, L. (1991). Stock price clustering and discreteness. *Review of Financial Studies*, 4, 389–415.
- Hasbrouck, J., & Saar, G. (2011). Low-latency trading. Working paper, New York University.
- Heimstra, C., & Jones, J. (1994). Testing for linear and nonlinear Granger causality in the stock price-volume relation. *Journal of Finance*, 49, 1639–1664.
- Hendershott, T., Jones, C., & Menkveld, A. (2011). Does algorithmic trading improve liquidity? *Journal of Finance*, 66, 1–33.
- Jarrow, R., & Protter, P. (2012). A dysfunctional role of high frequency trading in electronic markets. *International Journal of Theoretical and Applied Finance*, 15, 1250022-1–1250022-15.
- Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica*, 53, 1315–1336.
- Lamoureux, C., & Lastrapes, W. (1994). Endogenous trading volume and momentum in stock-return volatility. *Journal of Business and Economic Statistics*, 12, 253–260.
- Lee, E. (2015). High frequency trading in the Korean index futures market. *Journal of Futures Markets*, 35, 31–51.
- Li, J., & Wu, C. (2006). Daily return volatility, bid-ask spreads, and information flow: Analyzing the information content of volume. *Journal of Business*, 79, 2697–2739.
- Madhavan, A. (1995). Consolidation, fragmentation, and the disclosure of trading information. *Review of Financial Studies*, 8, 579–603.
- Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of Business*, 51, 61–65.
- Richardson, M., & Smith, T. (1994). A direct test of the mixture of distributions hypothesis: Measuring the daily flow of information. *Journal of Financial and Quantitative Analysis*, 29, 101–116.
- Rogers, L., & Satchell, S. (1991). Estimating variance from high, low and closing prices. *The Annals of Applied Probability*, 1, 504–512.
- Roll, R. (1984). A simple implicit measure of the bid/ask spread in an efficient market. *Journal of Finance*, 39, 1127–1139.

- Smith, T., & Whaley, R. E. (1994). Estimating the effective spread from times and sales data. *Journal of Futures Markets*, 14, 437–455.
- Tauchen, G., & Pitts, M. (1983). The price variability-volume relationship on speculative markets. *Econometrica*, 51, 485–505.
- Thomakos, D., & Wang, T. (2003). Realized volatility in the futures markets. *Journal of Empirical Finance*, 10, 321–353.
- Whaley, R. E. (2009). Understanding the VIX. *Journal of Portfolio Management*, 35, 98–105.
- Yang, D., & Zhang, Q. (2000). Drift-independent volatility estimation based on high, low, open, and close prices. *The Journal of Business*, 73, 477–492.