

Tail Wags Dog: Intraday Price Discovery in VIX Markets

Nicolas P.B. Bollen, Michael J. O'Neill , and Robert E. Whaley*

Beginning with VIX futures in 2004, followed by VIX options in 2006 and VIX ETPs in 2009, the daily open interest in volatility contracts is now in the tens of billions of dollars. Given this growth, it is important to develop a better understanding of price discovery and the supply/demand dynamics in each market. Some of the price relations are linked by arbitrage. Others are not. In particular, the relation between the VIX cash index and the VIX futures is not arbitrated, and we show that, where once VIX changes led VIX futures price changes, the VIX futures now leads. © 2016 Wiley Periodicals, Inc. *Jrl Fut Mark* 37:431–451, 2017

1. INTRODUCTION

More than twenty years ago, Whaley (1993) argued that the marketplace would benefit from the introduction of derivative contracts written on the newly created CBOE Market Volatility Index (VIX). The primary illustration centered on a stock index option market maker who, in the process of fulfilling his duty to transact, at times accumulates a sizable option position and a corresponding exposure to volatility risk. Whaley showed how volatility derivatives could be used to hedge this risk, driving down the market maker's inventory holding costs and ultimately transaction costs for index option end users.

Over the next ten years, the VIX grew from an esoteric measure of stock market risk to a widely disseminated and popularly followed gauge of investor fear regarding a downward jump in the stock market. The interest in the VIX motivated the CBOE to introduce VIX futures contracts in March 2004. Trading activity was meager until the launch of VIX options in February 2006. VIX calls, in particular, were an immediate success, presumably because they are a low cost mechanism for providing tail-risk insurance. The presence of an active options market, as it turns out, spurred trading activity in the VIX futures market. This is natural in the sense that VIX option market makers, who are generally short calls from supplying insurance, need to vega-hedge their inventories and the VIX futures are a convenient, low-cost

Nicolas P.B. Bollen is Frank K. Houston Professor of Finance in Vanderbilt University, Nashville, Tennessee. Michael J. O'Neill is Adjunct Associate Professor in Bond University, Queensland, Australia. Robert E. Whaley is Valere Blair Potter Professor of Finance and Director of the Financial Markets Research Center in Vanderbilt University, Nashville, Tennessee. Seminar participants at the Auckland University of Technology, the College of William and Mary, Southern Methodist University, and the University of Texas at Dallas provided helpful suggestions, as did participants at Vanderbilt's May 2015 Options and Volatility Conference. Discussions with Henry Chien, Bernard Dumas, Gary Gastineau, Joanne Hill, Neil Ramsey, Jacob Sagi, Tom Smith, Bill Speth, Hans Stoll, George Tauchen and Damon Walvoord and the research assistance of Daejin Kim are gratefully acknowledged. This research was supported by the Financial Markets Research Center at Vanderbilt University.

*Correspondence author, The Owen Graduate School of Management, Vanderbilt University, 401 21st Avenue South, Nashville, TN 37203. Tel: 615-343-7747, Fax: 615-376-8214, e-mail: whaley@vanderbilt.edu

Received December 2015; Accepted July 2016

mechanism for doing so. The next meaningful intervention in VIX trading activity occurred in January 2009 with the launch of VIX exchange traded products (ETPs). The introduction of VIX ETPs was motivated, at least in part, by un-satiated investor demand for volatility exposure. Investors such as retail customers, too small or unsophisticated to trade in the futures market, and institutions such as pension funds and endowments, barred from trading in futures market, were able to trade market volatility for the first time.

Given the proliferation of VIX-related instruments, and the access they provide to new classes of investors, a natural question to ask is whether volatility's "price discovery" has been affected. The purpose of this paper is to determine where information about volatility first appears using intraday price movements in four related markets: S&P 500 index options, VIX futures, VIX options, and VIX ETPs.

Prior literature studies the relation between the VIX, computed from S&P 500 index option prices, and VIX futures prices. Zhang and Zhu (2006) develop a futures pricing function based on Heston's (1993) stochastic volatility model. They find that the large deviations between model values and market prices are dramatically reduced when model parameters are re-estimated annually, suggesting one or more model constants actually time varies. Zhang, Shu, and Brenner (2010) present a similar model but specify a separate diffusion for the long-term mean of the volatility process to accommodate the time variation documented by Zhang and Zhu. Shu and Zhang (2012), who also study the lead/lag relation between the cash VIX and VIX futures prices, find the two time-series are largely contemporaneously related. Given the speed of today's markets, however, daily changes may not provide the granularity necessary to determine where traders trade first.

Our paper makes two main contributions. First, we use a more up-to-date sample, with data through April 2013, and incorporate a more comprehensive set of VIX products, including the most actively traded VIX ETP. These data provide additional insights regarding the evolution of trading in the related markets. Second, we use intraday trade and quote data to study lead/lag relations between the different VIX-related time-series. Frijns, Tourani-Rad, and Webb (2016), in an important departure from prior work, also use intraday data in their study of lead/lag relations between the VIX and VIX futures. Our study differs from theirs in that we use all trade prices, as opposed to regularly spaced bid-ask midpoints, to identify the transmission of information. In principle, new information enters the marketplace when trades occur. Consequently, we also use different econometric techniques that are designed explicitly to accommodate high frequency microstructure. We find that the lead/lag relations between the prices of VIX futures and VIX options, and between the VIX ETPs and the VIX futures, are short-lived as they are largely governed by arbitrage price relations. In contrast, the deviations between VIX futures and the VIX cash index cannot be arbitrated due to the high costs of trading the index options that comprise the cash VIX. One of the most interesting results in the paper stems from this market friction: while the VIX futures price lagged the VIX cash index in the first few years after it was launched, the VIX futures now leads the cash index. Put differently, the VIX futures has become the "go-to" market for hedging volatility risk in a timely manner.

The paper is organized into five sections. The first section describes the daily and intraday data used in our analyses. Section II uses the daily data to sketch the landscape of the volatility market during our sample period. Section III reviews the methodology we use for assessing the intraday lead/lag relation between VIX and VIX products. The methodology is relatively new in the sense that it does not involve sampling prices at fixed intervals during the day. Instead, all trade prices changes are used whenever the time period between trades of two prices series has overlap. Such a methodology is better able to determine the true nature of the lead/lag relation. The main results of the paper are presented in Section IV. Section V contains a brief summary of the study and our conclusions.

2. DESCRIPTION OF VIX PRODUCTS AND DATA SOURCES

We study five different VIX-related price series: (i) the VIX index itself, created from S&P 500 index option prices; (ii) the VIX futures written on the VIX index; (iii) the VIX options written on the VIX futures; (iv) one of the S&P 500 VIX futures indexes; and (v) one of the VIX ETPs benchmarked to S&P 500 VIX futures indexes. For each series, we describe the data and their sources. We use both daily and intraday time-series for all five series.

The CBOE Volatility Index (VIX) is based on the bid/ask price quotes of nearby and second nearby out-of-the-money and at-the-money S&P 500 index options and can be interpreted as a real-time market estimate of the expected stock market volatility over the next 30 calendar days.¹ When the CBOE introduced the VIX in 1993, its stated purpose was not only to provide a market consensus estimate of short-term stock market volatility but also to provide a “standard” upon which futures and options contracts on market volatility could be written. While the financial press and investment community were quick to adopt VIX as the “investor fear gauge,” the launch of exchange-traded derivatives on VIX took more than a decade to develop. VIX futures contracts were launched by the CBOE Futures Exchange (CFE) on March 26, 2004, and VIX options were launched by the CBOE on February 23, 2006.

In the analyses that follow, we compare the time-series of VIX levels with prices of VIX derivatives and VIX ETPs. Occasionally, we will refer to the VIX as “the VIX cash index” in order to distinguish it from VIX futures, VIX options, VIX futures indexes, and VIX ETPs. The daily closing levels of VIX were obtained from the CBOE website. Although the data are available from early 1993, we use only the data beginning with the VIX futures launch on March 26, 2004. The sample ends on April 30, 2013. The intraday data on VIX were obtained directly from the CBOE.

The VIX futures contract is written on the VIX cash index and has a denomination of \$1,000 times the index level.² Daily open-high-low-close-settlement VIX futures prices, trading volume, and open interest were obtained from the CBOE Futures Exchange (CFE) website. Trade and quote data were acquired from the CBOE’s Market Data Express. The VIX option contract is written on the VIX futures (not the cash VIX index) and has a multiplier of \$100, only one-tenth the size of the futures. Daily data for the VIX options were acquired from the WRDS OptionMetrics data files. Each record in the file provides daily market information for an option series. An option series is defined by the triplet (i) call-put; (ii) exercise price; and (iii) expiration date. The market data include the highest closing bid price, the lowest closing ask price, the number of shares traded during the day, and the end-of-day open interest. Intraday trade and quote data were acquired from the CBOE’s Market Data Express.

Standard and Poor’s computes four constant maturity VIX futures indexes. Their series date back to December 20, 2005. Two indexes have 30 days to expiration—the S&P 500 VIX short-term total return index (SPVXSTR) and the VIX short-term excess return index. The S&P 500 VIX short-term (ST) futures indexes, for example, track a strategy of holding long

¹The VIX uses nearby and second nearby index options with at least 8 days left to expiration and to imply a constant, 30-calendar day measure of the expected stock market volatility. When it was introduced in 1993, it was based on S&P 100 index option prices since they accounted for 75% of the index option trading volume. It was also based on at-the-money options as they were the most liquid. The methodology is provided in Whaley (1993). By the late-1990s, the S&P 500 index options took over as the highest trading volume contract. In addition, a new weighting scheme was devised to account for the prices of all available nearby and second nearby options out-of-the-money and at-the-money contracts. The current methodology for calculating VIX is provided on the CBOE’s website at <http://cboe.com/micro/vix/introduction.aspx>.

²See the appendix for details for contract details for VIX futures and options.

positions in the nearby and second nearby futures in proportions that create an average time to maturity of 1 month at the close of trading. The positions are held for 1 day recognizing any gains or losses, and then rebalanced or rolled to move the weighted average maturity from 29 days back to 30 (i.e., selling some of the nearby contracts and buying more of the second nearby contracts). The only difference between total return and excess return indexes is that the total return index includes the return on a short-term Treasury bill (as should be the case on a collateralized futures position). The S&P 500 VIX mid-term (MT) total return and excess return futures indexes are similarly rebalanced daily so as to maintain a 5-month constant maturity using the fourth, fifth, sixth, and seventh futures contracts.³ Our analyses use only the SPVXSTR.

The daily futures index level data since inception were obtained from Bloomberg. The intraday data dating back to the inception date were not readily available, requiring us to create the intraday SPVXSTR price series. To create the weights applied to each contract each day to create a futures index with a constant 30 days to expiration, we used the methodology described on Standard and Poor's website (www.us.spindices.com). The calculation involves taking the prices of the two nearest maturity VIX futures contracts and replicating a position that rolls the nearest month VIX futures to the next month on a daily basis in equal fractional amounts, with the roll and weighting files produced at the end of each day. In generating the intraday index levels, we mimicked the procedure used for the VIX index. Rather than recomputing the index whenever one of the options within the index trades, the CBOE computes the index every 15 seconds. We set the intraday SPVXSTR computations at the same frequency.

VIX exchange-traded products (ETPs) were launched on January 29, 2009. Selected attributes of the nine largest VIX ETPs traded in the United States are presented in Table I. The most active VIX ETP is VXX, which was one of the first two VIX ETPs launched and has enjoyed a first-mover advantage. VXX traded more than 35.4 million shares a day during the 3 months leading up to April 30, 2013 and had a market capitalization exceeding \$1.18 billion on April 30, 2013. VXX is benchmarked to the SPVXSTR futures index. VXX has a multiplier or leverage factor (L) of 1, which means that the ETP performance is benchmarked to one times the daily index return (less management fees and expenses). The second most active VIX ETP is XIV, which is benchmarked to the S&P 500 VIX ST ER futures index and has a multiplier of -1 . The third is UVXY, which is also benchmarked to the S&P 500 VIX ST ER futures index, but promises twice the daily futures index return. Two of the nine VIX ETPs are benchmarked to MT futures indexes, but account for only 4.1% of the \$2.7 billion in market capitalization. Of the \$2.7 billion pegged to ST indexes, 71.7% are direct investments (i.e., $L = 1$ and $L = 2$) and 28.3% are inverse investments (i.e., $L = -1$). Daily data for the VIX ETPs benchmarked to the VIX short-term indexes were obtained from Bloomberg. Intraday data for VXX, the largest and most active VIX ETP, were obtained from TAQ.

3. VOLATILITY LANDSCAPE DURING THE SAMPLE PERIOD

This section uses daily data for the various VIX-related time-series to provide some insights regarding the behavior of volatility during the sample period, March 26, 2004 through April 30, 2013. We document the events that drove extraordinary VIX movements during the period, show the incredible growth of the VIX product markets, and provide a sense for their importance on both absolute and relative scales.

³A detailed explanation of the methodology used to compute each of the indexes is provided on Standard and Poor's website www.us.spindices.com.

TABLE I
Attributes of the Nine Largest Direct and Inverse VIX ETPs Traded in the United States as of April 30, 2013

Symbol	Name	Date of Inception	Average Volume	Net Asset Value (\$)	Price per Share (\$)	Benchmark			Expense Ratio (%)
						ST/MT	TR/ER	Multiplier	
VXX	iPath S&P 500 VIX Short-Term Futures ETN	20090129	35,353,434	1,179,535,630	18.97	ST	TR	1	0.89
XIV	VelocityShares Daily Inverse VIX Short-Term ETN	20101129	9,752,718	654,115,600	22.90	ST	ER	-1	1.35
UVXY	ProShares Ultra VIX Short-Term Futures ETF	20111004	22,409,833	296,780,400	6.30	ST	ER	2	0.95
VIXY	ProShares VIX Short-Term Futures ETF	20111003	1,325,182	192,902,500	10.22	ST	ER	1	0.85
TVIX	VelocityShares Daily 2x VIX Short-Term ETN	20101129	4,265,345	163,072,000	2.80	ST	ER	2	1.65
SVXY	ProShares Short VIX Short-Term Futures ETF	20111004	503,856	76,483,000	89.98	ST	ER	-1	0.95
VIXM	ProShares VIX Mid-Term Futures ETF	20111003	84,311	55,851,250	24.55	MT	ER	1	0.85
VXZ	iPath S&P 500 VIX Mid-Term Futures ETN	20090129	265,752	53,313,260	20.02	MT	TR	1	0.89
VIIIX	VelocityShares VIX Short-Term ETN	20101129	682,198	16,128,000	10.24	ST	ER	1	0.89
Total				2,688,181,640					

The source of the ETP data is <http://efdb.com/etfdb-category/volatility/>. Average volume is average daily trading volume over the past 3 months. NAV and price per share are closing levels on April 30, 2013. The benchmarks are denoted by ST and MT, which represent the S&P 500 VIX short- and mid-term futures indexes, respectively. The notation TR and ER denotes the total return and the excess return versions of the indexes. The multiplier of the ETP is applied to the daily return of the futures index. A multiplier of 1 implies that the ETP promises the daily rate of return on the underlying futures index. A multiplier of 2 implies that the ETP promises two times the return of the index, and -1 implies that the ETP promises minus the daily return of the index. Expense ratio is annual management fee and is charged on a daily basis.

3.1. Behavior of VIX

The growth in the VIX products markets has been driven, in part, by the demand to hedge tail-risk, defined as losses generated by an unexpectedly large drop in the stock market. As noted by Black (1976a), equity volatility tends to spike when the stock market falls, hence a positive exposure to VIX can offset these losses. To illustrate the behavior of VIX during the sample period, consider Figure 1 which shows the daily closing level of VIX during the sample period. Two characteristics of VIX are immediately apparent. First, it is mean-reverting. Second, and more importantly perhaps, it is prone to periodic spikes when there is unexpected geopolitical or macroeconomic news.⁴ On October 17, 2008, for example, the VIX spiked to a level of 70.33 based on disappointing economic statistics. On October 27 of the same year, the VIX rose further to 80.08 as the United States Treasury injected \$125 billion into nine major U.S. banks. On November 20, the VIX hit 80.86 as the S&P 500 dropped to an 11.5 year low following signs of economic contraction. The increasingly popularity of VIX products results, at least in part, from the fact that it is an effective means of providing tail-risk insurance.

3.2. VIX Derivatives Trading Activity

VIX derivatives trading activity has had four distinct phases since the launch of VIX futures on March 26, 2004. Consider Figure 2. Phase 1 began with the launch of VIX futures. With no VIX options, the average daily trading volume of the VIX futures was only about 500 contracts. Phase 2 began with the introduction of VIX options on February 24, 2006. VIX options trading volume quickly exceeded VIX futures trading volume.⁵ VIX futures volume also increased dramatically during the second phase, with average daily trading volume rising to nearly 3,400 contracts, a 568% increase over the first phase. Phase 3 began on January 29, 2009 with the launch of VIX ETPs. Both futures and options trading volume increased sharply, with average daily VIX futures volume during the third phase in excess of 53,000 contracts, a 645% increase over the second phase. Finally, in Phase 4, VIX futures trading activity took on a life of its own. With no commensurate increase in trading activity of the VIX options or the VIX ETPs, VIX futures volume grew to over 111,000 contracts per day, a 339% increase over Phase 3, and now, reportedly, rivals SPX variance swap liquidity.⁶

The increase in trading activity for both futures and options had at least two contributing factors. First, the marketplace was finally afforded the opportunity to buy tail-risk insurance. The demand for tail-risk insurance grew quickly, with VIX option volume surpassing VIX futures volume almost immediately in Phase 2. Second, interrelations between markets generate trading volume. During Phase 2, for example, the simultaneous presence of the two complementary futures and options markets provides the opportunity for market makers in both markets to hedge their inventories. With competitive markets, bid/ask spreads are reduced, thereby promoting increased trading activity.⁷ Similarly, in Phase 3, issuers of VIX ETPs were required to hedge the short volatility exposure, and did so primarily

⁴The prominence of these spikes has led to theoretical work modeling jumps in volatility, with a recent application to VIX in Todorov and Tauchen (2011).

⁵In this figure, daily options trading volume is divided by 10 to account for the difference in the denomination of VIX futures (\$1,000 times the index) and VIX options (\$100 times the index).

⁶See Morgan Stanley *QDS Vega Times* March 28, 2011, page 6 and Mason and Onur (2014).

⁷We computed the average relative spread across all VIX futures price quotes for each of the four nearby contracts in each of the four phases. Its level was 0.983% in Phase 1, 0.896% in Phase 2, 0.369% in Phase 3, and 0.294% in Phase 4. Not surprisingly, these levels correspond to the trading volume numbers shown in Figure 2.

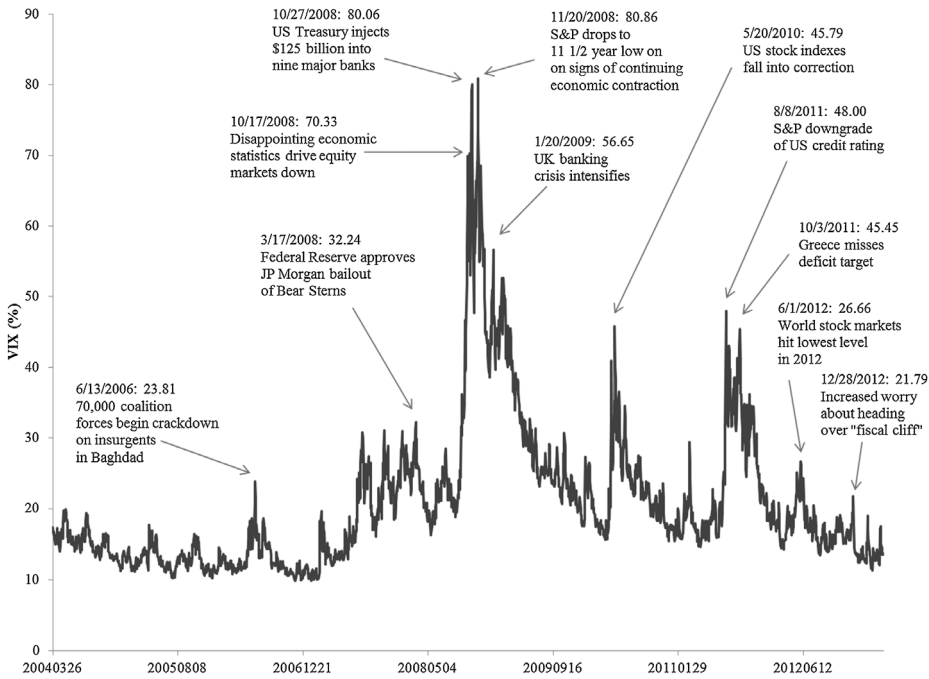


FIGURE 1
 Daily Closing Levels of VIX During the Sample Period March 2004 Through April 2013
 [Color figure can be viewed at wileyonlinelibrary.com]

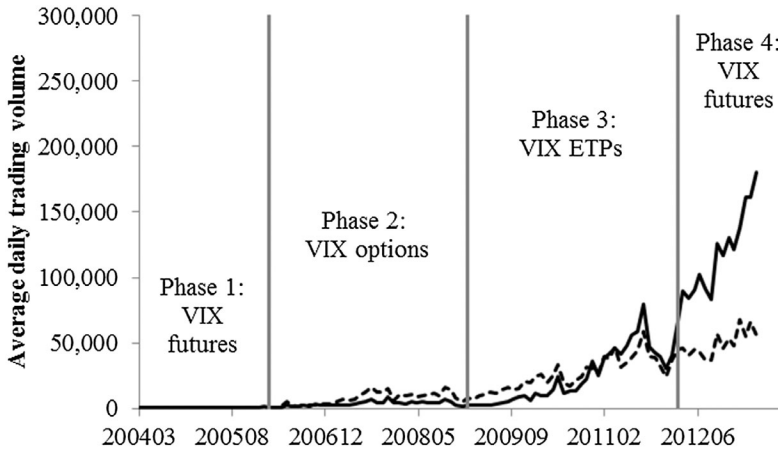


FIGURE 2
 Average Daily Trading Volume by Month for VIX Futures and Option Contracts During the
 Period March 2004 Through April 2013

Options volume is divided by a factor of 10 to account for the difference in the contract denomination of the futures (\$1,000 times index level) and the options (\$100 times index level). Futures volume is represented by the solid black line, and options volume by the dotted line. Phase 1 begins on March 26, 2004, when VIX futures were launched. VIX options were not yet traded. Phase 2 begins on February 24, 2006, when VIX options were launched. VIX ETFs were not yet traded. Phase 3 begins on January 29, 2009, the launch date of the first VIX ETPs. Phase 4 has no new product introduction. It begins February 29, 2012, when VIX futures trading raced ahead of VIX options trading. [Color figure can be viewed at wileyonlinelibrary.com]

in the VIX futures market. Hedging activity resulted in VIX futures trading volume surging ahead of VIX option volume by late 2010.

3.3. VIX ETP Trading Activity

VIX ETPs were launched in January 2009. The VIX ETPs are most like the VIX futures in the sense that they provide a direct long or short position in volatility. A distinction, however, is that the VIX ETPs are benchmarked to a constant maturity VIX futures index while the time to expiration of a VIX futures contract falls each day. There is a certain complementarity about the products, however. The VIX ETPs can be (and in some cases are) constructed from the VIX futures, hence the price relation is a no-arbitrage condition. Consequently, one would expect to see a strong contemporaneous relation between the trading activities in the two products. Figure 3 shows this to be the case. As the dollar market value of the VIX ETPs grew in early 2009, the dollar market value of the VIX futures grew in a parallel fashion. Indeed, for a short time in 2010, the dollar market value of the VIX ETPs surpassed that of the futures. This is possible because the price exposure of VIX ETPs can also be mimicked using VIX options. At one point in 2010, VIX ETPs had nearly 20% higher market value than the VIX futures. As time has passed, however, the market value of VIX ETPs flattened out while the VIX futures grew.

4. INTRADAY RETURN ANALYSIS TEST METHODOLOGY

This paper examines lead/lag relations between the VIX cash index (VIX) and financial contracts based on the VIX. In the finance literature, such empirical investigations have

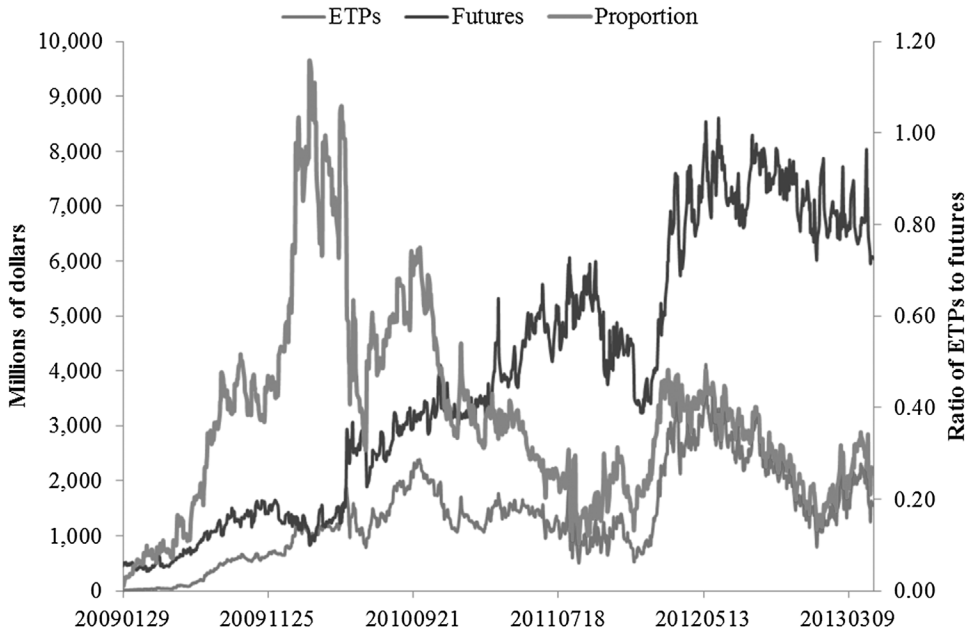


FIGURE 3

Dollar Market Value of VIX ETPs Linked to Short-Term S&P 500 VIX Futures Index and Dollar Value of Open Interest of VIX Futures During the Period January 29, 2009 Through April 30, 2013 [Color figure can be viewed at wileyonlinelibrary.com]

focused on cases in which the relation between the cash and derivatives markets follows a cost of carry relation such as

$$F = Se^{rT} \quad (1)$$

where F is the price of a futures contract with time to expiration T , S is the underlying cash price, and r is the risk-free interest rate. What ensures this relation holds is arbitrage. If the futures price exceeds its theoretical value by more than trading costs, risk-free profits can be earned by selling the futures and buying the underlying asset. If the futures price is below its theoretical level, risk-free profits can be earned by buying the futures and selling the underlying asset.

This paper represents a departure from past work because we focus primarily on a market where arbitrage between the futures and cash market is not possible in a cost-effective manner. Consequently, the cost of carry relation (1) does not hold and the nature of the relation depends on the supply and demand for futures contracts at different expiration dates. Whaley (2013), for example, documents the fact that the VIX futures markets is in contango about 80% of the time, meaning that long hedging demands for volatility outstrip short hedging demands, likely the result of the well-known inverse relation between volatility and equity returns. Whether the futures prices lead the cash prices, however, remains an open question.

The wide-ranging variety of VIX-related instruments shows that market participants are offered a choice regarding how to trade and gain exposure to market volatility. If one market is favored due to higher liquidity and lower trading costs, we might expect new information about volatility to be impounded there first. Furthermore, the markets are linked to varying degrees: those linked by arbitrage should reflect new information relatively synchronously, whereas others may do so at a delay. We review below our methodology for assessing resulting lead/lag relations among the different time-series.

Most studies of lead/lag relations between security returns rely on cross-correlation functions using regularly spaced time intervals. Stoll and Whaley (1990), for example, examine the relation between S&P 500 futures returns and the returns of the underlying index using the last trade price in 5-minute time intervals. De Jong and Nijman (1997) examine the same relation and present results using the last trade price in 10-, 5- and 1-minute intervals. Stephan and Whaley (1990) examine the relation between price changes of stock and stock options using the last trade price in 5-minute intervals and conclude that the stock market leads the option market by as much as 15 minutes. Chan, Chung, and Johnson (1993) confirm the Stephan and Whaley findings, but also show that the price changes are contemporaneous when bid/ask quote midpoints are used. Lo and MacKinlay (1990) study the relation in the returns of twenty size-sorted stock portfolios using daily, weekly, and monthly data. Hou, Barbaris, and Chen (2007) examine the relation between weekly returns calculated from Wednesday close prices, using size-ranked portfolios of NYSE, AMEX, and NASDAQ firms as well as portfolios grouped by industry classification.

Regularly spaced time intervals have also been used in analyses of lead/lag relations in other markets. Pan et al. (1996) investigate lead/lag relations between the currency spot market and the currency option market, using transaction data and dividing each trading day into eight 45-minute time intervals. Danielsson and Payne (2002) assess lead/lag relations for spot foreign exchange markets, between indicative and firm quotes using best limit bid and offer prices outstanding at the end of each interval, and transactions data using last trade price. Data are drawn from Reuters' FX broking systems and are sampled in time intervals of 5 minutes, 10 minutes, and 20 seconds. Crain and Lee (1995) assess whether there is an intraday lead-lag relation between spot and futures volatility for interest rate and foreign

exchange markets on macroeconomic announcement days, using tick data to prepare volatility data in hourly intervals. Granger, Huang, and Yang (2000) analyze causality relations between stock index prices and exchange rates using end of day close prices for nine Asian exchanges. Kaen and Hachey (1983) apply Granger causality tests to selected series of U.S. and U.K. currency denominated money market rates, using weekly and monthly intervals between observations. An alternative approach using irregularly spaced data has been to explicitly insert missing observations for lead/lag analyses and this has also been applied to currency markets (Covrig & Melvin, 2002; De Jong, Mahieub, & Schotman, 1998).

Applying this regular-interval sampling technique to high frequency data, however, can be problematic. Lead/lag relations that occur at higher frequencies than the sampling interval data cannot be assessed, and, when estimators of cross-correlation between return series are applied, spurious conclusions about lead/lag relations may result. Non-synchronous trading and its impact on correlation measures was first discussed by Scholes and Williams (1977). Epps (1979) showed that relatively high frequency observations can distort correlation measures. As the sampling frequency increases, covariance estimates tend to zero as a consequence of non-synchronous trading effects. As high frequency trading has increased in intensity, these issues have become more important.

We address non-synchronous trading in our analysis of intraday lead/lag relations between the VIX cash index, VIX futures prices, VIX option prices, and VIX ETP prices using a methodology developed by Hayashi and Yoshida (2005). Hayashi and Yoshida (HY) discuss the problem of non-synchronous trading in terms of the errors that can arise when randomly spaced observations like trade arrivals are sampled at regular time frequencies. The typical approach, as noted above, is to use only the last trade price at 1- or 5-minute intervals. In contrast, the HY cross-correlation estimator makes use of all data, regardless of the time intervals between samples, and thereby deals with the issues of non-synchronous trading and spurious correlations without a need to rely on sampling.⁸

The intuition underlying the HY methodology for estimating return correlation for two securities is straightforward. Rather than base returns on prices recorded over arbitrary windows, they record a return at every trade. A pair of returns enters the correlation estimate whenever the time period between trades has any overlap. When one series is lagged, the HY estimator can be used to measure the cross-correlations between the two series. Huth and Abergel (2014) and Hoffmann, Rosenbaum, and Yoshida (2013) use the lagged version of the HY estimator to devise a lead/lag ratio where the lags may be dependent and assess whether one variable leads or lags another. Based on Hoffmann et al., we define the lead/lag ratio (LLR) as

$$LLR = \frac{R_{Y,X-l}^2}{R_{X,Y-l}^2}, \quad (2)$$

where $R_{Y,X-l}^2$ is the R^2 from a regression of Y on lagged X , and $R_{X,Y-l}^2$ is the R^2 from a regression of X on lagged Y . If $R_{Y,X-l}^2 > R_{X,Y-l}^2$ or, equivalently, $LLR > 1$, the returns of X lead those of Y . If $R_{Y,X-l}^2 < R_{X,Y-l}^2$ or, equivalently, $LLR < 1$, the returns of Y lead those of X . We dub this the “ R^2 criterion.” This approach is consistent with the work of Fung, Lau, and Tse (2015) who take the sum of the squares of partial correlation coefficients in order to determine the

⁸Pomponio and Abergel (2013) acknowledge the standard approach of measuring lagged cross-correlation of returns to study asymmetry between positive and negative lags. They present a different but related approach which looks at trade-through events to assess which asset moves first. Christensen, Kinnebrock, and Podolskij (2010) also present a modulated realized covariance estimator which deals with non-synchronous trading, as well as using pre-averaging to reduce the impact of microstructure noise.

numerator and denominator in (2). We estimate the population R^2 of the projection of Y on lagged X using the general case

$$R_{Y,X-l}^2 = (C^{YX_l})^T (C^{X_l X_l})^{-1} (C^{YX_l}) \quad (3)$$

where C^{UV} denotes the matrix of population covariances between a matrix U and a matrix V . The Hayashi–Yoshida estimator is used to estimate all covariances in the Equation (3). The lead/lag relation is then determined using (2).

The next stage of the analysis is calculating associated confidence intervals for the lead/lag relations. Bibinger and Mykland (2013) provide confidence intervals for the Hayashi–Yoshida covariance matrix. There is no literature which expresses confidence bounds for the significance of lead/lag correlations, however, and it is far from obvious how to translate confidence intervals for the covariance matrix into confidence intervals for cross-correlation function and the *LLR*.

An alternative approach is to bootstrap the HY covariance matrix under the null using the empirical distribution. This allows the construction of confidence intervals that are customized to the specific empirical setup used in estimating *LLR*, for example, sample size and number of lags considered. A crude rule of thumb for bootstrapping appears to be that one needs between 25 and 200 replicates for a standard error and ten times this (up to max of 2,000) for a confidence interval.⁹ The time-series data are bootstrapped directly to create 2,000 resamples for the futures/cash analysis per day. The difficulty in performing the bootstrap is in preserving the correlation structure. In this paper, we overcome this by resampling in blocks. Each trading day is broken into 25 blocks of 20 minutes each and 2,000 replicates are drawn randomly each day with replacement. The dependence would be lost at the junctions, but 20 minutes would seem a reasonable block size in that end-effects would be small relative to information contained within each block. These bootstrapped time-series can then be used to calculate 2,000 estimates for the covariance function and the *LLR* per day confidence intervals for the *LLR*.

The HY covariance matrix can also be used to produce correlation functions for lead/lag relations which can be used to assess the magnitude and duration of lead/lag relations between two series. Huth and Abergel use cross-correlation functions to assess the duration of lead/lag relations between 33 stocks and 4 market indexes. They examine, for example, the lead/lag relation between the returns of the French equity index CAC 40 and the French stock Renault. They assess lags from -300 to $+300$ seconds finding that lags between pairs of assets generally do not persist meaningfully beyond ± 5 minutes on a given day. Hoffmann et al. (2013) use the contrast function defined by covariances to reach similar conclusions regarding the durations of lead/lag relations between the futures contract on the DAX index and the Euro-Bund future contract both with maturity December 2010, using all the trades for 20 days of October 2010.

5. INTRADAY RETURN ANALYSIS OF VIX AND VIX PRODUCTS

We now turn to applying the correlation mechanics to the returns of VIX and VIX products. Each day, we calculate log return series for the VIX cash index, the VIX ST futures index, nearest maturity VIX option-implied forward, and nearest maturity VIX futures. We then take pairs of return series and apply time shifts of various lengths to one of the two return series. From these series, we construct pairs of observations that share an overlap. From these pairs,

⁹See Stata guidelines for bootstrap samples.

we calculate the population R^2 for a projection of the returns of one series on lagged returns on the other series. At a given time shift, X leads Y if the R^2 of the regression of Y on lagged X is greater than the R^2 of the regression of X on lagged Y . In summarizing our results, we report the percentage of trading days for which X leads Y to assess the overall relation. Each day we also construct the cross-correlation functions for leads and lags of log returns to visually appraise how long the lags persist.

The nature of the lead/lag relations between the level of VIX and the prices of its products depends critically on the degree to which cost-effective arbitrage between pairs of prices is possible. The price series of the nearby VIX futures and nearby VIX options, for example, are actively arbitrated. In this case, we expect price series to be tightly linked. Similarly, the VXX ETP is benchmarked to the S&P 500 VIX ST futures index (SPVXSTR). Like VIX, it, too, has a constant 30 days to expiration; however, it involves daily rebalancing of only the nearby and second nearby VIX futures contracts each so as to maintain the constant maturity. Trading costs are modest. Since the nearby and second nearby VIX futures have active and liquid markets, we can expect the VXX and SPVXSTR price series to also be tightly linked. In contrast, arbitrage between the VIX and the VIX futures index is simply not possible since the VIX is not a tradable security from a practical standpoint. It consists of hundreds of nearby and second nearby S&P 500 index options, which must be partially rolled each day so as to maintain a constant 30 days to expiration. Trading costs would be prohibitive. Without active arbitrage, the lead/lag relation between the price series is likely to show significant lags, one way or another.¹⁰ We now investigate each of these relations in turn.

5.1. Analysis of VIX Futures Versus VIX Options

The relations examined in this subsection and the next are like the relations examined in past studies of lead/lag relations between adjacent markets in which arbitrage can be easily and inexpensively executed. Indeed, in many cases, it can be accomplished without human intervention through algorithmic trading.

Specifically, we now turn to the relation between the returns of the VIX futures and the VIX options. To compare apples to apples, we need to transform the prices of the call and the put on the VIX futures to an implied VIX futures (or forward) price. Since the options are European-style, the put-call parity relation arising from costless arbitrage between the futures and the options sharing the same time to expiration is

$$c - p = e^{-rT}(F - X), \quad (4)$$

where the call and the put share the same exercise price X . What this means is that the VIX futures price implied by the call and put prices is

$$F = X + (c - p)e^{rT}. \quad (5)$$

In using (5) to compute the implied futures prices, we are faced with the problem that there are multiple expirations for the futures and the options and multiple exercise prices for the options at different expirations. Since using all call/put price pairs is redundant, we must decide which specific contracts to use in generating the VIX futures and options return series. Unreported calculations show that the nearby futures and options contracts are the most

¹⁰In addition, even if it was tradable, the cash VIX and the VIX futures index are not perfect substitutes. One measure the expected future market volatility over the next 30 days, and the other measures the expected 30-day volatility 30 days from now.

TABLE II
 Summary of Results Lead/Lag Analysis of the VIX Futures with the VIX Options During
 the Period February 24, 2006 Through April 30, 2013

<i>Phase</i>	<i>No. of Futures</i>	<i>No. of Options</i>	<i>R² Criterion (%)</i>	<i>P < 0.05 (Futures Lead) (%)</i>	<i>P > 0.95 (Options Lead) (%)</i>
2	310	5,704	87	10	5
3	2,745	29,761	85	12	4
4	7,514	21,263	54	31	1

The sample period includes only the last three phases of the full sample period because VIX options did not trade in the first phase. Phase 2 begins on February 24, 2006, when VIX options were launched. Phase 3 begins on January 29, 2009, the launch date of the first VIX ETPs. Phase 4 begins February 29, 2012. The average number of records used each day are reported in the second and third columns. The fourth column reports the proportion of days within the phase that the R^2 criterion concludes that VIX options lead VIX futures index. The fifth and sixth columns report the proportion of days where the P -values of the bootstrapped lead/lag ratios are in the left or right tail of the distribution, respectively. For the bootstrap, each trading day is broken into 25 blocks of 20 minutes each and 2,000 replicates are drawn randomly with replacement. A P -value < 0.05 for a given day, for example, indicates that futures leads options with a probability of 95%, while a P -value > 0.95 indicates that options leads futures with probability 95%.

actively traded. Consequently, we use the nearby expiration for the futures and options return series. Similarly, the at-the-money options sharing the same exercise price are the most actively traded. Thus, for the purposes of calculating option implied forward prices, the most actively traded call and put options being the closest to the money at the beginning of each day are used. The implied forward prices for the day are updated with each call or put quote using the bid/ask price midpoints. This is in contrast to the approach of Barndorff-Nielsen and Shephard (2005) who calculate the forward price only after a new observation is available for both the call and put and time is reset.¹¹

Competing arguments can be offered regarding whether the VIX futures or the VIX options, as represented by the implied VIX futures price, will lead. On the one hand, the trading volumes shown in Figure 2 suggest that VIX futures should lead VIX options. The activity in the VIX futures market now dominates the options, indicating that VIX options now simply tag along. Microstructure theory suggests that informed traders will access the deepest market first, as market depth provides camouflage which enables the informed trader to maximize his trading profit, as in Kyle (1985). Furthermore, trading costs tend to be lower for deeper markets because trading volume allows dealers to amortize their costs over more transactions. Consequently, models of the bid/ask spread often include the inverse of trading volume as a determinant, as in Bollen, Smith, and Whaley (2004). In addition, the implied VIX futures price series is slowed by the fact that the VIX call trading activity dominates VIX puts as a result of demand for tail-risk insurance. In other words, the price change in the implied VIX futures price series may be slowed by the inactivity of VIX puts. On the other hand, VIX calls are a popular and inexpensive form of tail-risk insurance, and institutions quickly flock to this type of insurance when markets become jittery about bad economic or geopolitical news. Buying the futures would expose these same institutions to downside risk in the event that they were wrong in their predictions and volatility dropped. In summary, the direction of the lead/lag relation between the VIX futures and VIX options is uncertain. The length of the lead, however, should be short in light of the fact that these are actively arbitrated markets.

¹¹An alternative approach would have been to use the Black (1976b) call option-implied VIX futures price. This imposes a specific model on the relation between VIX futures and option prices, which we wanted to avoid. Our main thesis is that the lead/lag relation depends critically on the presence or absence of arbitrage.

The HY procedure is used to evaluate the lead/lag relation between the actual and implied VIX futures price series. The HY covariance matrix is calculated daily and is used to derive the R^2 of the returns on VIX option implied futures prices with lagged returns on VIX futures and vice versa. Covariances are calculated using lags of between 0 and 60 units. The differential lags considered here are in increments of 1 second given our expectation that lead/lag relations are likely to have short durations.

Table II summarizes the proportion of days in each phase during the sample period where options leads futures, according to the R^2 criterion (column 4). In this instance, no Phase 1 results appear since VIX options were yet to be launched. In Phase 2, options (“implied futures”) lead futures in 87% of the days using the R^2 criterion. For Phases 3 and 4, the VIX options retains a clear lead, in 85% of the days of Phase 3 and 54% of the days in Phase 4.

Table II also contains the bootstrap results. The table summarizes the P -values of the lead/lag relations between the two series. The P -values are calculated as the point in the distribution function where the LLR crosses one, using 2,000 replicates each day. The numbers in the table are the proportion of days in each phase with a specified P -value for the relation with one series lagged against the other. In Phase 2, the P -values are less than or

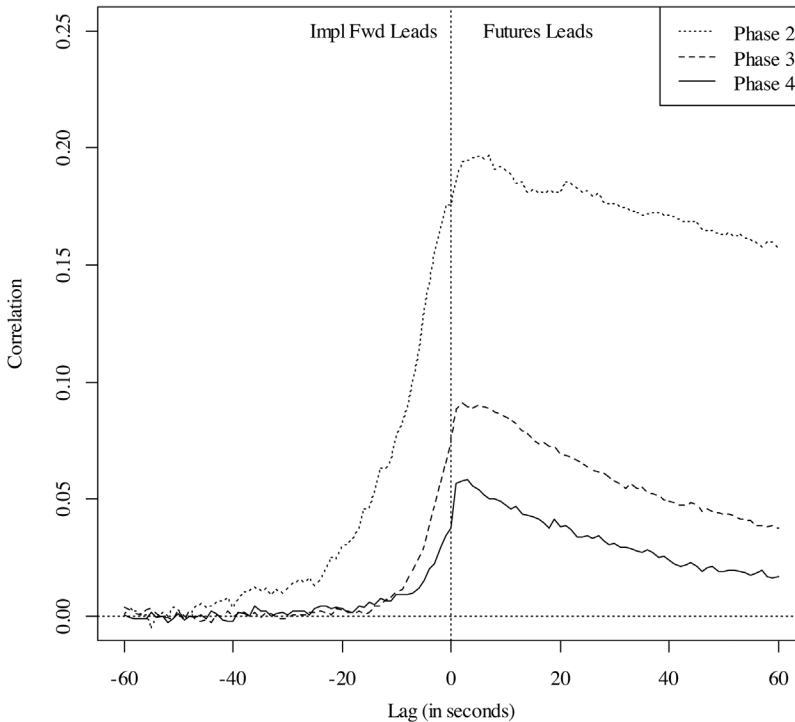


FIGURE 4

Cross-Correlation Function Between Nearest Maturity Option-Implied Forward and Nearest Maturity VIX Futures During the Sample Period February 24, 2006 Through April 30, 2013

Figure summarizes the cross-correlation function of returns of the nearest maturity option-implied forward prices (“Impl Fwd”) against leads and lags in returns of the nearest maturity VIX futures (“Futures”). Correlations are calculated using the Hayashi and Yoshida (2005) covariance matrix, which is generated each day after time-shifting returns of the futures with positive and negative lags in increments of 1 second from -60 to $+60$ seconds. The corresponding vector of correlations has 121 elements. The median correlation at each lag is shown for each of three phases during the sample period. Phase 2 begins February 24, 2006, Phase 3 begins January 29, 2009, and Phase 4 begins February 29, 2012. Phase 1 results do not appear because VIX options were not yet traded.

equal to 0.05 in 10% of days studied. This means that the lead/lag relation with futures leading options is significant at the 95% confidence level for 10% of days in Phase 2. In the same phase, options leads the futures with 95% confidence in 5% of days. The results are similar for Phase 3 with futures leading options in 12% of days at the 95% confidence level, and options leading futures in 4% of days. There is a material change in the lead/lag relation in Phase 4 that begins on February 29, 2012 and ends on April 30, 2013, where futures leads options in 31% of the days at 95% confidence, while the cash leads the futures in only 1% of days. From these results, we conclude that futures lead options in all three phases, although the result is much stronger in Phase 4, where VIX futures trading activity takes on a life of its own.

Figure 4 shows the cross-correlation function in the different phases of the sample period. The asymmetry of the cross-correlation function also seems to support the case that futures more frequently leads options, with any lead/lag relations diminishing beyond 40 seconds in Phase 4. The average duration of the lag in futures with maximum correlation occurs at 19 second in Phase 2, 13 second in Phase 3, and 11 seconds, in Phase 4.

5.2. Analysis of VXX Versus VIX Short-Term Futures Index

We now focus on the relation between the VIX ETPs and VIX futures markets. Specifically, we use iPath's VXX to represent the VIX ETP market since it has the highest market capitalization and greatest trading volume. Its benchmark is the SPVXSTR index, whose construction was described earlier. Again this is an arbitrage relation, so the expected lead/lag relation should be short lived. In light of this, we recalculate the SPVXSTR at a higher frequency to analyze this relation; rather than reset in 15-second intervals we reset continuously at every futures trade.

The lead/lag relation between VXX and SPVXSTR is unlikely to be stationary, however. Recall Figure 3. Where initially the hedging activity of ETPs took the lion's share of the futures open interest, it is now only a small sliver (even smaller for trading volume which is not shown in the figure). In addition, the demand for VIX ETPs has subsided, as shown by the market value of the shares outstanding in the figure. Presumably this is as a result of investors becoming more informed about the nature of the product and its inherent "contango trap."

To assess the time-series relation between the VIX ETPs and the constant maturity futures index, we again apply the HY methodology. The covariance matrix is calculated daily

TABLE III
Summary of Results Lead/Lag Analysis of the VXX with the VIX ST Futures Index
During the Period January 29, 2009 Through April 30, 2013

<i>Phase</i>	<i>No. of Futures</i>	<i>No. of ETPs</i>	<i>R² Criterion (%)</i>	<i>P < 0.05 (Futures Lead) (%)</i>	<i>P > 0.95 (ETP Leads) (%)</i>
3	4,747	76,377	75	6	36
4	13,146	113,202	59	15	23

The sample period includes only the last two phases of the full sample period because VIX ETPs did not trade until the third phase. Phase 3 begins on January 29, 2009, and Phase 4 begins February 29, 2012. The average number of records used each day are reported in the second and third columns. The fourth column reports the proportion of days within the phase that the R^2 criterion concludes that VXX leads the VIX ST futures index. The fifth and sixth columns report the proportion of days where the P -values of the bootstrapped lead/lag ratios are in the left or right tail of the distribution, respectively. For the bootstrap, each trading day is broken into 25 blocks of 20 minutes each and 2,000 replicates are drawn randomly with replacement. A P -value < 0.05 for a given day, for example, indicates that futures leads VXX with a probability of 95%, while a P -value > 0.95 indicates that VXX leads futures with probability 95%.

using lags of between 0 and 60 seconds. The differential lags considered here are again in increments of 1 second up to 1 minute, since any lead/lag relation should be short-lived.

Table III summarizes the results of the VXX/SPVXSTR lead/lag investigation. According to the R^2 criterion, VXX leads the futures index in 75% of the days in Phase 3, and in 59% of the days in Phase 4. These results suggest that the direction of causation is clear, VXX leads futures. We depend on the bootstrap results, also reported in Table III to draw formal conclusions. In Phase 3, the P -values are greater than or equal to 0.95 in 36% of days studied. This means that the lead/lag relation with VXX leading futures is significant at the 95% confidence level for 36% of days in Phase 3. In the same phase, futures leads VXX with 95% confidence in only 6% of days. In Phase 4, the results are weaker, however, with VXX leading futures in 23% of days at the 95% confidence level, and futures leading VXX in 15% of

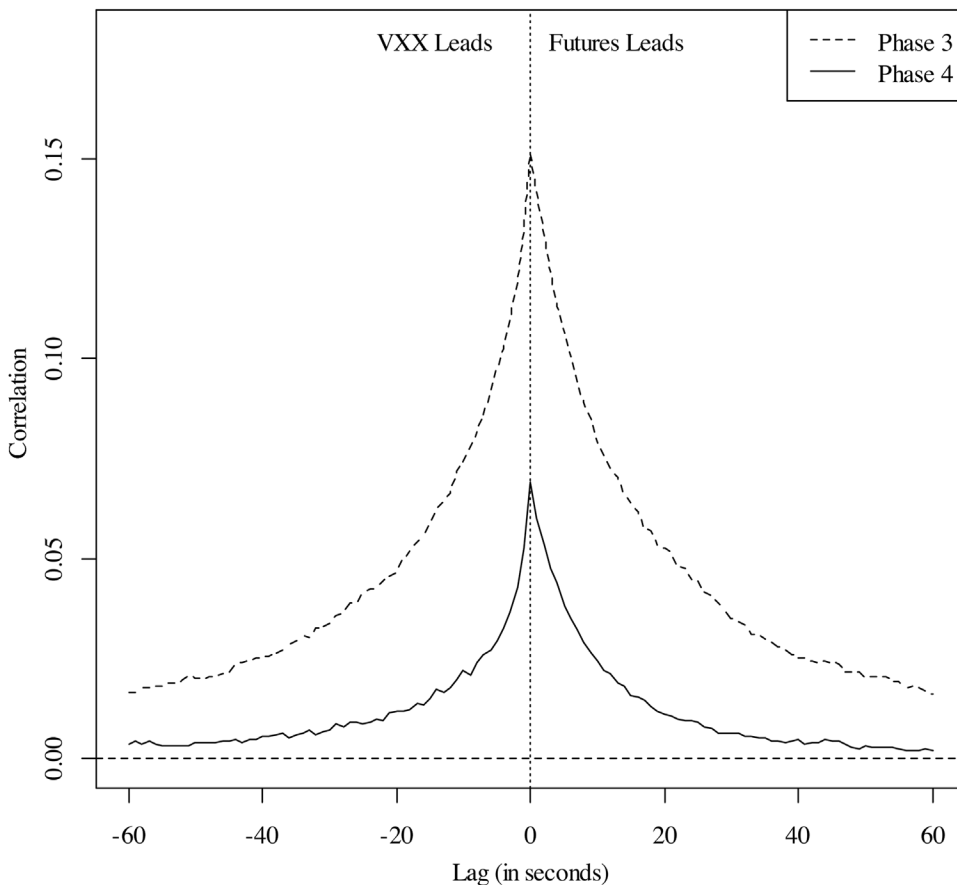


FIGURE 5

Cross-Correlation Function Between VXX and VIX ST Futures Index During the Sample Period
January 30, 2009 Through April 30, 2013

Figure summarizes the cross-correlation function of log returns of the VXX against leads and lags in returns of the VIX ST futures index (“Futures”). Correlations are calculated using the Hayashi and Yoshida (2005) covariance matrix, which is generated each day after time-shifting returns of the futures with positive and negative lags in increments of 1 second from -60 to $+60$ seconds. The corresponding vector of correlations has 121 elements. The median correlation at each lag is shown for each of two phases during the sample period. Phase 3 begins January 29, 2009, and Phase 4 begins February 29, 2012. Phase 1 and 2 results do not appear because VXX was not yet traded.

days. This coincides with significant increases in futures trading in Phase 4. From these results, we conclude that VXX leads the VIX futures index.

Figure 5 shows the cross-correlation function between the VXX and SPVXSTR returns. The function is symmetrical, showing that the price movements are largely contemporaneous. The lead/lag correlations do not persist much beyond 25 seconds. On average the maximum correlation of lead/lags occurs at ± 3 seconds in Phase 3 and ± 2 seconds in Phase 4.

5.3. Analysis of VIX Cash Index Versus VIX Short-Term Futures Index

In many ways, the time-series analysis of the relation between the VIX and SPVXSTR is the most telling. The theoretical arguments laid out earlier suggest that the markets will not be closely linked, although we have not argued whether the VIX futures leads the VIX cash or vice versa. An important clue in developing such an argument, however, is shown in the VIX futures trading volumes of Figure 2. When VIX futures were initially launched, they appeared to be a non-starter. Daily trading volume in the first few years was meager. During this period, it is highly unlikely that VIX futures price movements had any influence on the movements in the VIX index. The VIX, which is based on the prices of a highly active and liquid S&P 500 index options market, in all likelihood drove the movements of the VIX futures. In recent years, however, VIX futures trading volume and open interest have soared to unprecedented levels. They now rival the SPX options market in the volatility (or “vega”) risk exposures that the contracts create. The aim is to assess the time-series relation between the VIX futures and VIX cash index using high frequency data and determine whether the relation has changed over time.

We now turn to the relation between VIX and SPVXSTR returns each day. The investigation period begins December 20, 2005, when Standard and Poor’s began reporting S&P 500 VIX futures indexes. Each day the return data are aligned at every overlapping observation, with all observations considered regardless of arrival time. Time shifts of between 0 and 60 units of time are used. The differential lags considered here are in increments of 15 seconds as dictated by the return series. First the population R^2 of the

TABLE IV
Summary of Results Lead/Lag Analysis of the VIX with the VIX ST Futures Index
During the Period December 20, 2005 Through April 30, 2013

<i>Phase</i>	<i>No. of Futures</i>	<i>No. of Cash Levels</i>	R^2 <i>Criterion (%)</i>	$P < 0.05$ <i>(Futures Lead) (%)</i>	$P > 0.95$ <i>(Cash Leads) (%)</i>
1	1,612	1,612	46	11	7
2	1,627	1,627	68	2	1
3	1,635	1,635	37	26	0
4	1,608	1,608	25	31	0

The sample period is divided into four phases. Phase 1 begins on March 26, 2004, when VIX futures were launched. VIX options were not yet traded. Phase 2 begins on February 24, 2006, when VIX options were launched. VIX ETFs were not yet traded. Phase 3 begins on January 29, 2009, the launch date of the first VIX ETPs. Phase 4 has no new product introduction. It begins February 29, 2012, when VIX futures trading raced ahead of VIX options trading. Within each phase, the average number of records used each day are reported in the second and third columns. The fourth column reports the proportion of days within the phase that the R^2 criterion concludes that VIX leads the VIX futures index. The fifth and sixth columns report the proportion of days where the P -values of the bootstrapped lead/lag ratios are in the left or right tail of the distribution, respectively. For the bootstrap, each trading day is broken into 25 blocks of 20 minutes each and 2,000 replicates are drawn randomly with replacement. A P -value < 0.05 for a given day, for example, indicates that futures leads cash with a probability of 95%, while a P -value > 0.95 indicate that cash leads futures with probability 95%.

projection of VIX futures index returns on the lagged returns of the VIX cash index level, $R_{F,C-l}^2$, is calculated as described in Equation 3, along with the R^2 of the projection of VIX cash index returns on the lagged returns of the VIX futures index, $R_{C,F-l}^2$. The results of the R^2 criterion are reported in column 4 of Table IV, summarized in terms of the number of days in each period for which the futures index with lagged cash, $R_{F,C-l}^2$, exceeds that for cash with lagged futures index, $R_{C,F-l}^2$. If so, then VIX futures leads VIX.

Table IV reveals a number of interesting results. During Phase 1 the number of observations is few since the Phase 1 period begins March 26, 2004 but SPVXSTR data are not available until December 20, 2005. In this phase, the lead/lag relation between the cash and the futures was weak, likely as a result of the sparse trading volume in the VIX futures market. During Phase 2, when the VIX options were launched but before VIX ETPs were

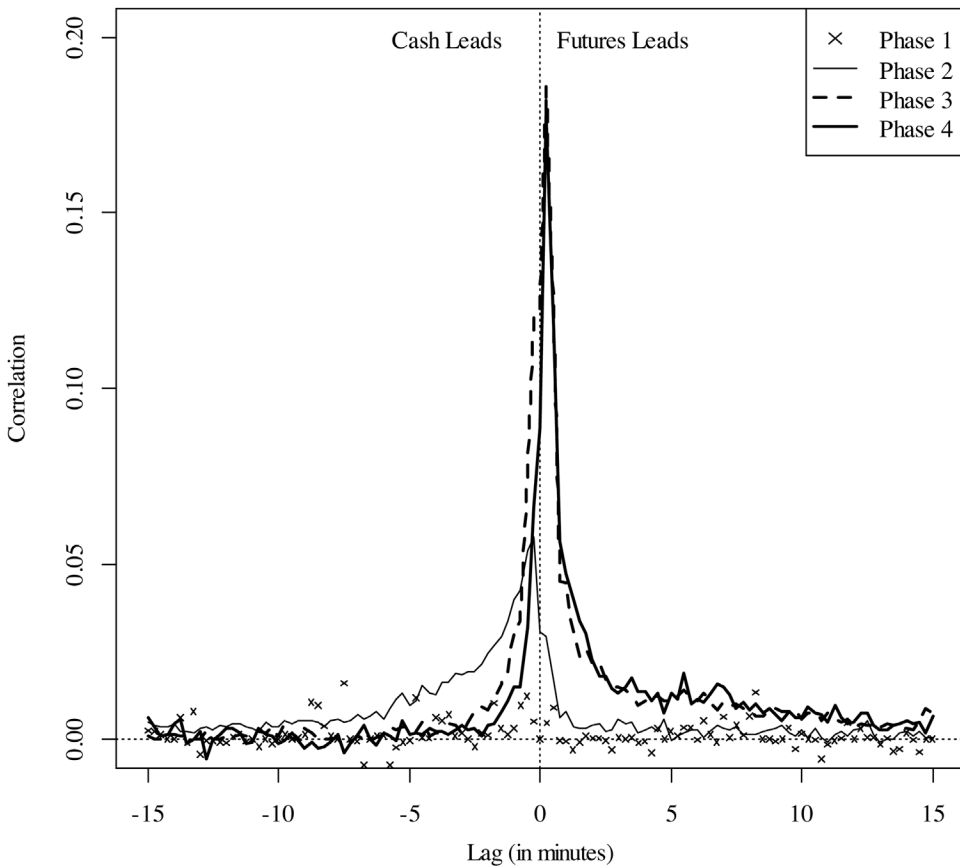


FIGURE 6

Cross-Correlation Function Between VIX Cash Index and VIX ST Futures Index During the Sample Period December 20, 2005 Through April 30, 2013

Figure summarizes the cross-correlation function of returns of the VIX cash index (“Cash”) against leads and lags in returns of the VIX ST futures index (“Futures”). Correlations are calculated using the Hayashi and Yoshida (2005) covariance matrix, which is generated each day after time-shifting the returns of the futures with positive and negative lags in increments of 15 seconds from -15 to $+15$ minutes. The corresponding vector of correlations has 121 elements. The median correlation at each lag is shown for each of the four phases during the sample period. Phase 1 begins March 26, 2004, Phase 2 begins February 24, 2006, Phase 3 begins January 29, 2009, and Phase 4 begins February 29, 2012.

introduced, the results are more definitive. In 68% of the days in this phase, VIX led SPVXSTR using the R^2 criterion. The introduction of VIX ETPs at the beginning of Phase 3 appears to have created an intervention. During Phase 3, cash leads futures in only 37% of the 772 days. In other words, futures leads cash 63% of the time. During Phase 4, futures leads cash 75% of the time.

The bootstrap results are also contained in Table IV. In phase 1, the P -values are less than or equal to 0.05 in 11% of days studied. This means that the lead/lag relation with futures leading cash is significant at the 95% confidence level for 11% of days in Phase 1. In the same phase, cash leads the futures with 95% confidence in 7% of the days. Interestingly and importantly, note the profound change in the lead/lag relation that has occurred through time. In Phase 4 that begins on February 29, 2012 and ends on April 30, 2013, the futures leads cash in 31% of the days at 95% confidence, while the cash leads the futures in none.

For the period where cash leads, the average duration of the lag with maximum correlation occurs at 5 minutes and 40 seconds in Phase 1 and 5 minutes in Phase 2, suggesting that cash leads by more than 5 minutes. For the period where futures leads, the average duration of the lag with maximum correlation occurs at 1 minute and 43 seconds in Phase 3 and 1 minutes and 11 seconds in Phase 4, suggesting that futures leads by more than 1 minutes. Figure 6 shows the cross-correlation function in the different phases of the sample period. For Phases 3 and 4, the right hand side of the cross-correlation function highlights that futures leads cash, with the correlation remaining meaningful for up to 5 minutes. In Phase 2, the left hand side of the function highlights that cash leads futures, with a lag which also appears to persist for up to 5 minutes.

In summary, the evidence clearly points to the fact that the tail is now wagging the dog. The VIX was created to be a barometer of market fear and a reference asset to use in the development of new exchange-traded derivatives. But, it has become a victim of its own success. VIX futures now predict the direction in which the VIX will move. This does not imply that trading profits can be earned. Buying or selling the VIX index to capture a predicted move is much too costly. Further research could involve lead/lag analysis of the complete volatility surface calculated from the S&P 500 index options which make up the VIX.

6. SUMMARY AND CONCLUSIONS

A decade ago, trading volatility on exchanges was in its infancy. Daily VIX futures trading volumes numbered in the hundreds of contracts. Now they trade in tens of thousands. Since that time, other exchange-traded volatility products have been introduced—VIX options in 2006 and VIX ETPs in 2009. Each has commanded a significant trading activity by institutions and retail customers alike. The purpose of this paper is to try to understand the supply of and demand for volatility through careful analyses of the lead/lag time-series relations between the returns of different volatility products.

The study uncovers several important empirical facts. First, the price relation between the VIX futures and VIX options are linked by put-call parity, and our analysis shows that the duration of any lead/lag relations is short-lived. Second, we find that VIX ETPs and the VIX futures are largely contemporaneous. They, too, are governed by an arbitrage price relation. The length of the lead/lag is trivial, presumably as a result of both markets being highly active. Finally, the relation between the VIX futures and the VIX cash index has changed through time. Because the relation between the prices is not arbitrated, contemporaneous price moves are not expected. The direction of the causality, however, remained an open question. We find that the VIX futures price lagged the VIX cash index in the first few years after it was launched, but the VIX futures now leads the VIX. The VIX is determined by S&P 500 index

options, and, when the VIX futures market was launched, the dominant supply of exchange-traded volatility contracts was the S&P 500 index option market. Through time, the interest in VIX futures grew. Now, it has apparently supplanted the SPX options.

APPENDIX. CONTRACT DETAILS FOR VIX FUTURES AND OPTIONS.

The CFE lists for trading up to nine near-term serial months and 5 months on the February quarterly cycle for the VIX futures contract. The minimum price interval per tick is 0.01 points or \$10.00 per contract. The close of trading is the day before the final settlement date. The final settlement date is the Wednesday that is 30 days prior to the third Friday of the calendar month immediately following the month in which the contract expires. The final settlement value for VIX futures (VRO) is a “special opening quotation” of VIX calculated from the sequence of opening trade prices of the S&P 500 index options used to calculate the index on the settlement date. The opening price for any series in which there is no trade is the average of that option’s opening bid and ask price quotes. The final settlement value is rounded to the nearest \$0.01. Open contract positions are settled in cash on the business day immediately following the final settlement date. On the days prior to the settlement date, the VIX futures trade from 8:30 AM to 3:15 PM Central Time (Chicago time).

VIX options are European-style. One index point equals \$100, and the minimum tick for series trading below \$3 is 0.05 (\$5.00), and above \$3 is 0.10 (\$10.00). The CBOE lists for trading up to six contract months, provided that the time to expiration is no greater than 12 months. The minimum exercise price intervals of the options are (i) \$0.50 where the exercise price is less than \$15; (ii) \$1.00 where the exercise price is less than \$200; and (iii) \$5 where the exercise price is greater than \$200. When a new contract month is first listed, exercise prices for an in-, at- and out-of-the-money strike prices are introduced, but new exercise prices are added as the index moves up or down. The VIX options expire together with the VIX futures on the Wednesday that is 30 days prior to the third Friday of the calendar month immediately following the expiring month. The last day of trading is the preceding Tuesday. The exercise-settlement value for the VIX options is the same as for the VIX futures.

REFERENCES

- Barndorff-Nielsen O. E., & Shephard, N. (2005). Measuring the impact of jumps in multivariate price processes using bipower covariation. Discussion paper, Nueld College, Oxford University.
- Bibinger, M., Mykland, & Per A. (2013). Inference for multi-dimensional high-frequency data: Equivalence of methods, central limit theorems, and an application to conditional independence testing. *arXiv 1301.2074* online at <http://arxiv.org/pdf/1301.2074.pdf>.
- Black, F. (1976a). Studies of stock price volatility changes. Proceedings of the Business and Economics Section of the American Statistical Association, 177–181.
- Black, F. (1976b). The pricing of commodity contracts. *Journal of Financial Economics*, 3, 167–179.
- Bollen, N., Smith, T., & Whaley, R. (2004). Modeling the bid/ask spread: Measuring the inventory-holding premium. *Journal of Financial Economics*, 72, 97–141.
- Chan, K., Chung, P., & Johnson, H. (1993). Why option prices lag stock prices: A trading-based explanation. *Journal of Finance*, 48, 1957–1967.
- Christensen, K., Kinnebrock, S., & Podolskij, M. (2010). Pre-averaging estimators of the ex-post covariance matrix in noisy diffusion models with non-synchronous data. *Journal of Econometrics*, 159, 116–133.
- Covrig, V., & Melvin, M. (2002). Asymmetric information and price discovery in the FX market: Does Tokyo know more about the yen? *Journal of Empirical Finance*, 9, 271–285.
- Crain, S. J., & Lee, J. H. (1995). Intraday volatility in interest rate and foreign exchange spot and futures. *Journal of Futures Markets*, 15, 395–422.
- Danielsson, J., & Payne, R. (2002). Real trading patterns and prices in spot foreign exchange markets. *Journal of International Money and Finance*, 21, 203–222.

- De Jong, F., & Nijman, T. (1997). High frequency analysis of lead-lag relationships between financial markets. *Journal of Empirical Finance*, 4 (2–3), 259–277.
- De Jong, F., Mahieub, R., & Schotman, P. (1998). Price discovery in the foreign exchange market: An empirical analysis of the yen/dmark rate. *Journal of International Money and Finance*, 17, 5–27.
- Epps, T. W. (1979). Comovements in stock prices in the very short run. *Journal of the American Statistical Association*, 74, 291–298.
- Frijns, B., Alireza, T.-R., & Webb, R. I. (2016). On the intraday relation between the VIX and its futures. *Journal of Futures Markets*, 36, 870–886.
- Fung, J.K.W, Lau, F., & Tse, Y. (2015). The impact of sampling frequency on intraday correlation and lead/lag relations between index futures and individual stocks. *Journal of Futures Markets*, 35, 939–652.
- Granger, C.W.J, Huang, B.-N., & Yang, C.-W. (2000). A bivariate causality between stock prices and exchange rates: Evidence from recent Asian flu. *Quarterly Review of Economics and Finance*, 40, 337–354.
- Hayashi, T., & Yoshida, N. (2005). On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli*, 11, 359–379.
- Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6, 327–343.
- Hoffmann, M., Rosenbaum, M., & Yoshida, N. (2013). Estimation of the lead-lag parameter from non-synchronous data. *Bernoulli*, 19, 426–461.
- Hou, K., Barberis, N., & Chen, X. (2007). Information diffusion and asymmetric cross-autocorrelations in stock returns. Working paper, <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.8.6147&rep=rep1&>.
- Huth, N., & Abergel, F. (2014). High frequency lead/lag relationships—Empirical facts. *Journal of Empirical Finance*, 26, 41–58.
- Kaen, F. R., & Hachey, G. A. (1983). Eurocurrency and national money market interest rates: An empirical investigation of causality. *Journal of Money, Credit and Banking*, 15, 327–338.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53, 1315–1336.
- Lo, A. W., & MacKinlay, A. C. (1990). When are contrarian profits due to stock market overreaction. *Review of Financial Studies*, 3, 175–205.
- Mason, S. & Onur, E. (2014). Volatility derivatives in practice: Activity and impact, CFTC research report.
- Pan, M., Hocking, R. T., & Rim, H. K. (1996). The intertemporal relationship between the currency spot market and the currency option market. *Journal of Business Finance and Accounting*, 23 (9–10), 1307–1317.
- Pomponio, F., & Abergel, F. (2013). Multiple-limits trades: Empirical facts and application to lead/lag measures. *Quantitative Finance*, 13, 783–793.
- Scholes, M., & Williams, J. (1977). Estimating betas from nonsynchronous data. *Journal of Financial Economics*, 5, 309–327.
- Shu, J., & Zhang, J. (2012). Causality in the VIX futures market. *Journal of Futures Markets*, 32, 24–46.
- Stephan, J. A., & Whaley, R. E. (1990). Intraday price change and trading volume relations in the stock and stock options markets. *Journal of Finance*, 45, 191–220.
- Stoll, H. R., & Whaley, R. E. (1990). The dynamics of stock index and stock index futures returns. *Journal of Financial and Quantitative Analysis*, 25, 441–468.
- Todorov, V., & Tauchen, G. (2011). Volatility jumps. *Journal of Business & Economic Statistics*, 29, 356–371.
- Whaley, R. E. (1993). Derivatives on market volatility: Hedging tools long overdue. *Journal of Derivatives*, 1, 71–84.
- Whaley, R. E. (2013). Volatility trading: At what cost? *Journal of Portfolio Management*, 40, 95–108.
- Zhang, J., Shu, J., & Brenner, M. (2010). The new market for volatility trading. *Journal of Futures Markets*, 30, 809–833.
- Zhang, J., & Zhu, Y. (2006). VIX futures. *Journal of Futures Markets*, 26, 521–531.