Volatility Derivatives in Practice: Activity and Impact

Scott Mixon*
Esen Onur*

June 2015

Abstract:
We use unique regulatory data to examine open positions and activity in both listed and OTC volatility derivatives. Gross vega notional outstanding for index variance swaps is similar in size to that of S&P 500 options, with dealers short vega in order to supply the long vega demand of asset managers, but VIX futures are the dominant instrument for maturities less than one year. We find two distinct channels by which VIX trading activity manifests itself into futures prices: a risk-based level effect and an order flow effect. We estimate that the long volatility bias of asset managers puts upward pressure on VIX futures prices. Hedge funds have offset this potential impact by actively taking a net short position in nearby contracts. We also find that the front part of the curve is steepened by the order flow effect. In our 2011-2015 sample, the net impact added less than half a volatility point, on average, to nearby VIX futures contracts but added approximately one volatility point for contracts in less liquid, longer-dated parts of the curve. We find no evidence that this price impact forces VIX futures outside no-arbitrage bounds.

---

* Office of the Chief Economist, Commodity Futures Trading Commission, 1155 21st Street, N.W., Washington, D.C. 20581. Mixon: smixon@cftc.gov, (202) 418-5771. Onur: eonur@cftc.gov, (202) 418-6146. We thank seminar participants at the CFTC, SEC, and the Federal Reserve Board of Governors for helpful comments.

Abstract:
We use unique regulatory data to examine open positions and activity in both listed and OTC volatility derivatives. Gross vega notional outstanding for index variance swaps is similar in size to that of S&P 500 options, with dealers short vega in order to supply the long vega demand of asset managers, but VIX futures are the dominant instrument for maturities less than one year. We find two distinct channels by which VIX trading activity manifests itself into futures prices: a risk-based level effect and an order flow effect. We estimate that the long volatility bias of asset managers puts upward pressure on VIX futures prices. Hedge funds have offset this potential impact by actively taking a net short position in nearby contracts. We also find that the front part of the curve is steepened by the order flow effect. In our 2011-2015 sample, the net impact added less than half a volatility point, on average, to nearby VIX futures contracts but added approximately one volatility point for contracts in less liquid, longer-dated parts of the curve. We find no evidence that this price impact forces VIX futures outside no-arbitrage bounds.

1. Introduction

Are VIX-linked products distorting the VIX? Some commenters have questioned whether VIX-linked products (Exchange Traded Funds or Exchange Traded Notes) have grown so popular that the hedging activities of the issuers dominate and distort the VIX futures market. Pessimists ask if liquidity-driven disruptions in VIX ETN markets are being transmitted to the VIX futures markets; optimists point to increased (indirect) participation in futures markets providing better liquidity and more informative pricing. Analysis is complicated by the fact that related volatility derivatives are actively traded in the OTC market, and little is known about the trading activity in that space.

We begin our analysis by presenting a novel comparison of activity in the listed and OTC markets for volatility; we conclude that the majority of the transactions are in the VIX futures market. We subsequently focus on VIX futures and quantify two distinct channels by which VIX trading activity manifests itself into futures prices. First, we identify a risk-based effect that we identify with end-users tending to be net long VIX futures and dealers tending to be short VIX futures. Consistent with Gârleanu, Pedersen, and Poteshman's (2009) notion that dealers provide liquidity in option products and charge for the unhedgeable risks they assume, we find that a larger net long position for end-users causes the futures price to be higher relative to the spot price. Second, we identify an order flow, or "price pressure" effect: as end-users transact large quantities (primarily in the front two contracts), the prices of those contracts exhibit temporary price changes that are subsequently corrected. We go on to find that these statistically significant price effects are economically small by some measures.

We use unique regulatory data on both OTC variance swaps and exchange-listed VIX futures to evaluate the activity and impact of volatility derivatives trading. This study is the first to use the market-wide regulatory data on variance swap transactions mandated by the Dodd-Frank Wall Street Reform and Consumer Protection Act. Both the swaps and futures data we examine include information on the participant holding any particular reported position; we are therefore able to link precisely the participant type with the positions and need not estimate such relationships.

We find that gross vega notional outstanding for index variance swaps is over USD 2 billion, with USD 1.5 billion in S&P 500 products (similar in size to that of exchange-traded S&P 500 options[1]) and nearly USD 500 million across five other major indexes across the globe. Dealers are net short vega in order to supply the long variance swap demand of unlevered asset managers. For maturities less than one year, VIX futures are far more actively traded and have twice as much notional vega outstanding than S&P 500 variance swaps.

To the extent that dealers take on risk when facilitating trades, we estimate that the long volatility bias of asset managers acts to put upward pressure on VIX futures prices. Hedge funds have offset this potential impact because they have actively taken a net short position in nearby contracts. In our 2011-2014 sample of VIX futures, the net impact added less than half a volatility point in nearby contracts, on average, but between one and two volatility points for contracts in less liquid, longer-dated parts of the curve. The impact of non-dealer positioning has evolved over time as the demands of the various participant types have changed. In 2011 and 2012, for example, unlevered asset managers were long and hedge funds were not active, and we conclude that this positioning increased front-month VIX futures approximately one volatility point. By 2013, hedge funds took sizeable short positions and the estimated impact fell to roughly zero.

We also find significant results when we interpret the data from a term structure perspective: the active rolling of contracts by some large asset managers in order to keep a target maturity appears to have a small but significant impact on the slope of the term structure. The price of the front month contract has been depressed by this persistent selling, on average, while the price of the second month contract has been increased. In both cases, this term structure effect averages less than half a volatility point. Throughout the entire period analyzed, we find no evidence that these price impacts forced VIX futures outside no-arbitrage bounds.

[1] See, for example, Barclays Bank's *The VIX Compass*, 4 June 2014, produced by Maneesh Deshpande and team.

The outline of the paper is as follows. First, we quantify activity in the OTC market and contrast the findings with those for the VIX futures markets. We use transaction-level regulatory data to gauge typical activity in index variance swaps in various major indexes. Second, we quantify the open commitments of variance swaps. In both exercises, we break out the transactions due to interdealer trading and the ones due to dealer/end-user trading. We also evaluate the net positioning of various types of market participants across products and across the term structure for the S&P 500. Given our finding that the most active part of the volatility derivative space is the listed futures market, we focus on VIX futures for the remainder of the analysis. We model the impact on VIX futures prices of futures positions and end-of-day position changes for both the buy-side and sell-side.

The papers closest in spirit to our work are those of Bollen and Whaley (2004), Bollen, O'Neill, and Whaley (2013), and Garleanu, Pedersen, and Poteshman (2009).[2] These papers provide strong empirical evidence that dealer capacity and buying pressure are key determinants of derivatives prices. In particular, Bollen, O'Neill, and Whaley (2013) examine the buying pressure of VIX ETNs on VIX futures and the VIX; they conclude that the long volatility bias of asset managers has strongly impacted VIX futures pricing. A portion of our work examines that question, but we use daily regulatory data to examine the impact of asset managers. The richness of this data allows us to identify the actual quantity (not an estimate) of futures demanded by asset managers as well as the potentially offsetting short demand by hedge funds. We separately identify the impact of both asset managers and hedge funds on the VIX futures pricing and argue that both impacts should be analyzed together to understand the dynamics of this market.

2. ACTIVITY IN VARIANCE SWAPS AND VIX FUTURES

[2] A variety of papers have examined (synthetic) variance swap prices and VIX futures prices, including Carr and Wu (2009), Egloff, Leippold, and Wu (2010), Nossman and Wilhemsson (2009), and Simon and Campasano (2014). Other papers have examined the empirical link between index implied volatility and macroeconomic, risk, or sentiment variables, e.g., Corradi, Distaso, and Mele (2013), David and Veronesi (2014), Glatzer and Scheicher (2005), Han (2008), Mixon (2002), Mixon (2007), and Vähämaa and Äijö (2011). These empirical analyses do not directly link observed trading activity and positions to market prices.

The type of OTC data we examine was not available to regulators or market participants prior to the current reporting regime and has not been analyzed before. While individual participants knew their own transactions and may have had second hand knowledge of other transactions, there are no publicly available sources of OTC transaction data for these swaps. Presently, large participants in the US swaps market must register with the Commodity Futures Trading Commission (CFTC) and report swap transactions to Swap Data Repositories (SDR). Basic details on transaction activity are reported to the public by SDRs, while comprehensive details identifying the counterparties and their trade directions are available to regulators.

The repositories consequently offer an unprecedented, transaction-level view of derivatives markets. The data is not just time and sales type information; it is data in the same spirit as the CFTC's Large Trader Reporting System data, which provides position-level information on participants. The aggregation of derivatives transaction data into SDRs provides a powerful tool to regulators and market participants. Confidential information from individual participants remains confidential, but regulators can present the aggregated SDR data to the public in order to promote sound risk management while preserving the liquidity and price discovery role of derivatives markets.

The remainder of this section elaborates on our description of activity and positioning in the variance space. We find that new swaps are predominately initiated in a handful of indexes, with nearly half of the new activity in the S&P 500. In turn, a large fraction of the S&P activity is dealer/dealer trading. The majority of activity in the front part of the curve is in VIX futures trading, but variance swap activity extends to maturities out a decade or more. Vega outstanding aggregates to USD 2 billion, with USD 1.5 billion in S&P 500 swaps. Roughly two thirds of the S&P 500 notional outstanding has dealers on both sides of the trade. Dealers, on net, are significantly short variance across the major indexes.

2.1 Transaction Activity

We begin with a brief description of variance swap market terminology to facilitate the description of our summary statistics. The terminal valuation of a variance swap is given by the expression $N_{var}\ (\sigma_{realized}^2 - \sigma_{strike}^2)$, where N_{var} is the variance notional of the transaction, $\sigma_{realized}$ is 100 times the annualized realized volatility of the reference asset over the life of the swap, and σ_{strike}^2 is the variance strike. The variance strike is thus the fixed leg of the swap and the realized variance is the floating leg. The floating leg of the variance swap is defined by the equation $\sigma_{realized} = 100 \times \sqrt{\frac{252}{n}\sum_{i=1}^{n} r_i^2}$, where r_i is the daily log price change in the reference asset over the n observation dates.

Industry convention is to set the variance notional of the swap using the equation $N_{var} = \frac{N_{vol}}{2\sigma_{strike}}$, where N_{vol} is the desired vega notional of the swap. This vega notional is approximately the value of the swap if, at the valuation date of the contract, the realized volatility and the square root of the variance strike (the volatility strike) differ by one vol point. Consistent with industry usage of this intuitive notional value, commonly referred to as "dollars per vol point", we also reference the vega notional of the swaps throughout this paper.

We rely on SDR data on open swap positions at the beginning of June 2014 for our cross-sectional analysis of variance swaps outstanding. We use the same dataset to identify new swaps initiated during the month of May 2014, which is a straightforward measure of transaction activity in this market. Futures statistics are computed using publicly available end of day reports for May 2014. The raw swap data contains approximately 19,000 records, and we filter it in the following ways. We remove roughly 6,000 transactions between affiliated entities within the same corporate family (e.g., a transaction between "Dealer A – New York" and "Dealer A – London"). In many instances, these transactions occur virtually simultaneously as a transaction by one of the dealer entities with an end-user on the other side of the trade. For the purposes of this study, they are taken to be more of an accounting exercise than a market-facing transaction, and they are deleted from the dataset. This practice is also followed by the Bank for International Settlements in its surveys of the global derivatives market. We also eliminate obvious transcription errors. The resulting dataset of open swaps contains 12,314 records and the dataset of new trades contains 1,384 records.

Table 1 displays summary statistics for new variance swaps initiated during the sample period. Several observations are immediately apparent from the table. First, trade in S&P 500 contracts represents the most activity across products, with approximately half of the USD 160 million vega initiated in these instruments. There were 21 regular trading days during the sample period, yielding an average daily volume of USD 4 million in S&P 500 variance; transactions were observed on each day in the sample. Nearly USD 2 million SPX vega was transacted with end-users, on average, per day during the sample. Transactions on the EuroStoxx 50 were second most prominent, with transactions taking place with end-users on most, but not all, days of the sample. Activity on other products is often far less than on these top two.

A second observation is that transactions between swap dealers are typically larger in size and long-dated than transactions between swap dealers and end-users. For example, the median trade size (not shown in the table) for S&P 500 swaps was USD 100,000, whether measured across all trades or for dealer/dealer trades or dealer/end-user trades. The average size is USD 132,000 for dealer/dealer trades and 112,000 for dealer/end-user trades, reflecting some large trades between dealers. Similarly, dealer/dealer trades in the SPX have a tenor about six months longer than dealer-end-user trades (1.3 years vs. 0.8 years). These facts are consistent with typical dealer activity documented by Naik and Yadav (2003), with dealers facilitating many trades by end-users and hedging the entire book in the inter-dealer market.

Figure 1 breaks out the new trades during the sample period by the tenor of the trades.[3] The vertical axis represents the USD vega transacted, and the horizontal axis represents the valuation dates for the swaps. The height of the line at a given valuation date therefore displays the cumulative activity in swaps maturing before or on that date. For example, approximately USD 35 million of vega was transacted in S&P 500 swaps maturing in December 2014 or earlier. Another USD 50 million in vega was transacted in swaps maturing between that date and December 2017, with a small amount transacted in dates beyond December 2017. The same

3 Other papers that explore the term structure of variance swaps are Egloff, Leippold, and Wu (2010), Dew-Becker, Giglio, Le, and Rodriguez (2014), and Ait-Sahalia, Karaman, and Mancini (2014).

pattern shows up for the other indexes presented, with the levels of transactions and tenors of transactions less than those for the S&P.

For context, the chart also shows May 2014 trading volume for VIX futures. VIX futures activity dwarfs that of the variance swap market, but the activity is concentrated in the first few contracts. VIX futures trading totals USD 3 billion of vega over the month (an average of 144 million vega per day). We recognize that notional futures trading volume is not perfectly comparable with the notional value of new swaps positions (and that VIX futures and variance swaps are not perfect substitutes). The goal of the juxtaposition, however, is to illustrate the vastly different characters of the two markets. The futures market exhibits many transactions for a given interval of time, whereas the swaps market is more about larger, more infrequent trades. The futures market also focuses on the short end of the curve, with contracts only extending to December 2014 and the majority of volume occurring in the first few contracts.

2.2 Open Transactions

Table 2 provides a summary of the variance positions outstanding at the end of May 2014. As expected, the vast majority of positions are in the S&P 500. Of the USD 2.2 billion vega outstanding in the SDR data, about 1.5 billion is in the SPX. Of the SPX positions, about 2/3 are between dealers, with the remaining 428 million vega representing open dealer/end-user positions. The EuroStoxx 50 is the second largest set of positions in the market, with a bit more than 100 million of positions between dealers and a bit less than 100 million representing positions with end-users. The original tenor of the swaps currently open averaged 3.9 years for the SPX, representing a blend of interdealer positions averaging 4.3 years and positions with end-users averaging 2.9 years. Tenors for other indexes are shorter but show the same pattern.

The set of open positions also can be used to quantify the concentration of the positions across dealers. In this computation, the concentration is assessed by aggregating positions in different subsidiaries of a parent company. (For example, positions held by Dealer A - New York and Dealer A – London are aggregated into one set of positions.) As with many dealer markets, the top few dealers capture the lion’s share of positions. The concentration is such that 70-95% of positions are with the top five dealers in this data. While this is true for swaps on each individual index, the concentration for the top five dealers across all of the data is slightly lower at 68%, suggesting that it is not the same five dealers who hold the majority of positions in every index.

The CFTC's Commitments of Traders report has provided transparency on positioning across various market participant types for decades. Figure 2 provides a similar view on net positioning across participant types for VIX futures and the major variance swaps.[4] With respect to the buy side, the chart suggests that asset managers are net long vega in each of the indexes, with the largest net positions in S&P 500 variance and VIX futures (totaling around USD 150 million in the two of them together). Leveraged money is generally short vega, with the largest position, by far, short VIX futures. Dealers are short vega, with a large short position in S&P 500 variance.

Figures 3 and 4 focus on S&P 500 exposure across VIX futures and variance swaps of different tenors. Figure 3 displays the gross notional amount outstanding, while Figure 4 displays net positioning across market participant types. Both figures break out S&P 500 positions across tenors, showing the consistency of positioning across the term structure. The displays are for a single bucket covering VIX futures (which all expire in less than a year) and four buckets for SPX variance covering the zero to one year, one to five year, five to ten year, and ten to twenty year tenors. With respect to gross notional, Figure 3 illustrates that the VIX futures footprint is more than twice as large as that of S&P variance (USD 470 million vs. USD 200 million) in the zero to one year bucket. The S&P variance buckets for longer-dated tenors are larger than the VIX futures buckets, however. With respect to net notional, Figure 4 makes clear that the positioning in variance is consistent across tenors for the participant types. Asset managers are consistently long volatility but there is a pronounced long position in the five to ten year bucket. Dealers are short vega across the term structure and are especially short in the five to

[4] The participant types are broadly the same as those used for the Commitments of Traders (COT) report. The "Asset Manager" category includes investors such as pension funds, endowments, insurance companies, and mutual funds. Participants labeled "Leveraged Funds" are typically hedge funds. We reclassify some entities from the COT mapping to facilitate our empirical analysis in the second part of the paper, which identifies dealers as passive liquidity providers. For example, we would classify a dealer entity managing a retail product book in the "Asset Manager" category. Swap data parties are identified by Legal Entity Identifier (LEI) and are classified manually for this project.

ten year bucket in a mirror image of the asset manager positioning. Leveraged money is significantly short VIX futures, but the positions in variance are relatively small irrespective of the tenor bucket. The chart also highlights the dominance of VIX futures positioning in the most active part of the curve, the portion that expires in one year or less.

3. IMPACT OF TRADING ON MARKETS

The evidence thus far confirms one point that most market participants already believed: the major instrument for volatility trading is the S&P 500. For the remainder of this paper, we focus solely on the S&P 500 and on the major listed volatility future, the VIX future. Given this narrow focus, we can explore very specific, sharp evidence on the impact of positioning and trading in VIX futures. The main finding from this section is that the positions of the various market participants do, in fact, move the futures price. Net positive end-user demand acts to raise the price, while significant selling (buying) of contracts acts to lower (raise) the price over short horizons such as one day.

3.1 Data and Theory

The data examined in this section are from the Commodity Futures Trading Commission's Large Trader Reporting System. The sample covers the four year period from April 2011 to March 2015. The data are daily market settlement prices for the first six futures contracts and for the VIX, as well as daily levels and changes in levels of net positions by market participant type. Longer dated contracts exist, but they are thinly traded.[5]

Positions are aggregated across traders and accounts into participant types for each contract outstanding on a given date. Market participants are categorized in this section as "Dealers", "Asset Managers", or "Leveraged Funds". This categorization is generally similar to that used by the CFTC's Commitments of Traders (CoT) report, although we classify some entities in different roles due to our slightly different analytical focus. For example, we aggregate some participants into the "Leveraged Funds" category that might otherwise be classified as "Financial – Other" in

[5] We limit our focus to the more active, shorter maturities in order to avoid inadvertent revelation of confidential information. We also limit our sample to a recent period when the various categories of market participant types include enough diversity to minimize any concerns regarding confidential information.

a CoT report; this categorization choice has no major influence on our general findings. We also ensure that traders known to be associated with ETN issuance are classified as Asset Managers, even if they are within an entity that is otherwise classified as a Dealer.

The theoretical framework utilized has been carefully articulated by Gârleanu, Pedersen, and Poteshman's (2009) work. They discuss a logically consistent theory in which dealer capacity constraints act to drive a wedge between market prices and textbook no-arbitrage prices. In their empirical implementation, they find that the level and variation of the typical gap between implied and realized volatility for S&P 500 index options can be explained by the net positive non-dealer demand for optionality. Dealers facilitate the trades but cannot hedge perfectly, so they charge a premium price based on the amount of risk they take on. Models that explicitly incorporate such financial intermediation appear to be very important in explaining option prices.

While the foundation of the empirical analysis is that the net stock of dealer vega positions outstanding at a given time might impact derivatives pricing, we also allow that order flow of positions can have impact on derivatives pricing. We generally refer to the first effect as a "level effect" and the second as a "roll effect", reflecting the popular notion that the systematic rolling of passive index positions to maintain a constant maturity impacts the shape of the VIX futures term structure.

3.2 Empirical Results

3.2.1 Risk effects.

Table 3 provides evidence on the dealer capacity issue with respect to the pricing of volatility. The first test is whether the price of a VIX future varies according to the level of dealer positions in VIX futures. Their capacity to take on short volatility positions is finite, and theory suggests that they charge higher prices to take on more risk. The test is conducted separately for each contract *i* and is based on the regression:

$$v_i - v_0 = \alpha_i + \beta_i v_0 + \gamma_i (Dealer_i) + \varepsilon_i. \tag{1}$$

In this setup, v_i is the daily settlement price for contract *i*, v_0 is the VIX spot price, and $Dealer_i$ is the net position (the number of contracts) of dealers in contract *i*. We follow Gârleanu *et al.* and assume that end-user demand is inelastic and exogenous, with dealer prices set such that markets clear at each date. The maintained hypothesis is that the VIX is exogenously determined relative to VIX futures prices (we perform robustness tests for this assumption later in the text).[6] Allowing the level of volatility to enter as an explanatory variable allows us to focus explicitly on futures pricing. We are not attempting to relate the overall level of volatility to futures trading; our goal is to isolate the impact of futures pricing related to futures trading. The tested hypothesis is that the coefficient γ_i is significantly negative, signifying that the futures price is negatively related to the level of positioning by dealers. A lower than average (i.e., more net short) dealer position is related to a higher futures price (equivalently, a higher futures basis).

Panel A of Table 3 displays the regression results for the first six contracts. For five of the six contracts, the coefficient is significantly negative, with Newey West (1987) standard errors delivering t-statistics well above conventional significance levels. The results strongly suggest that, at the margin, market prices of VIX futures are higher the more dealers are asked to be short volatility.

The second test is an elaboration of the first test and uses the fact that dealer positioning is equal and opposite in sign to net end-user demand.[7] The idea is that end-user demand (i.e., the buy-side demand for vega) differs across market participant types. In particular, we regress the futures price on the level of VIX futures positions held by two separate segments of the buy side: asset managers and leveraged funds (i.e., hedge funds). The hypothesis to test is whether the futures price is positively related to the net positions of the two types of end-user demand.

---

[6] For our sample, the evidence from Dickey-Fuller and KPSS tests is consistent with stationarity for the futures basis (futures minus VIX), but the results for futures levels are not consistent with stationarity. Similarly, the evidence suggests that the levels of the net position variable for Dealers, for Asset Managers, and for Leveraged Funds are each stationary.

[7] The reported data do not exactly offset in practice because not all positions are reportable, but the discrepancies are small and, in aggregate, average about 1,000 contracts for each date in this sample.

For notation, let AM_i be the net position (in contracts) for Asset Managers in the *i*th maturity contract, and let LF_i be the net position for Leveraged Funds in the *i*th maturity contract. As before, the maintained hypothesis is that the VIX is exogenous to the VIX futures price. The regression is estimated separately for each contract with the specification as follows:

$$v_i - v_0 = \alpha_i + \beta_i v_0 + \gamma_{1i}(AM_i) + \gamma_{2i}(LF_i) + \varepsilon_i. \quad (2)$$

Panel B of Table 3 displays the results of the regressions. For the six $\gamma_{1\mathrm{i}}$coefficients, the evidence strongly supports a positive relation between the level of Asset Manager positions and the level of VIX futures, with all of the coefficients highly significant at conventional levels. For the $\gamma_{2\mathrm{i}}$ coefficients, the evidence strongly supports a positive relation between futures prices and Leveraged Funds positions, with five of the six coefficients easily significant at conventional levels. The basic message from these regressions is that the size of the positions for end-users has a significant impact on the level of VIX futures: the data supports the dealer capacity theory.

To address potential concerns about endogeneity among the relevant variables, we also tested a variant of equations 1 and 2. We relax the exogeneity assumption on the VIX by replacing it on the right hand side with the closing value of the VSTOXX. The VSTOXX is computed using the same methodology as the VIX, but it is computed on the EURO STOXX 50 index. The value is computed at the close of European trading, which is mid-day U.S. time. We also replace the contemporaneous values of Dealer, Asset Manager, and Leveraged Funds with the value lagged by a week (5 trading days), e.g.,

$$v_i - v_0 = \alpha_i + \beta_i VSTOXX + \gamma_{1i}\left(AM_{i,t-5}\right) + \gamma_{2i}\left(LF_{i,t-5}\right) + \varepsilon_i. \quad (2')$$

The results are virtually unchanged using these specifications, confirming that the results are not distorted by the statistical assumptions. While the R^2 values and t-statistics predictably decline, the relevant t-statistics are all well beyond conventional significance levels. The full set of results is displayed in the Appendix.

Researchers such as Bollen, O'Neill, and Whaley (2013) have concluded that the significant VIX futures positions held by asset managers impacts the price of VIX futures, and the worry is that these impacts are distorting prices, potentially due to trading by uninformed noise traders. The regression results discussed above are in line with the conclusion that the large, long positions of Asset Managers have acted to increase VIX futures prices above what they might be in absence of those positions. On the other hand, the regressions include the large, short positions of Leveraged Funds, which act to decrease the futures price, everything else held constant.

Because futures contracts are in net zero supply, the large net long demand by Asset Managers could be offset in equilibrium by a) Leveraged Funds getting net short, b) dealers facilitating the demand and taking short positions, or c) a combination of the two. Asset Managers have displayed large long positions in this sample, averaging +60,000 contracts, and Leveraged Funds have displayed large short positions, averaging -62,000 contracts. Net non-dealer demand has averaged -1,200 contracts over the sample, representing a much smaller net position (and the opposite sign) than Asset Managers alone. We will further explore the economic impact of the positions of various market participant types in a later section.

3.2.2 Price Pressure Effects.

Are there measurable "price pressure" or order flow effects in this market? The next test gauges whether net position changes by the market participants are related to changes in futures prices. This question is motivated by the observation that passive Asset Managers in VIX futures often target a constant maturity and mechanically buy and sell contracts to maintain this target. As the positions have grown in size over the past few years, market participants have speculated that the constant rolling of contracts by passive managers (selling shorter dated contracts to buy longer dated ones) has put more curvature into the VIX futures term structure or acted to steepen it in the region around one month to maturity (the target maturity for many funds).

The first differenced regression tested is

$$\Delta(v_i - v_0) = \alpha_i + \beta_i \Delta v_0 + \delta_{1i}(\Delta AM_i) + \delta_{2i}(\Delta LF_i) + \varepsilon_i. \quad (3)$$

This specification is simply a first difference version of the levels regression used above. The hypothesis to test is whether the change in position sizes is positively related to the change in the futures basis. Examination of Table 4 suggests that there is a systematic relation between the change in the futures basis and the change in the net positions for Asset Managers. The evidence is pervasive and shows up for the first, second, and third contracts, and the t-statistics are well into the significant range. There is less evidence for an impact for the trading by Leveraged Funds, although the coefficient for the second contract is significant.

The models described thus far are designed to examine the contemporaneous impact of non-dealer position size and order flow on futures prices for a given day. We now explore a more dynamic version of this order flow relationship, not only to gain a better understanding of how order flow is associated with prices, but also to mitigate concerns about misspecification or reverse causality for the price change regressions. In particular, we provide evidence that the significant correlations uncovered in the first difference regressions illustrate a causal, transient "price pressure" effect and do not reflect informationally-motivated trading.

The "price pressure" view is more likely to be valid if we observe a price reversal following a change in prices that is contemporaneous with a shock to order flow. If price changes associated with order flow are found to be permanent, the significant correlations uncovered in the first difference regressions might be more properly attributed to a factor that simultaneously drives both prices and trading. A third potential situation is that feedback effects exist and lead to an amplification of shocks to prices. For example, a positive shock to prices might generate an increase in net longs by asset managers that leads to further price increases, and so on.

We apply three different tests to explore the "price pressure" effect further. Our first test looks at the unexpected component of end-user order flow to test whether it is new information that is associated with contemporaneous price change on a daily basis. The second and third tests are more dynamic in nature; the second test explores how changes in end-user flow affect the

VIX futures basis on different horizons. Finally, the third test traces out the impact of an unexpected shock to end-user order flow on futures prices.

To examine the price pressure hypothesis, we first refine the regression to relate futures price changes to the unexpected component of end-user order flow. The analysis confirms that the regression is picking up a positive correlation between prices and the component of flow specific to the current day. It is not reflecting, for example, a lingering effect of order flow from prior days. As in the previous literature relating inflows to mutual funds and asset prices (e.g., Ben-Rephael, Kandel, and Wohl (2011), Edelen and Warner (2001), Warner (1995)), we find that the significant contemporaneous correlation between prices and flows, where it exists, is mostly due to the correlation between prices and the unexpected component of flows. We discuss the results in the main text; the detailed numerical results, including tables presenting the coefficients from our regressions, are available in the unpublished appendix.

Specifically, for each contract expiry, we perform the following procedure. We use regressions to estimate the expected and unexpected components of flows for both Asset Manager and Leveraged Funds (one regression for each type of flow), as is done in the mutual fund flow literature. The regressions include 5 lagged values of VIX changes and 5 own lags of flows. Predicted values from these first stage regressions are taken to be the "Expected" component, and the residuals are taken as the "Unexpected" component. Given this decomposition, we estimate the following regression for each contract expiry:

$$\Delta(v_i - v_0) = \alpha + \beta_1 \Delta\mathrm{AM}_t^{UNEXPECTED} + \beta_2 \Delta\mathrm{LF}_t^{UNEXPECTED} + error_t. \tag{4}$$

The first two contract expiries are expected to have the most relevance, based on the discussion above. The contemporaneous correlation of futures basis changes with the unexpected Asset Manager flow (β_1 in equation 4) has a t-statistic of 3.50 for the nearby contract, and it is 3.07 for the second nearby contract, but none of the t-statistics suggest significance for the unexpected Leveraged Fund flows. We conclude that the contemporaneous correlation between futures order flow by Asset Managers and changes in the futures basis is predominately due to the unexpected component of order flow and not due to a persistent effect that originated in prior days. While this evidence makes it clear that that the new information available on a given day is associated with the price change, it does not rule out simultaneous or dynamic effects on prices and volumes. We turn to those next.

We begin the examination of more dynamic models by including lagged variables on the right hand side of regression (3), thus relating today's futures basis change to both today's and yesterday's change in the VIX, change in Asset Manager net position, and the change in Leveraged Fund net position. If the lagged values of the flows show significantly positive coefficients, it would be evidence for "positive feedback". Negative coefficients suggest a reversal of the impact, and a zero coefficient would suggest no intertemporal relation at all.

Full details of the regressions for each of the six contracts are in the appendix; we focus on the most salient coefficients in this discussion. We continue to see significant coefficients on contemporaneous Asset Manager flows for the first three contracts. For the first contract, there is strong evidence for a price reversal effect related to the previously identified contemporaneous impact: the coefficient on the lagged Asset Manager flow is negative and about half the magnitude of the contemporaneous coefficient, and the t-statistic is -2.20. The analogous coefficient is negative and has a t-statistic of -1.85 for the second contract. There is little evidence for a relation, contemporaneous or lagged, for Leveraged Fund flows. We conclude from these regressions that there is no evidence of an explosive feedback effect, but there is evidence for some dynamic effects between net position changes and the futures basis changes, consistent with the price pressure hypothesis.

We further explore this finding by using vector autoregressions to trace out the dynamic impact of an unexpected shock to order flow on futures prices. For each contract expiry, we estimate four variable vector autoregressions that include three lags of each variable. The variables are 1) the VIX, 2) the daily flow of Asset Manager positions, 3) the daily flow of Leveraged Fund positions, and 4) the VIX futures basis (futures price minus VIX). We examine the impulse responses and variance decompositions for these VARs and confirm that there is little if any feedback from order flow to prices and back. Comprehensive VAR results, including impulse response charts for all of the variables and expiries, and variance decomposition results, are displayed in the unpublished appendix.

Given the previous regression results, we next address the dynamic response of the futures basis to an unexpected shock to the Asset Manager flow. Figure 5 displays the impulse response functions, computed by placing the Asset Manager flow second in the ordering, behind the VIX. For a one standard deviation shock to the Asset Manager flow the contemporaneous impact is 0.06 volatility points, and this impact quickly dies out. By the next day, the 95% Monte Carlo confidence interval around the impact includes zero (but just barely). Placing the Asset Manager flow lower in the ordering produces even less evidence for any lingering effect of order flow.

We also use the VARs to compute a variance decomposition for the VIX futures basis. The relative lack of importance of the flow variables to explain the dynamics of the basis is evident in the results. The VIX accounts for 60 to 80% of the variance for the front two contracts and 90%+ in the contracts further out. The futures basis accounts for nearly 40% of the variance for the front month contract and declines for longer dated contracts. The net position changes, summed across the two participant types, accounts for less than 3% of the variance at most. The numbers do not dramatically change as one moves further out the forecasting horizon.

We have explored the relation between net position changes and the VIX futures basis from several angles. Based on the analysis described above, we conclude that the regression model incorporating only the contemporaneous shock to flows is appropriate for further analysis. We see no compelling evidence for an explosive “feedback effect”; the compelling evidence suggests that the positive, contemporaneous correlation between Asset Manager flows and the futures basis is not permanent. The relation damps out quickly, which is more supportive of a quickly vanishing “price pressure” effect than an information effect.

3.2.3 Final Model.

Based on the results of these regressions, we are motivated to include both level and difference effects of end-user positions in our final model. Table 5 displays the results of the model, which has the specification

$$v_i - v_0 = \alpha_i + \beta_i v_0 + \gamma_{1i}(AM_i) + \gamma_{2i}(LF_i) + \delta_{1i}(\Delta AM_i) + \delta_{2i}(\Delta LF_i) + \varepsilon_i. \quad (5)$$

The fit of the models compares favorably with the levels regressions from Table 3, with adjusted R^2 values approximately the same or slightly higher. The significance of the levels variables (coefficients shown under the header "Positions") generally obtains for both Asset Managers and Leveraged Funds. There is strong evidence for an "Asset Manager" effect for all six contracts, and the evidence is strong for the first four contracts for "Leveraged Funds". The significance of the first differenced, or flow, variables (coefficients shown under the header "Position Changes") shows up only for the front part of the curve, which is where conventional wisdom puts the impact. Interestingly, this specification yields significance for the Leveraged Fund position change variables near the front end, but not for the Asset Manager variables. The negative coefficients for Leveraged Funds, which are on average shorting the contracts that the Asset Managers are going long is consistent with the hypothesized positive relation between order flow and the futures price.[8] The regressions strongly support a link between contract rolls and the steepness of the VIX futures term structure. In the next section, we explore the economic relevance of this model by examining counterfactuals implied by it.

4. ECONOMIC RELEVANCE

On average, the VIX futures term structure sloped steeply upward during our sample period. The settlement price of the front month contract averaged 18.20, with each more distant contract displaying a higher average value out to the sixth contract, which averaged 21.70 (VIX futures prices can be interpreted as annualized volatility levels). The worry by some is that this steep upward slope is abnormal and is driven by, or at least exacerbated by, the large long

[8] The unpublished appendix provides summary statistics on the two position change variables. For example, the correlation between the two types of flows is -0.73 and -0.80 for the front two contracts, respectively.

positions of Asset Managers. Based on the estimated models, we ask “How much of this steep upward slope is explained by positioning?”

4.1 Counterfactual Analysis

We compute the counterfactual futures prices by assuming zero positions and zero position changes by end-users at each date in the sample (*i.e.*, the counterfactual is computed as $v_i - v_0 = \alpha_i + \beta_i v_0$). Table 6 displays the average counterfactual values implied by the model, the average effects due to the positioning variables, and the resulting average VIX term structure over the sample. To be clear, we are not asserting that the counterfactuals are “correct” prices and the observed prices are “incorrect”, nor are we making any claims that the observed prices are distorted. The goal is simply to use the estimated model to isolate the component of futures prices related to specific types of transaction activity.

The observed average value for each contract is displayed in the final column of Table 6. The counterfactual price (displayed in the second column) also slopes upward. The counterfactual term structure is more concave, the observed term structure is more linear, and the two have the same values in the middle of the curve. More specifically, the counterfactual front and second month contracts are below the observed values (shown in the last column of the table), the peak of the counterfactual term structure is at contract 4, which is virtually identical to the observed value, and the longest dated counterfactuals are below the observed average values by approximately one volatility point.

The middle two columns of the table (“Level Effect” and “Roll Effect”) break out the estimated impact of the actual positions for each contract. The “Level Effect” value is the average, combined effect of the Asset Managers and the Leveraged Funds positions at each date. The “Roll Effect” is the average, combined effect of the position changes of the two participant types. For convenience, the sum of these two values is shown in the next to last column (“Net Effect”). For each date, this predicted net effect is the sum of the products of the estimated coefficients multiplied by the observed variables.

We find that the Level Effect is generally larger than the Roll Effect, with the Roll Effect less than 0.1 vol points for the third through the sixth contract. While the Level Effect appears to be statistically significant at a very high confidence level, the Roll Effect is not significant for the contracts furthest out.

This breakout highlights the relative impact of the Level Effect versus the Roll Effect: even when the rolling of positions impacts the futures price in the front two contracts, the impact is small compared to the Level Effect of aggregate position sizes. For the front month contract, the Level Effect acts to increase the futures price by 0.36 vol points, but the Roll Effect partially offsets this by 0.14 vol points, with the Net Effect just 0.22 vol points. A larger impact is actually observed in the fifth and sixth contracts, with the model suggesting impacts of approximately one volatility point.

We use the estimated model to attribute the differences between the observed data and the counterfactual data in a second breakout. Whereas the previous table attributed the differences to a Level Effect and a Roll Effect, Figure 6 displays the average estimated net impact attributed to non-dealer participant types. It also shows, in addition to these two impacts, the net impact of these two.

The results tell a simple story for this sample. The average long positions of Asset Managers acted to increase the futures price, at the margin, but the average short positions of Leveraged Funds acted to decrease the futures price for every contract. The net impact varies according to the strength of the two effects, which reflects the net size of the two positions. In the front month, Asset Managers dominated slightly, leading to the net impact of +0.22 vol points noted previously. In the third contract, very few net positions are held within these two groupings, leading to a minimal impact on the curve. In longer dated contracts, where liquidity is undoubtedly lower, the modest Asset Manager net long positions dominated, leading to a predicted average impact of approximately one volatility point.

We compare our results with those of Bollen *et al.* (2013). Our model suggests an estimated upward pressure of impact of on the front month contract, due to Asset Managers, of just under half a vol point. This is partially offset by short positons of Leveraged Funds, with an average net impact of approximately a quarter of a volatility point.. Bollen *et al.* find that the distribution of VIX futures prices relative to VIX peaks at 2.5% since VIX ETNs were introduced. Given an average VIX level of 22.9 for their sample, their analysis suggests that long demand for VIX futures by ETP hedgers added 0.6 vol points (22.9 x 2.5%) to the price of VIX futures. By that comparison, our results are quite similar to those of Bollen *et al*. We find, and they find, a statistically significant impact of non-dealer long futures demand on prices that is approximately half a vol point on average in recent years. We also find, however, an offsetting impact by non-dealer short futures demand, which they do not analyze. The net short demand by non-dealers becomes most influential for the third contract, where Asset Manager's net demand is the smallest. Further, we find even larger net impacts on the longer-dated contracts expiring five to six months out.

4.2 Relation to No-Arbitrage Bounds

We have not evaluated the economic significance of the estimated impacts described so far. Is half a vol point of price pressure economically large? One measure of the significance is to compare it to the bid-ask spread for VIX futures. Over the sample period, the median end-of-day bid-ask spread has been 0.05 vol points for all six contracts, which might make one conclude that the impact is quite large.

On the other hand, we argue that the more relevant action is to compare the estimated impact to the no-arbitrage bounds for the futures price. Carr and Wu (2006) show that the VIX futures should be bounded below by the at-the-money-forward implied volatility and above by the volatility strike of the forward variance swap. Both of these variables are observable.

Figure 7 displays the width of this range and the predicted impact from the models for the front month contract. The no-arbitrage range is computed using the one- and two-month constant maturity S&P 500 implied volatilities and the one- and two-month indicative variance swap rates, both from Bloomberg. For clarity, the chart displays 20-day moving averages of the impact and the bounds.

Our first conclusion from examining the chart is that the estimated impact due to non-dealer positions is usually well within the no-arbitrage range. The range is far wider than the 0.05 vol point bid-ask spread, and it varies over time (averaging around two vol points wide during this period). When measured by this standard, the impact due to net non-dealer futures demand is not large.

Our second conclusion from examining the chart is that the impact itself varies over time as the relative positions of the two participant types changes. The conclusion is facilitated by examination of Figure 8, which displays the net positions of the three participant categories. The price impact declined sharply in August 2011 as Asset Managers rapidly reduced their positions and dealers increased their positions as they absorbed the other side. It fluctuated around one vol point during the September 2011 – September 2012 period, when Asset Manager longs were larger in magnitude than Leveraged Funds shorts (with dealers averaging -19,000 contracts short as a consequence). It was approximately zero over the September 2012-December 2014 period when Leveraged Funds increased their short positions and dealers moved net long as a consequence.

5. CONCLUSION

VIX futures represented nearly 20% of CBOE's annual transaction revenue in 2013, according to publicly available data, and contract volumes have grown at a dramatic pace each of the last several years. Variance swaps, a related product, are traded OTC and have historically featured the opaqueness common to smaller OTC markets. Volatility derivatives used to be science fiction, but today they are clearly a major part of derivatives activity. With this prominence comes concern that the activity is somehow distorting the market signals that the product should provide. We use unique regulatory data in this paper to examine open positions and activity in both the listed and OTC volatility space. We then go on to evaluate the impact of trading and positions in the most active volatility derivative of the S&P 500: VIX futures contracts.

We find that the stock of variance swaps is large for long-dated contracts, with S&P 500 instruments exhibiting USD 1.5 billion vega outstanding (a third of which is between dealers and end-users) and transactions occurring daily. Dealers are short vega in order to supply the long vega demands of asset managers. The market for variance swaps on other underlying indexes is somewhat smaller, with activity concentrated among the next five or so indexes before trailing off to less active markets. Total notional vega outstanding is over USD 2 billion, and we find over USD 150 million initiated in new swaps in a recent sample month. VIX futures are far more actively traded, with USD 3 billion of vega over the month, and all of the activity is in the front end of the curve.

We quantify two distinct channels by which VIX trading activity manifests itself into futures prices. First, we identify a risk-based effect that we identify with end-users tending to be net long VIX futures and dealers tending to be short VIX futures. To the extent that dealers take on risk when facilitating trades, we estimate that the long volatility bias of asset managers acts to put upward pressure on VIX futures prices. Hedge funds offset this theoretical impact because they have actively taken a net short aggregate futures position in our sample. For the front-month contract, the net upward impact of long non-dealer activity on prices, on average, is approximately a quarter of a volatility point. At the longer end of the futures curve (e.g., contracts expiring in 5-6 months), non-dealer positioning has a long bias that results in upward pressure of approximately one volatility point. Second, we identify an order flow, or "price pressure" effect: as end-users transact large quantities (primarily in the front two contracts), the prices of those contracts exhibit transient price changes that are quickly reversed. Nonetheless, we find no evidence that the aforementioned price impacts force VIX futures outside of their no-arbitrage bounds.

We conclude that end-user demand for volatility products exerts a measurable influence on VIX futures prices. The impact at a given point in time is dependent on non-dealer supply and demand for VIX futures, which each vary over time. In our 2011-2015 sample period, the net impact added approximately a quarter of a volatility point in nearby contracts, on average, but approximately one volatility point for contracts in less liquid, longer-dated parts of the curve.

Table 1.
Summary Statistics for Variance Swap Activity (May 2014).

	S&P 500	Euro-Stoxx 50	Russell 2000	FTSE 100	Nikkei 225	Kospi 200	Other	Total
Total New Swaps (USD Vega Millions)	**85.2**	**23.4**	**4.8**	**11.2**	**4.7**	**8.7**	**21.5**	**159.4**
SD/SD	47.9	11.1	2.8	0.5	1.8	1.6	3.3	68.9
SD/non-SD	37.3	12.3	2.0	10.7	2.9	7.1	18.2	90.5
New Swaps Total (Count)	**694**	**193**	**57**	**125**	**44**	**96**	**175**	**1,384**
SD/SD	362	69	28	5	12	12	50	538
SD/non-SD	332	124	29	120	32	84	125	846
Original Tenor (Years)	**1.1**	**0.6**	**1.2**	**0.3**	**0.8**	**0.9**	**0.8**	**0.9**
SD/SD	1.3	0.8	1.4	0.4	1.2	1.3	1.2	1.2
SD/non-SD	0.8	0.4	1.1	0.3	0.7	0.9	0.6	0.7
Average Swap Notional (USD Vega Thousands)	**123**	**121**	**84**	**90**	**106**	**90**	**123**	**115**
SD/SD	132	161	99	99	149	130	65	128
SD/non-SD	112	99	70	89	90	85	146	107

Figure 1.
Variance Swap Activity (May 2014), cumulative by tenor.

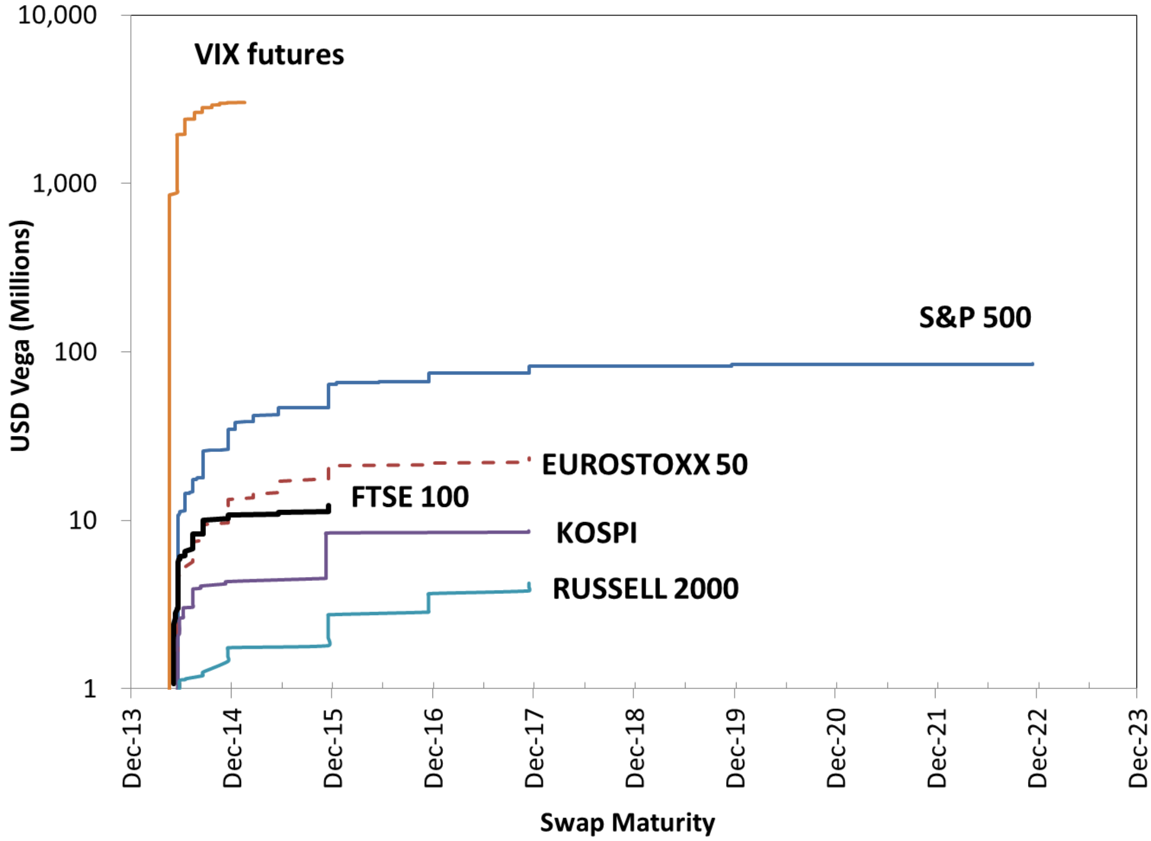

Table 2.
Summary Statistics for Variance Swaps Outstanding (May 2014).

	S&P 500	Euro-Stoxx 50	Russell 2000	FTSE 100	Nikkei 225	Kospi 200	Other	Total
Total Outstanding (USD Vega Millions)	**1,512**	**205**	**81**	**66**	**51**	**81**	**166**	**2,164**
SD/SD	1,084	112	57	40	15	37	65	1,411
SD/non-SD	428	93	24	26	37	44	102	753
Open Swaps Total (Count)	**7,729**	**1,146**	**583**	**568**	**514**	**537**	**1,237**	**12,314**
SD/SD	5,293	532	450	326	168	165	553	7,487
SD/non-SD	2,463	614	133	242	346	372	684	4,827
Original Tenor (Years)	**3.9**	**2.0**	**4.0**	**2.3**	**1.3**	**1.8**	**1.9**	**3.2**
SD/SD	4.3	2.3	4.3	3.0	1.5	2.9	2.7	3.9
SD/non-SD	2.9	1.7	2.8	1.4	1.2	1.3	1.2	2.2
Top 5 Dealer Share by Vega (SD/non-SD trades)	**70.4%**	**80.6%**	**83.6%**	**86.1%**	**78.8%**	**92.7%**	**79.7%**	**68.0%**

Figure 2.
Net Vega Notional Outstanding by Participant Type and Underlying Index. The data represent positioning at the end of May 2014.

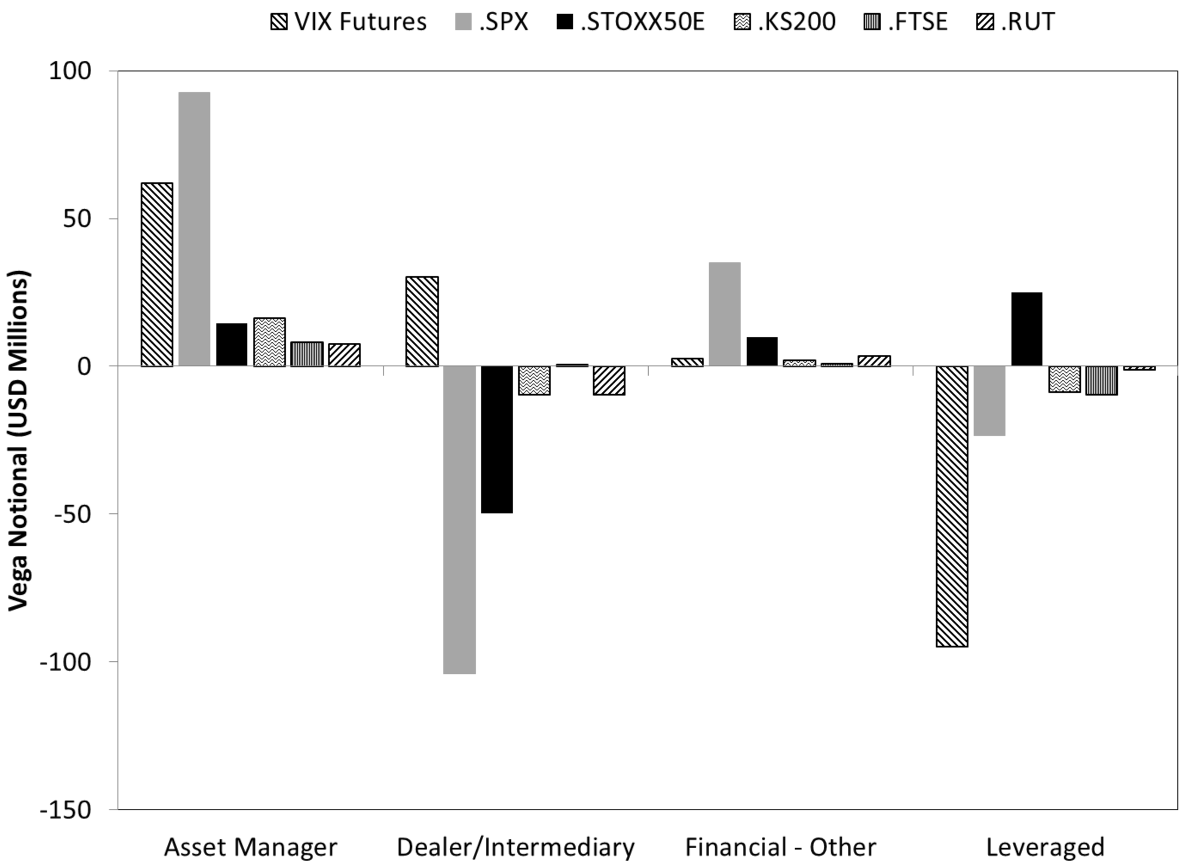

Figure 3.
Vega Notional Outstanding, May 2014 (VIX Futures and S&P 500 Variance Swaps). The data represent positioning at the end of May 2014.

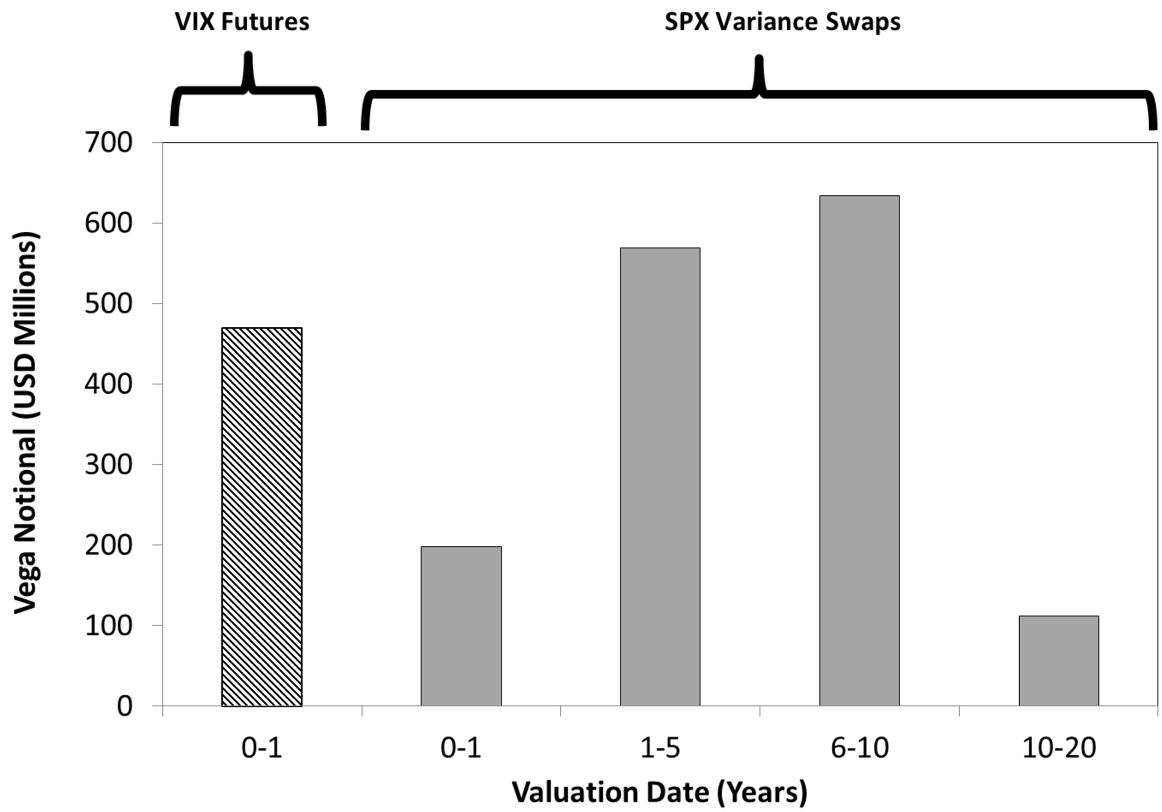

Figure 4.
Net Vega Notional Outstanding, VIX futures and S&P 500 Variance Swaps, by Participant Type and Time to Maturity. The data represent positioning at the end of May 2014.

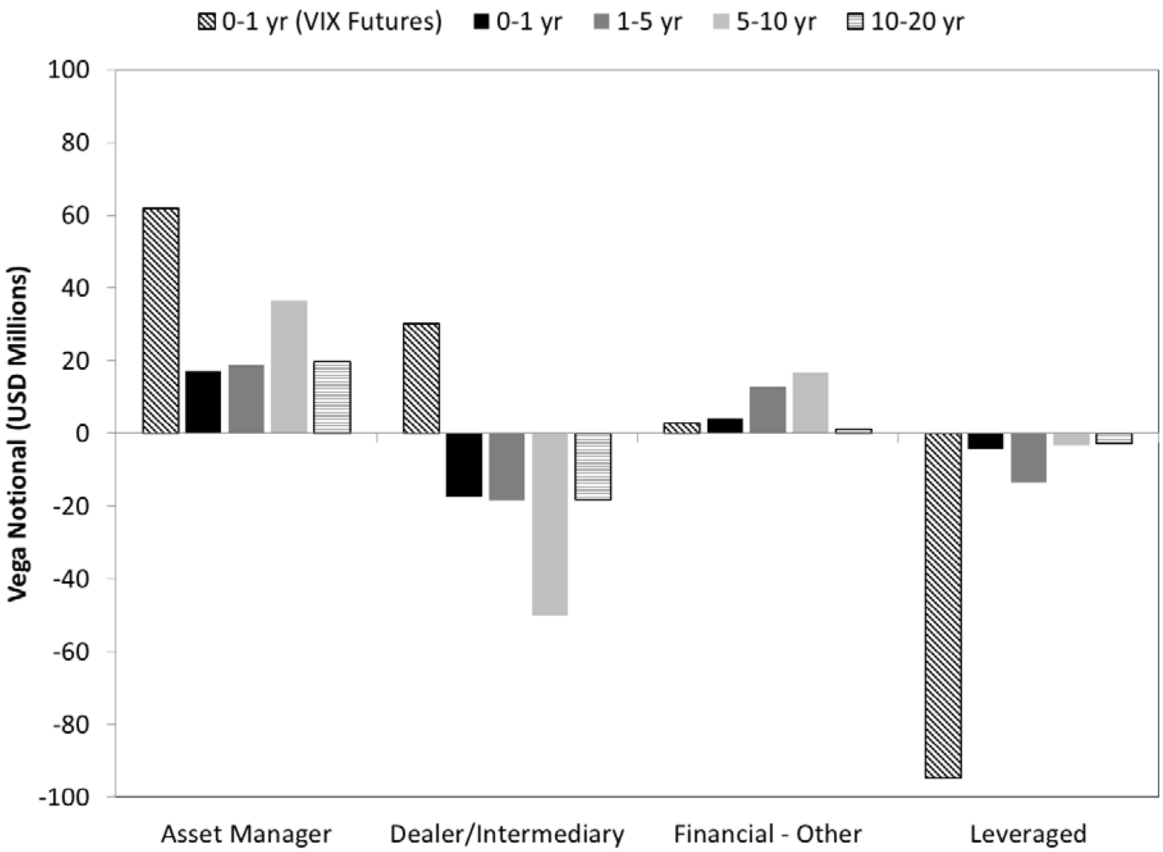

Table 3.
Results from regressing VIX futures basis on VIX and dealer and non-dealer VIX futures positions.
The table displays estimation results for the regressions

$$v_i - v_0 = \alpha_i + \beta_i v_0 + \gamma_{1i}(Dealer_i) + \varepsilon_i \text{ (in Panel A)}$$

$$v_i - v_0 = \alpha_i + \beta_i v_0 + \gamma_{1i}(AM_i) + \gamma_{2i}(LF_i) + \varepsilon_i \text{ (in Panel B).}$$

Regressions are estimated separately for each contract (*i* = 1 to 6). Shaded boxes represent statistical significance at 5% level. T-statistics are based on Newey-West (1987) standard errors with 3 lags. Columns marked "AM" reflect coefficients on Asset Manager variables. Columns marked "LF" reflect coefficients on Leveraged Fund variables. The regressions are estimated on daily data spanning the period April 2011 to March 2015.

	Panel A			Panel B			
Contract Expiry	VIX	Dealer	Adj. R^2 (%)	VIX	AM	LF	Adj. R^2 (%)
1	-0.118	-0.022	23.6	-0.115	0.034	0.021	31.5
	(-5.09)	(-4.84)		(-5.08)	(6.99)	(4.85)	
2	-0.256	-0.052	42.2	-0.246	0.057	0.051	42.7
	(-7.49)	(-6.51)		(-7.23)	(6.86)	(6.42)	
3	-0.318	-0.097	50.1	-0.307	0.089	0.088	47.9
	(-9.40)	(-7.20)		(-9.04)	(3.15)	(6.09)	
4	-0.400	-0.108	49.3	-0.371	0.247	0.070	53.6
	(-10.25)	(-5.03)		(-10.81)	(7.17)	(3.37)	
5	-0.368	-0.001	42.5	-0.392	0.315	0.002	56.3
	(-10.51)	(-0.01)		(-13.16)	(8.83)	(0.06)	
6	-0.392	-0.092	44.4	-0.415	0.426	-0.010	61.9
	(-11.29)	(-3.15)		(-13.41)	(13.64)	(-0.49)	

Table 4.
Results from regressing VIX futures basis changes on VIX and non-dealer VIX futures position changes.

The table displays estimation results for the regressions

$$\Delta(v_i - v_0) = \alpha_i + \beta_i \Delta v_0 + \delta_{1i}(\Delta AM_i) + \delta_{2i}(\Delta LF_i) + \varepsilon_i.$$

Regressions are estimated separately for each contract (*i* = 1 to 6). Shaded boxes represent statistical significance at 5% level. T-statistics are based on Newey-West (1987) standard errors with 3 lags. Columns marked "AM" reflect coefficients on Asset Manager variables. Columns marked "LF" reflect coefficients on Leveraged Fund variables. The regressions are estimated on daily data spanning the period April 2011 to March 2015.

Contract Expiry	ΔVIX	(t-stat)	ΔAM	(t-stat)	ΔLF	(t-stat)	Adj. R^2 (%)
1	-0.369	(-10.09)	0.027	(3.34)	0.012	(1.73)	58.7
2	-0.538	(-14.37)	0.024	(3.67)	0.020	(3.14)	81.7
3	-0.639	(-25.91)	0.034	(3.06)	0.002	(0.25)	90.4
4	-0.700	(-29.47)	0.011	(0.70)	-0.019	(-1.42)	93.0
5	-0.747	(-30.43)	0.039	(1.54)	-0.017	(-1.06)	94.1
6	-0.768	(-36.32)	0.045	(1.86)	-0.014	(-0.70)	94.9

Figure 5.

Impulse response analysis of VIX futures basis (futures minus VIX) for front two contracts.

The VARs are estimated separately for each contract and include the VIX, Asset Manager net position change, Leveraged Fund net position change, and the VIX futures basis (in that order). VARs include three lags of each variable and are estimated using daily data spanning the period April 2011 to March 2015. Monte Carlo 95% confidence intervals are also shown. Impulse response steps are daily.

Response to

Contract Expiry | **Shock to Asset Manager Flow** | **Shock to Leveraged Fund Flow**

1

2

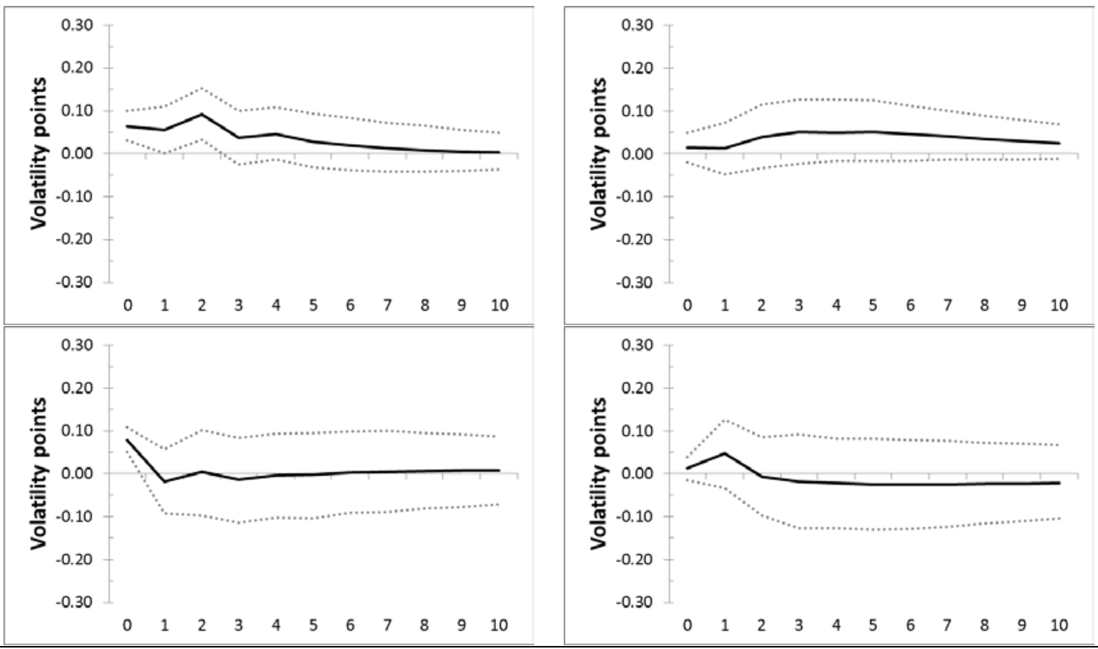

Table 5.
Results from regressing VIX futures basis on VIX and non-dealer VIX futures positions and position changes.
The table displays estimation results for the regressions

$$v_i - v_0 = \alpha_i + \beta_i v_0 + \gamma_{1i}(AM_i) + \gamma_{2i}(LF_i) + \delta_{1i}(\Delta AM_i) + \delta_{2i}(\Delta LF_i) + \varepsilon_i.$$

Regressions are estimated separately for each contract (*i* = 1 to 6). Shaded boxes represent statistical significance at 5% level. T-statistics are based on Newey-West (1987) standard errors with 3 lags. Columns marked "AM" reflect coefficients on Asset Manager variables. Columns marked "LF" reflect coefficients on Leveraged Fund variables. The regressions are estimated on daily data spanning the period April 2011 to March 2015.

	Positions				Position Changes				
Contract Expiry	AM	(t-stat)	LF	(t-stat)	ΔAM	(t-stat)	ΔLF	(t-stat)	Adj. R^2 (%)
1	0.031	(6.69)	0.016	(3.72)	-0.011	(-0.74)	-0.076	(-5.55)	33.6
2	0.051	(6.18)	0.049	(6.17)	0.026	(0.91)	-0.065	(-2.35)	44.5
3	0.090	(2.98)	0.088	(5.97)	-0.143	(-1.85)	-0.162	(-2.60)	47.9
4	0.243	(6.91)	0.070	(3.28)	-0.354	(-2.56)	0.112	(1.21)	54.0
5	0.306	(8.23)	-0.008	(-0.23)	0.164	(0.80)	0.155	(1.18)	56.3
6	0.428	(13.45)	-0.010	(-0.46)	-0.240	(-1.85)	-0.065	(-0.56)	62.0

Table 6.
Estimated non-dealer net demand impact on average VIX futures prices.
Average effects and their significance are estimated by computing the relevant regression coefficients multiplied by the sample average position and change in position data, when the averages are assumed observed without error. Shaded boxes represent statistical significance at 5% level. T-statistics are based on Newey-West (1987) standard errors with 3 lags. The regressions are estimated on daily data spanning the period April 2011 to March 2015.

Contract	Counterfactual Average Price	Level Effect (t-statistic)	Roll Effect (t-statistic)	Net Effect (t-statistic)	Observed Average Price
1	17.97	0.36 (7.51)	-0.14 (-5.93)	0.22 (4.64)	18.20
2	18.76	0.10 (1.65)	0.21 (5.99)	0.31 (5.06)	19.11
3	20.01	-0.06 (-4.10)	-0.07 (-2.66)	-0.13 (-4.17)	19.94
4	20.62	-0.13 (-1.57)	0.04 (2.34)	-0.09 (-1.01)	20.55
5	20.36	0.79 (7.07)	-0.02 (-1.15)	0.78 (7.16)	21.16
6	20.55	1.11 (12.67)	0.00 (1.71)	1.11 (12.65)	21.70

Figure 6.

Estimated average impact, by contract expiry, of VIX futures positions and daily position changes by Asset Managers, Leveraged Funds, and the net impact of the two. The estimates are constructed using the final model in equation (5). For each contract, the impact figure displays the sum of the relevant regression coefficients multiplied by the position levels and daily net change values. The regressions and the average values are computed using daily data spanning the period April 2011 to March 2015.

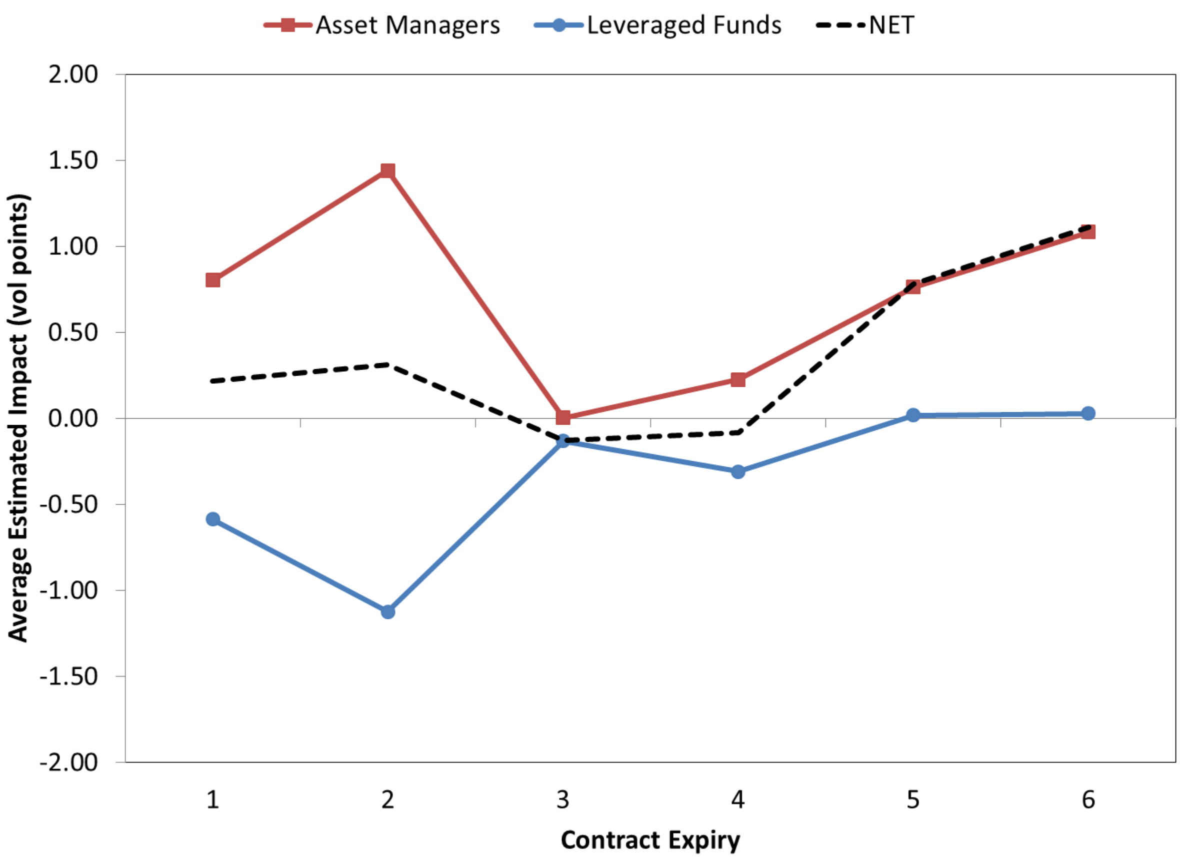

Figure 7.

Estimated impact of non-dealer net demand on VIX futures front month contract prices and estimated arbitrage bounds. The estimated impact is computed using the final model in equation (5), and the arbitrage bounds are constructed using daily S&P 500 variance swap and at-the-money implied volatilities from Bloomberg. All series are smoothed using 20 trading day moving averages to improve clarity of the figure.

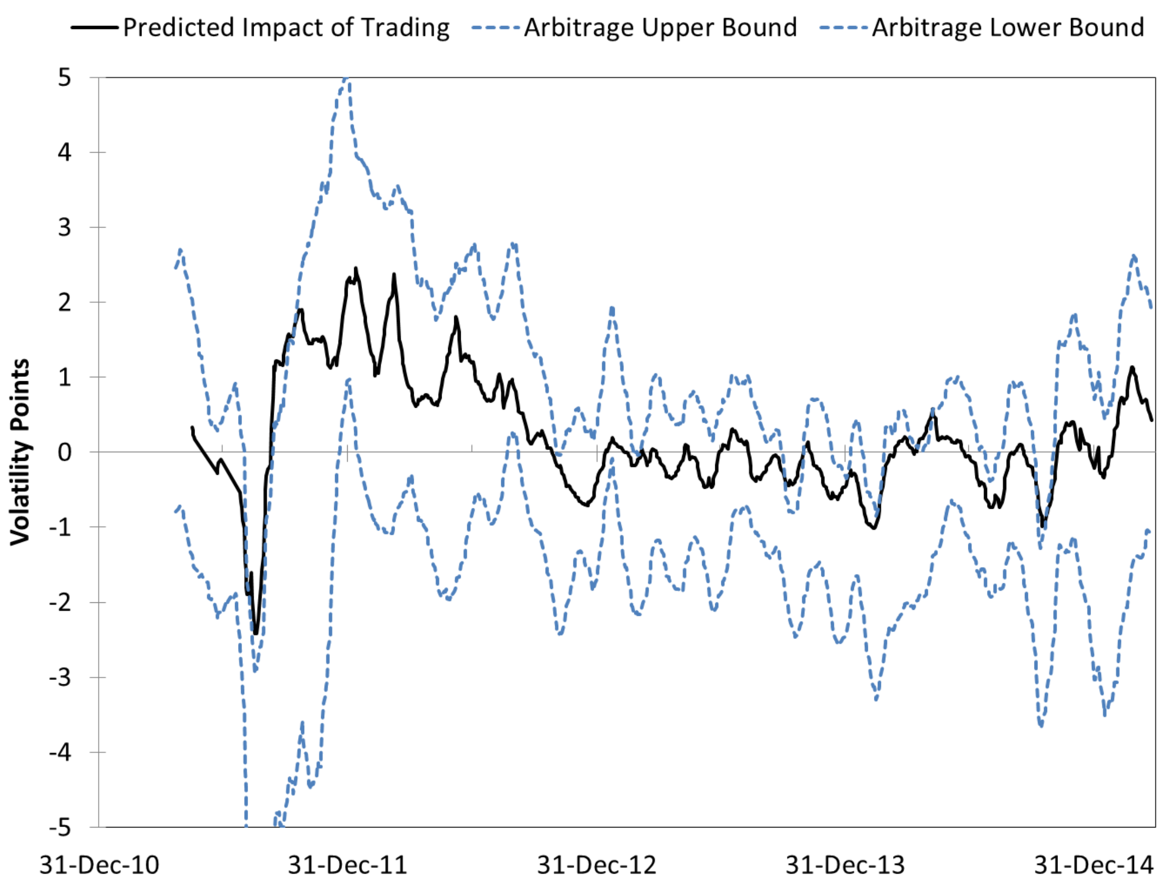

Figure 8.

Net positioning in VIX futures by Asset Managers, Leveraged Funds, and Dealers. The figure displays the net positions, aggregated within each of the Asset Manager, Dealer, and Leveraged Fund + other categories. Aggregates are constructed using the first six contracts and use daily data from April 2011 to March 2015.

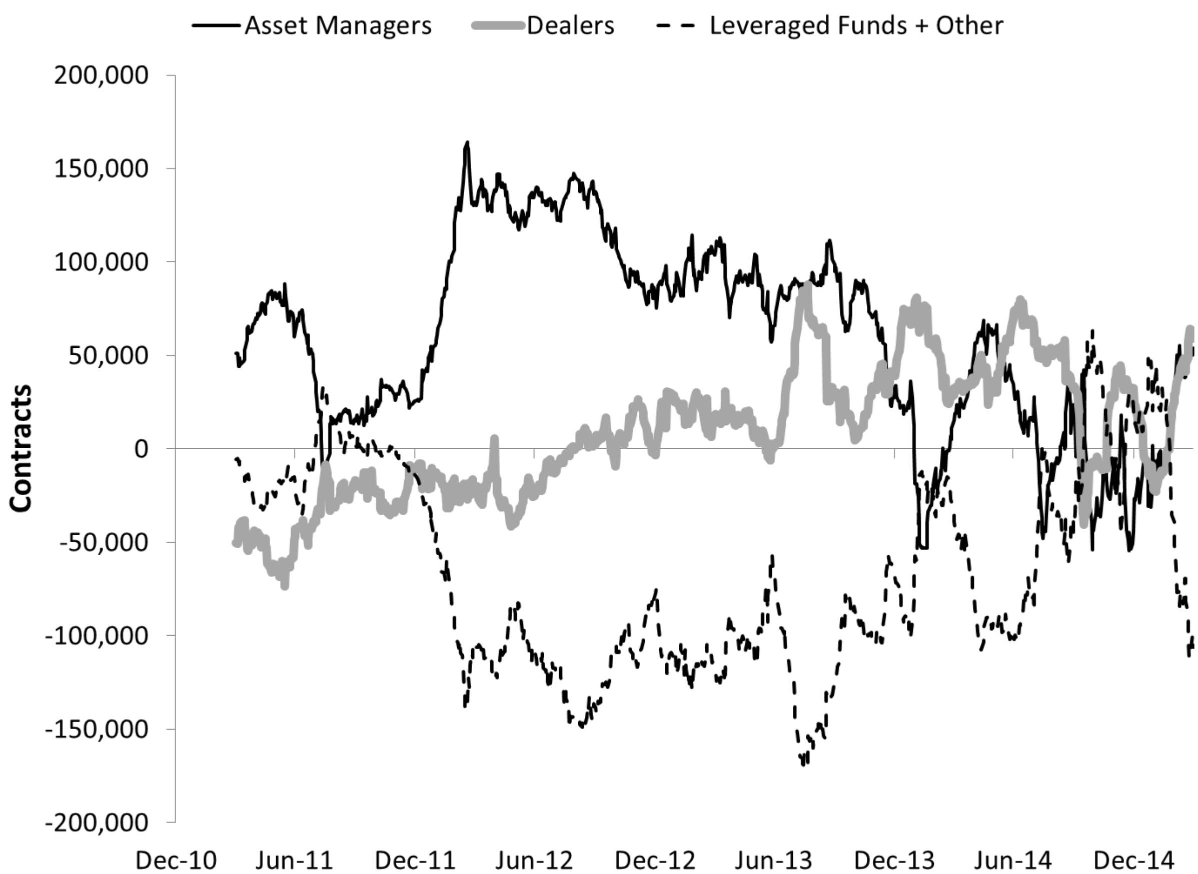

REFERENCES

Ait-Sahalia, Yacine, Mustafa Karaman, and Loriano Mancini "The Term Structure of Variance Swaps, Risk Premia and the Expectations Hypothesis," 2014. Working paper.

Ben-Rephael, Azi, Shmuel Kandel, and Avi Wohl, "The Price Pressure of Aggregate Mutual Fund Flows." Journal of Financial and Quantitative Analysis, Vol. 46, No. 2, (2011), pp. 585-603.

Bollen, N., M. J. O'Neill, and R. E. Whaley. "On the Supply of and Demand for Volatility." (2013) Working paper Vanderbilt University.

Bollen, N., and R. E. Whaley. "Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?" Journal of Finance, Vol. 59 (2004), pp. 711-753.

Carr, Peter, and Liuren Wu. "A Tale of Two Indices." Journal of Derivatives, Vol. 13, No. 3 (2006), pp. 13-29.

Carr, Peter, and Liuren Wu. "Variance Risk Premiums." Review of Financial Studies Vol. 22, No. 3 (2009), pp. 1311-1341.

Corradi, Valentina, Walter Distaso, and Antonio Mele. "Macroeconomic Determinants of Stock Volatility and Volatility Premiums." Journal of Monetary Economics, Vol. 60, No. 2 (2013), pp. 203-220.

David, Alexander, and Pietro Veronesi. "Investors' and Central Bank's Uncertainty Embedded in Index Options." Review of Financial Studies Vol. 27, No. 6 (2014) pp. 1661-1716.

Edelen, Roger, and Jerold Warner, "Aggregate Price Effects of Institutional Trading: A Study of Mutual Fund Flows and Market Returns." Journal of Financial Economics, Vol. 59 (2001), pp. 195-220.

Egloff, Daniel, Markus Leippold, and Liuren Wu. "The Term Structure of Variance Swap Rates and Optimal Variance Swap Investments." Journal of Financial and Quantitative Analysis Vol. 45, No. 5 (2010), pp. 1279-1310.

Garleanu, Nicolae, Lasse Heje Pedersen, and Allen M. Poteshman. "Demand-Based Option Pricing." Review of Financial Studies, Vol. 22, No. 10 (2009), pp. 4259-4299.

Glatzer, Ernst, and Martin Scheicher. "What Moves the Tail? The Determinants of the Option-Implied Probability Density Function of the DAX Index." Journal of Futures Markets, Vol. 52, No. 6 (2005), pp. 515-536.

Han, Bing. "Investor Sentiment and Option Prices." Review of Financial Studies, Vol. 21, No. 1 (2008), pp. 387-414.

Mixon, Scott "Factors Explaining Movements in the Implied Volatility Surface." Journal of Futures Markets, Vol. 22, No. 10 (2002), pp. 915-937.

Mixon, Scott. "The Implied Volatility Term Structure of Stock Index Options." Journal of Empirical Finance, Vol. 14, No. 3 (2007), pp. 333-354.

Naik, N.K., and Pradeep Yadav. "Risk Management with Derivatives by Dealers and Market Quality in Government Bond Markets." Journal of Finance, Vol. 58, No. 5 (2003), pp. 1873-1904.

Newey, W.K., and K.D. West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." Econometrica Vol. 55 (1987), pp. 703-708.

Nossman, M., and Wilhemsson, A. "Is the VIX Futures Market Able to Predict the VIX Index? A Test of the Expectation Hypothesis." The Journal of Alternative Investments, Vol. 12, No. 2 (2009), pp. 54-67.

Simon, D.P., and J. Campasano "The VIX Futures Basis: Evidence and Trading Strategies." Journal of Derivatives Vol. 21, No. 3 (2014), pp. 54-69.

Vähämaa, Sami, and Janne Äijö. "The Fed's Policy Decisions and Implied Volatility." Journal of Futures Markets, Vol. 31, No. 10 (2011) pp. 995-1010.

Warther, Vincent A. (1995) "Aggregate Mutual Fund Flows and Security Returns." Journal of Financial Economics Vol. 39 (1995) pp. 209-35.

APPENDIX
(NOT FOR PUBLICATION)

Table A1.
Results from regressing VIX futures basis on VSTOXX and lagged dealer and non-dealer VIX futures positions.
The table displays estimation results for the regressions

$$v_i - v_0 = \alpha_i + \beta_i VSTOXX + \gamma_{1i}\left(Dealer_{i,t-5}\right) + \varepsilon_i \text{ (in Panel A)}$$

$$v_i - v_0 = \alpha_i + \beta_i VSTOXX + \gamma_{1i}\left(AM_{i,t-5}\right) + \gamma_{2i}\left(LF_{i,t-5}\right) + \varepsilon_i \text{ (in Panel B).}$$

Regressions are estimated separately for each contract (*i* = 1 to 6). Shaded boxes represent statistical significance at 5% level. T-statistics are based on Newey-West (1987) standard errors with 3 lags. Columns marked "VSTOXX" reflect coefficients on the VSTOXX, the VIX-equivalent for the EURO STOXX 50 index, computed at the European market close. Columns marked "AM_{t-5}" reflect coefficients on Asset Manager variables, lagged 5 days. Columns marked "LF_{t-5}" reflect coefficients on Leveraged Fund variables, lagged 5 days. Data are daily and span the period April 2011 to March 2015.

	Panel A			Panel B			
Contract Expiry	VSTOXX	$Dealer_{t-5}$	Adj. R^2 (%)	VSTOXX	AM_{t-5}	LF_{t-5}	Adj. R^2 (%)
1	-0.069	-0.017	11.4	-0.070	0.020	0.018	11.5
	(-3.82)	(-3.95)		(-3.87)	(4.02)	(4.07)	
2	-0.165	-0.045	24.6	-0.145	0.050	0.035	28.3
	(-5.35)	(-5.12)		(-4.85)	(5.63)	(4.03)	
3	-0.213	-0.124	38.3	-0.203	0.108	0.118	35.8
	(-6.59)	(-8.14)		(-6.26)	(3.29)	(7.12)	
4	-0.279	-0.117	32.4	-0.244	0.291	0.064	39.2
	(-6.98)	(-4.09)		(-9.61)	(6.92)	(2.38)	
5	-0.224	0.059	26.1	-0.255	0.285	-0.057	41.1
	(-7.25)	(1.03)		(-9.61)	(5.36)	(1.09)	
6	-0.253	-0.107	27.6	-0.291	0.463	0.008	46.0
	(-6.96)	(-3.24)		(-8.93)	(11.16)	(0.33)	

Table A2.
Summary statistics for non-dealer, daily net position changes. Data are daily and span the period April 2011 to March 2015.

Contract Expiry	**Asset Managers** Mean (Standard Error)	**Leveraged Funds** Mean (Standard Error)	Correlation
1	-2,358.8 (3,760.1)	2,174.4 (3,963.5)	-0.73
2	2,184.9 (4,273.6)	-2,341.9 (4,017.0)	-0.80
3	127.6 (1,163.5)	298.9 (1,508.6)	-0.53
4	-129.3 (704.8)	-21.6 (921.6)	-0.49
5	-25.7 (504.7)	-72.5 (753.2)	-0.37
6	-9.2 (509.7)	2.4 (611.4)	-0.25

Table A3.
Results from regressing VIX futures basis changes on unexpected VIX and non-dealer VIX futures position changes.
For each VIX futures contract (i = 1 to 6), the regression (A1) is estimated:

$$\Delta(v_i - VIX_t) = \alpha + \beta_1 \Delta \mathrm{AM}_t^{UNEXPECTED} + \beta_2 \Delta \mathrm{LF}_t^{UNEXPECTED} + error_t. \tag{A1}$$

The unexpected components of flows used in (A1) are the residuals from first-stage regressions (A2) and (A3):

$$\Delta \mathrm{AM}_t = \alpha_1 + \sum_1^5 \delta_i^1 \Delta \mathrm{VIX}_{t-i} + \sum_1^5 \delta_i^2 \Delta \mathrm{AM}_{t-i} + error_t^1 \tag{A2}$$

$$\Delta \mathrm{LF}_t = \alpha_2 + \sum_1^5 \delta_i^3 \Delta \mathrm{VIX}_{t-i} + \sum_1^5 \delta_i^4 \Delta \mathrm{LF}_{t-i} + error_t^2. \tag{A3}$$

Shaded boxes represent statistical significance at 5% level. T-statistics are based on Newey-West (1987) standard errors with 3 lags. Columns marked "ΔAM" reflect coefficients on Asset Manager variables. Columns marked "ΔLF" reflect coefficients on Leveraged Fund variables. Data are daily and span the period April 2011 to March 2015.

Contract Expiry	ΔVIX	(t-stat)	$\Delta\mathrm{AM}^{UNEXPECTED}$	(t-stat)	$\Delta\mathrm{LF}^{UNEXPECTED}$	(t-stat)	Adj. R^2 (%)
1	-0.373	(-8.27)	0.034	(3.42)	0.010	(1.12)	62.5
2	-0.544	(-12.04)	0.018	(2.93)	0.002	(0.30)	83.7
3	-0.643	(-21.54)	0.026	(2.13)	0.000	(0.05)	92.1
4	-0.705	(-24.41)	-0.005	(-0.26)	-0.015	(-1.07)	94.1
5	-0.755	(-25.35)	0.024	(0.87)	-0.017	(-0.99)	95.1
6	-0.774	(-30.30)	0.011	(0.50)	-0.031	(-1.30)	95.8

Table A4.
Results from regressing VIX futures basis changes on contemporaneous and lagged VIX and non-dealer VIX futures position changes.
The table displays results from the regressions

$$\Delta(v_{i,t} - VIX_t) = \alpha_i + \sum_{j=0}^{1} \beta_j \Delta VIX_{t-j} + \sum_{j=0}^{1} \delta_{1i,t-j} \Delta AM_{i,t-j} + \sum_{j=0}^{1} \delta_{2i,t-j} \Delta LF_{i,t-j} + \varepsilon_i. \quad \text{(A4)}$$

I.e., the regression for each contract (*i* = 1 to 6) includes the contemporaneous and lagged value of the independent variables: the VIX, the daily change in net Asset Manager positions, and the daily change in net Leveraged Fund positions. Shaded boxes represent statistical significance at 5% level. T-statistics are based on Newey-West (1987) standard errors with 3 lags. Columns marked "ΔAM" reflect coefficients on Asset Manager variables. Columns marked "ΔLF" reflect coefficients on Leveraged Fund variables. Data are daily and span the period April 2011 to March 2015.

Contract Expiry	Lag	ΔVIX	(t-stat)	ΔAM	(t-stat)	ΔLF	(t-stat)	Adj. R^2 (%)
1	0	-0.358	(-10.09)	0.027	(3.34)	0.009	(1.26)	60.9
	1	0.046	(2.93)	-0.018	(-2.20)	-0.003	(-0.42)	
2	0	-0.532	(-13.67)	0.024	(3.57)	0.013	(1.96)	82.5
	1	0.027	(2.22)	-0.010	(-1.85)	0.010	(1.74)	
3	0	-0.635	(-24.36)	0.031	(2.77)	0.004	(0.47)	90.7
	1	0.014	(1.59)	0.002	(0.15)	-0.005	(-0.69)	
4	0	-0.696	(-28.15)	-0.002	(-0.10)	-0.026	(-2.02)	93.3
	1	0.016	(1.96)	-0.006	(-0.37)	0.008	(0.70)	
5	0	-0.744	(-29.58)	0.025	(1.06)	-0.018	(-1.14)	94.4
	1	0.015	(1.71	-0.008	(-0.29)	-0.019	(-1.29)	
6	0	-0.765	(-34.86)	0.040	(1.78)	-0.015	(-0.69)	95.1
	1	0.017	(2.06)	0.006	(0.27)	0.027	(1.31)	

Figure A1.

Impulse responses of VAR models of VIX futures prices.

Each VAR consists of four variables: the VIX, the daily net position change (flow) of Asset Managers, the daily net position change (flow) of Leveraged Funds, and the VIX futures basis (futures price minus VIX) price for that contract (1 to 6). Three lags for each variable are included in the models. VARs are estimated separately for each contract expiry (from 1 to 6), and are estimated using daily data spanning the period April 2011 to March 2015. Impulse responses of Asset Manager net position changes and Leveraged Fund net position changes are cumulative over the forecast horizon. Impulse responses are displayed with 95% Monte Carlo confidence intervals.

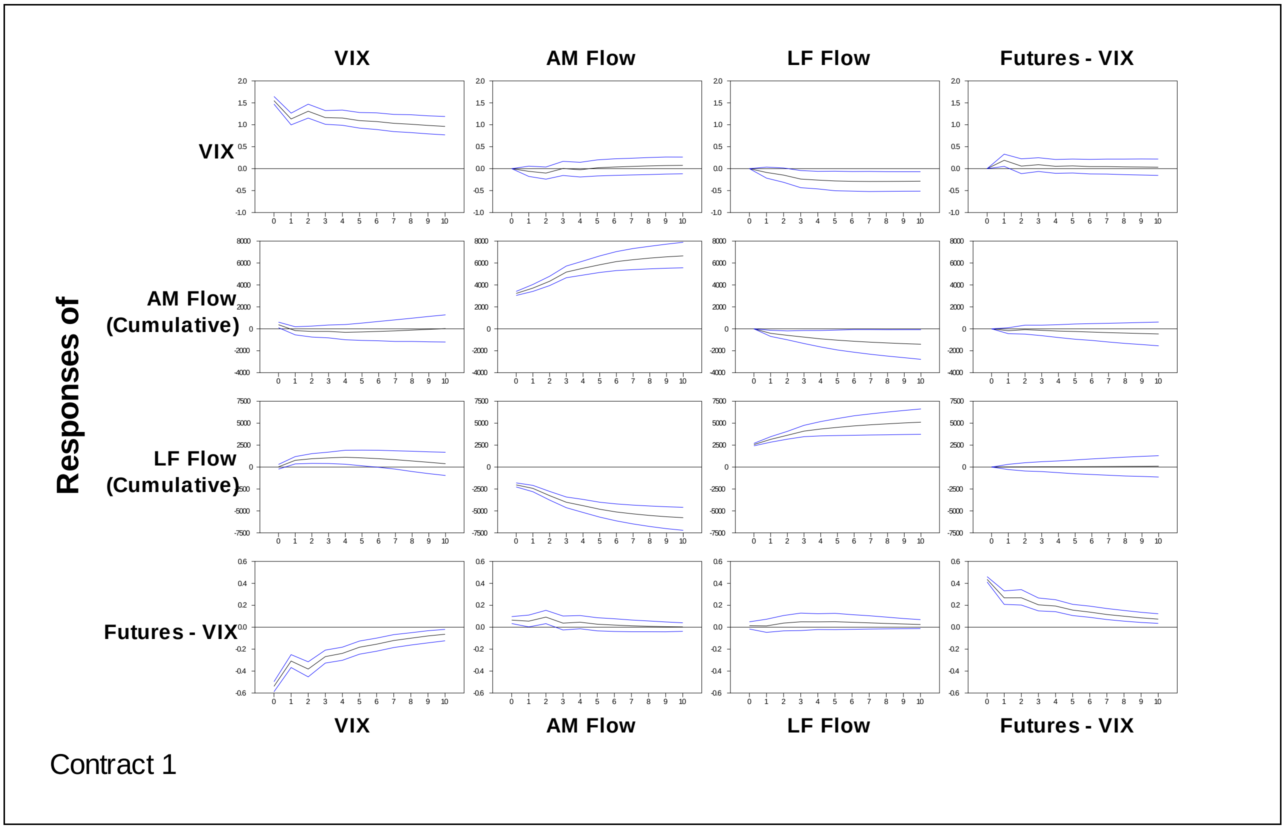

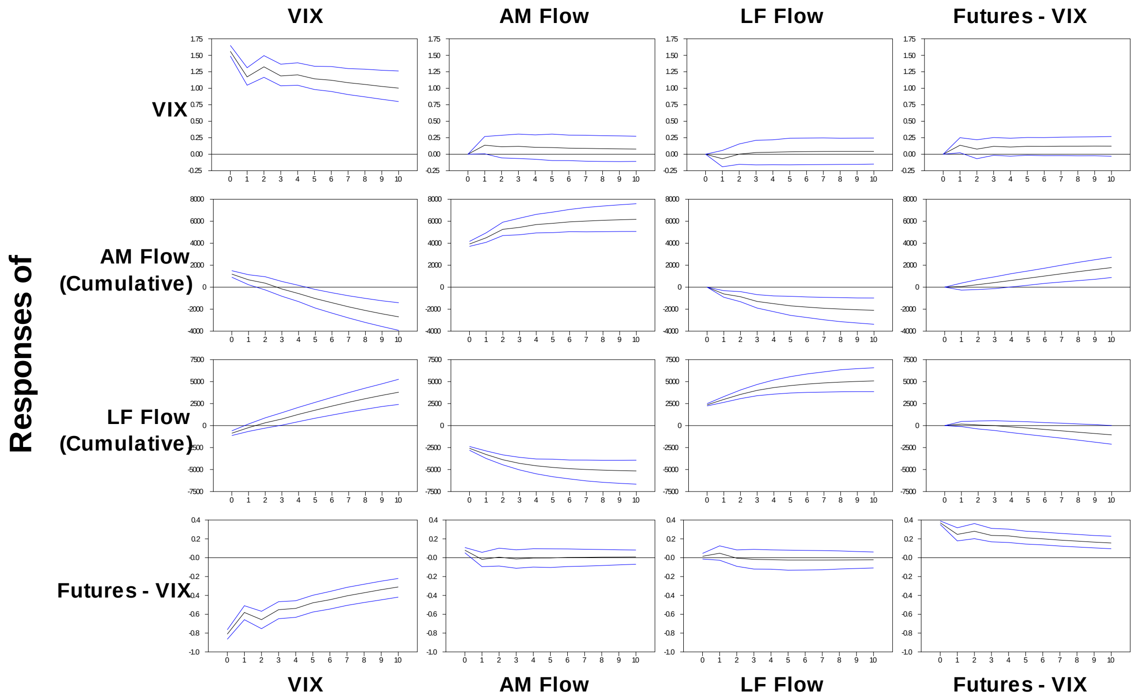

Contract 2

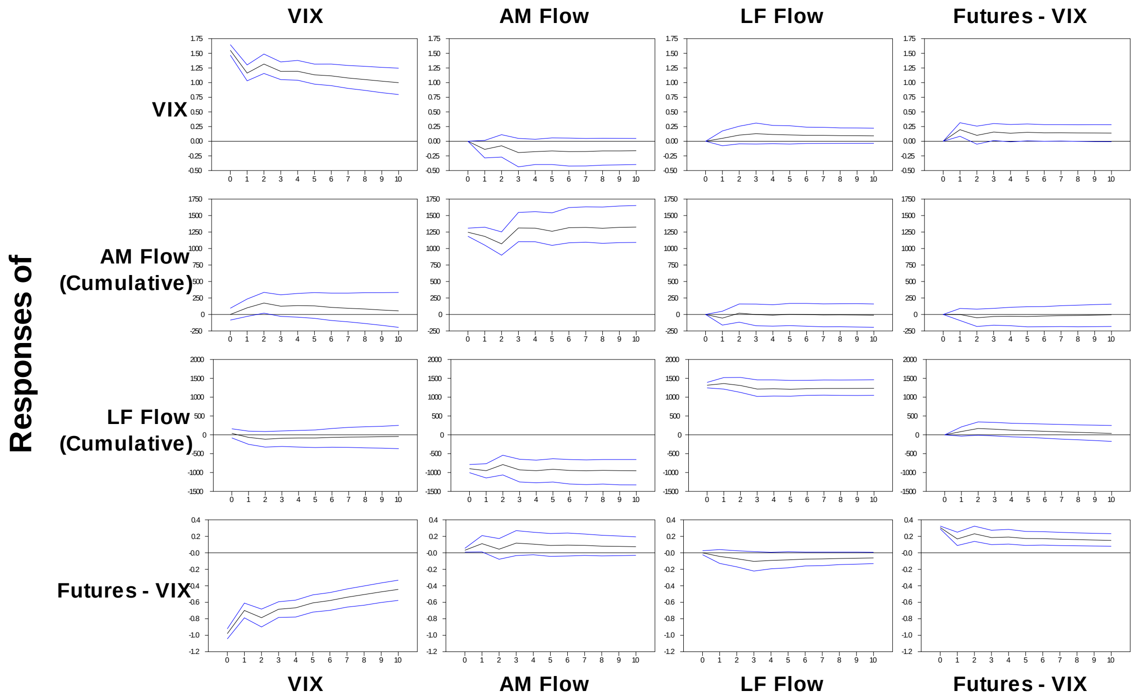

Contract 3

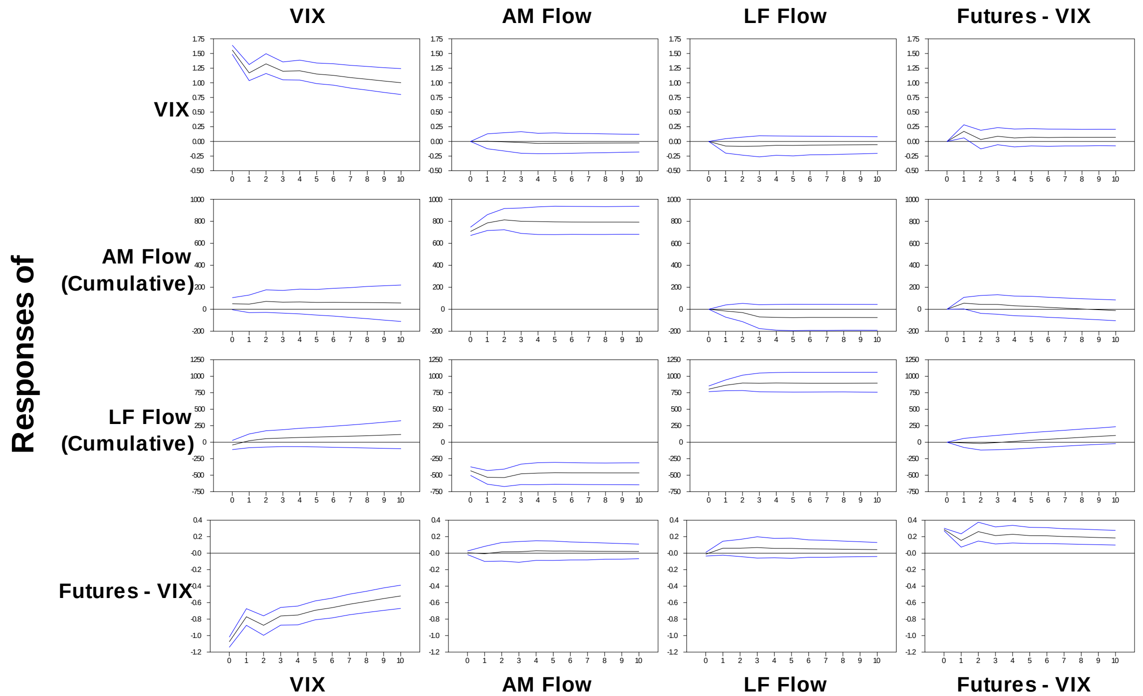

Contract 4

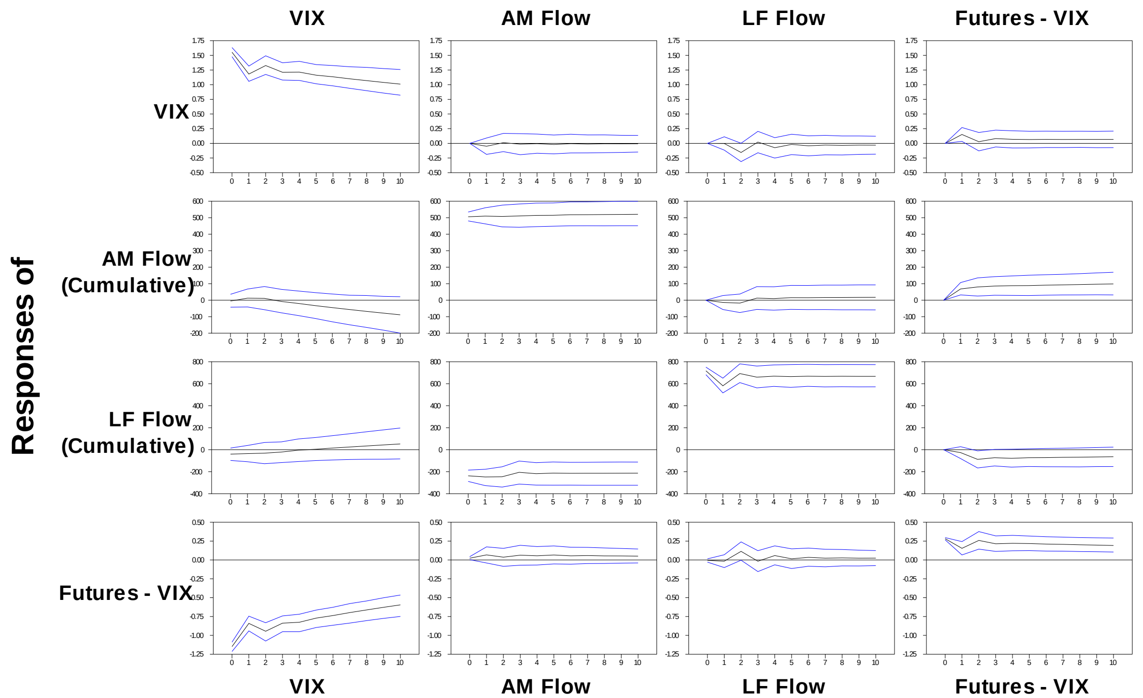

Contract 5

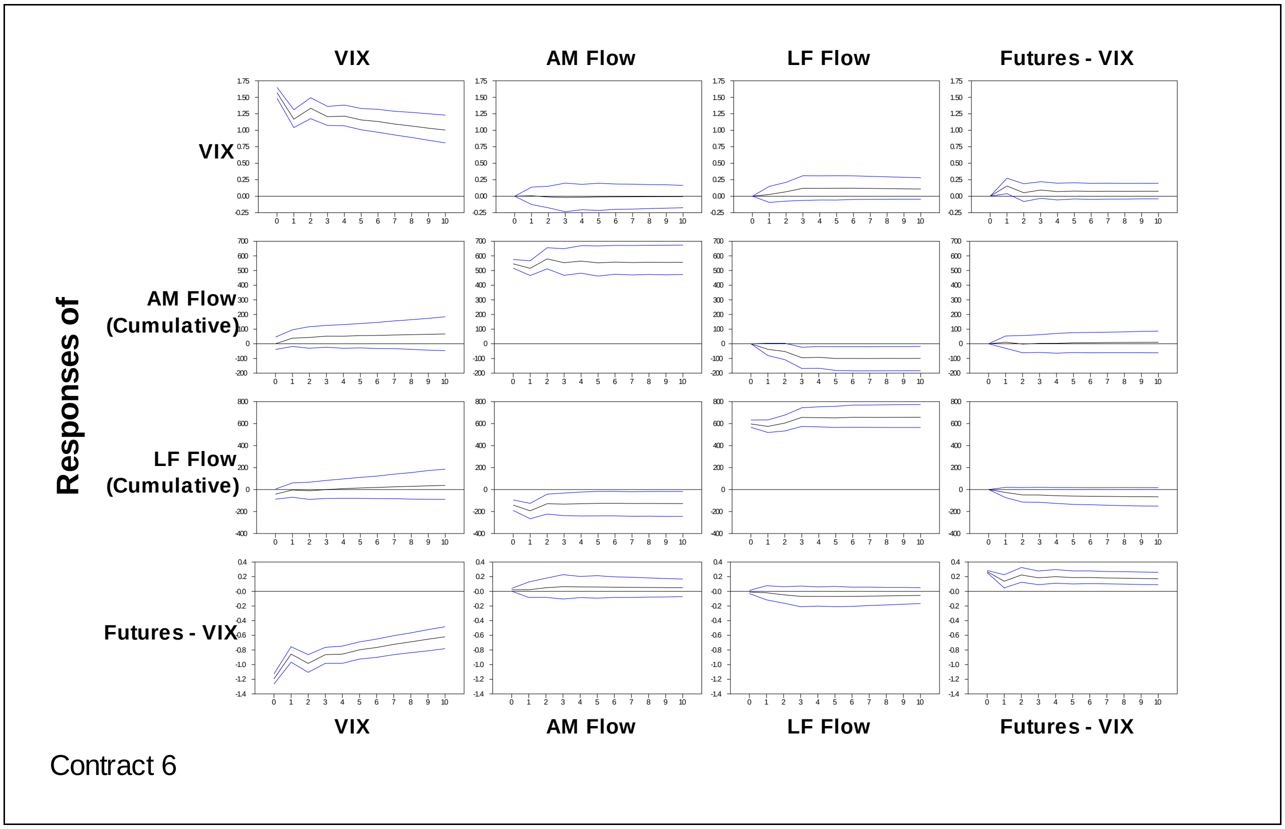

Responses of
VIX
AM Flow
LF Flow
Futures - VIX
VIX
AM Flow (Cumulative)
LF Flow (Cumulative)
Futures - VIX
VIX
AM Flow
LF Flow
Futures - VIX
Contract 6

Table A5.
Variance decomposition for VAR models of VIX futures basis.
The variance decompositions are computed from the VARs described in Figure A1. Each VAR consists of VIX, futures contract flows, and the VIX futures basis. VARs are estimated using daily data spanning the period April 2011 to March 2015. Each entry in the table is the proportion of the forecast variance for a given contract (at the specified forecast horizon) that is explained by the particular factor. The entries with the header "CONTRACT VOLUME" represent the sum of the combined variance decomposition values associated with changes in asset manager contract volume and changes in leveraged fund contract volume.

	VIX			CONTRACT VOLUME (Asset Manager + Leveraged Funds)			Futures basis (Futures price minus VIX)		
	Horizon			Horizon			Horizon		
Contract Expiry	**1**	**5**	**10**	**1**	**5**	**10**	**1**	**5**	**10**
1	58.9	59.9	58.9	1.1	2.5	2.8	39.9	37.6	38.3
2	82.6	83.5	83.0	0.8	0.4	0.4	16.6	16.1	16.6
3	91.3	90.3	89.5	1.0	2.2	2.6	7.7	7.5	7.8
4	94.1	92.4	91.5	0.2	0.5	0.5	5.7	7.1	7.9
5	94.9	93.3	92.6	0.2	0.7	0.7	4.9	6.0	6.7
6	95.8	94.8	94.1	0.1	0.6	0.8	4.1	4.6	5.1