The Law of One Price in Equity Volatility Markets

Peter Van Tassel*

December 2019

Abstract

This paper documents law of one price violations in equity volatility markets. While tightly linked by no-arbitrage restrictions, the prices of VIX futures exhibit significant deviations relative to their option-implied upper bounds. Static arbitrage opportunities occur when the prices of VIX futures violate their upper bounds. The deviations widen during periods of market stress and predict the returns of VIX futures. A simple trading strategy that exploits the no-arbitrage deviations earns significant returns with minimal exposure to traditional risk factors. There is evidence that systematic risk and demand pressure contribute to the variation in the deviations over time.

Keywords: Limits-to-Arbitrage, Volatility, VIX Futures, Variance Swaps, Return Predictability

JEL Classification: G12, G13, C58

*Federal Reserve Bank of New York (email: peter.vantassel@ny.frb.org): I thank Richard Crump, Bjorn Eraker, David Lando, David Lucca, Lasse Pedersen, Ivan Shaliastovich and seminar participants at the Wisconsin-Madison Junior Finance Conference, the Federal Reserve Bank of New York, and the Copenhagen Business School for comments. I thank Charles Smith for excellent research assistance. The views expressed in this paper are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the author.

1 Introduction

The law of one price is a fundamental concept in economics and finance. While studies differ in their conclusions for how risk is related to return, the vast majority of analysis assumes the absence of arbitrage which implies the law of one price. The law states that assets with identical payoffs must have the same price. In competitive markets with limited frictions, traders will exploit any deviations so that financial markets behave as if no arbitrage opportunities exist.

Given its powerful role for theory and practice, a large literature studies the limits of arbitrage and law of one price. Challenging the traditional paradigm, well-documented anomalies illustrate cases in which assets with closely related payoffs trade at significantly different prices for prolonged periods of time. Understanding these anomalies improves our knowledge of financial markets by identifying the frictions and inefficiencies that prevent even sophisticated traders and institutions from being able to take advantage of seemingly profitable situations (Lamont and Thaler (2003), Gromb and Vayanos (2010)).

This paper documents new evidence of systematic law of one price violations in equity volatility markets. These violations matter because equity volatility markets are among the largest and most actively traded derivatives markets in the world. Since the financial crisis, rapid growth in the S&P 500 index options and VIX futures markets has led to the development of separate venues where investors can hedge and speculate on stock market volatility. These markets provide a testing ground for the law of one price because they offer redundant securities upon which arbitrage pricing places tight restrictions (Merton (1973), Ross (1976a)). In practice, however, trading desks at banks and hedge funds tend to focus on specific products, using different models for hedging and valuing different derivatives (Longstaff et al. 2001). This risk management approach makes it difficult to determine when relative valuations and risk exposures are accurate, leaving open the possibility of observing arbitrage violations.

This paper studies relative pricing in equity volatility markets by computing an upper bound for the prices of VIX futures from S&P 500 index options. The upper bound follows from a straightforward application of Jensen's inequality. The payoff to a VIX futures contract is the difference between the futures price and the value of the VIX index at maturity. Because the VIX index is defined as the square-root of a one-month variance swap, the upper bound for a VIX futures contract is a one-month variance swap forward rate expressed in volatility units, not variance units (Carr and Wu 2006). The square-root adjustment in the definition of the VIX is convenient because it expresses the VIX in implied volatility units that are familiar to traders who use the Black-Scholes-Merton model, but the definition also introduces a wedge between the prices of VIX futures and variance swap forwards. The paper applies this observation to define a no-arbitrage deviation measure. The deviation measure is equal to the price of a VIX futures contract minus its upper bound which is the corresponding one-month variance swap forward rate.

The paper finds that VIX futures exhibit significant deviations relative to their upper bounds. Static arbitrage opportunities in which the prices of VIX futures violate their upper bounds occur for 19% of contract-date observations from 2004 to 2018 for the front six VIX futures contracts. There is at least one upper bound violation on 56% of the days in the sample. When the upper bound violations occur, traders can sell VIX futures and pay fixed in a specific quantity of variance swap forwards derived in the paper to lock in an arbitrage profit. Moreover, these violations are not merely a reflection of measurement error in estimating variance swap forward rates. Large violations of more than .50% occur less frequently but are still present for 5% of contract-date observations. To put this magnitude in context, a mispricing of .50% or half of one VIX-point is more than 10 times that typical bid-ask spread of .05% in the VIX futures market.

The paper finds that the upper bound is also relatively tight. The average bias and standard deviation of VIX futures prices around the upper bound are similar in magnitude and both around .80%. Large deviations below the upper bound indicate that VIX futures are cheap relative to index option prices. Combining the upper bound with a convexity adjustment from a term-structure model, the paper estimates a lower bound and finds that VIX futures prices are below the lower bound on 21% of contract-date observations. Taken together, the evidence suggests that there are large and significant deviations between the prices of VIX futures and their corresponding no-arbitrage bounds implied by the index option market.

The paper then explores whether the deviation measure predicts the returns of VIX futures hedged with variance swap forwards. The results indicate that the deviation measure significantly predicts returns across contracts, sample periods, and horizons while remaining robust to competing predictors like the variance risk premium. The predictability also holds using a variety of different datasets to estimate variance swap forward rates, showing that the results are not driven by using a particular variance swap dataset. A simple trading strategy that sells (buys) VIX futures and hedges with variance swap forwards when the deviation measure is large and positive (negative) earns an annualized Sharpe ratio of 3.0 from 2004 to 2018 with minimal stock market exposure. For comparison, the stock market earned a Sharpe ratio of .5 over the same period of time as measured by CRSP value-weighted returns. The large Sharpe ratio from the relative value trading strategy and significant return predictability results are consistent with the law of one price interpretation and notion of trading against an arbitrage spread.

What drives the no-arbitrage violations and return predictability results, the VIX futures or index options market? Since the analysis is relative, it is challenging to provide a sharp answer to this question. One hypothesis is that the larger and more established index options market would serve as a fair-value measure for VIX futures. Alternatively, given the high liquidity in the VIX futures market and the fact that the contracts are traded directly, rather than being estimated from option portfolios across maturities, one might hypothesize that VIX futures would be fairly priced with index option prices driving the deviations.

The paper tests these hypotheses by running return predictability regressions in which VIX futures and variance swap forward returns are hedged with the stock market returns. For example, to the extent that the deviations are driven more by mispricing in the VIX futures market, one would expect larger predictability results for the deviation measure when VIX futures are hedged with stock market returns as opposed to variance swap forward returns. The results indicate that the deviation measure remains most significant at predicting VIX futures returns, particularly during the post-crisis period from 2010 to 2018. In contrast, the deviation measure is less significant at predicting the returns of variance swap forwards. This distinction provides some evidence that VIX futures are driving the deviations and predictability. That said, the strongest results occur when trading VIX futures against variance swap forwards, a relative value finding that is silent on the market driving the results.

Having documented the existence of arbitrage opportunities and the return predictability of the deviation measure, the paper then explores what drives the law of one price violations over time. On one hand, the results are surprising because the index options and VIX futures markets are large and liquid exchange-based markets with competitive traders and limited short-sale constraints. On the other hand, the literature on anomalies shows that frictions like limited arbitrage capital, demand shocks, and financial constraints can lead to return predictability and law of one price deviations (Shleifer (1986), Shleifer et al. (1990), Liu and Longstaff (2003), Adrian and Shin (2013)). In addition, asynchronous observations of VIX futures and index option prices may impact the measurement of the law of one price deviations and transaction costs may prevent traders from implementing the arbitrage in practice. Which, if any, of these hypotheses can explain the observed no-arbitrage deviations?

The paper examines this question by exploring whether a range of variables can explain the variation in the deviation measure. VAR analysis and panel regressions indicate that risk and demand variables are most significant. When stock market returns are negative or when realized or implied volatility increases, the deviation measure tends to decline, with VIX futures cheapening relative to variance swap forwards. There is also evidence that the deviation measure is increasing in various proxies for VIX futures demand pressure. One explanation for these results is that hedgers take profit on long positions in VIX futures when risk increases, leading to demand shocks that are correlated with systematic risk. A trivariate VAR with the deviation measure, VIX index, and dealer net positions in VIX futures from the CFTC's Commitment of Traders (CoT) Report provides some evidence in favor of this hypothesis. To the extent that dealer positions in the CoT report are a veil for retail demand as argued by Dong (2016), the results are consistent with retail hedgers taking profit on hedges after risk increases, but also chasing momentum in the VIX futures market as evidenced by increases in long positions after the deviation measure increases.

The remainder of the paper proceeds as follows. Section 2 provides a literature review. Section 3 describes equity volatility markets. Section 4 defines the no-arbitrage deviation measure. Section 5 presents the return predictability and trading strategy results. Section 6 relates the deviation measure to risk and demand factors. Section 7 concludes. The Appendix includes additional results and robustness checks.

2 Literature Review

The results in this paper contribute to the literature on anomalies and the limits of arbitrage. Examples of related studies focusing on law of one price deviations in other markets include: closed-end funds (Lee et al. 1991), "negative stubs" or situations where the market value of a company is less than its subsidiary (Mitchell et al. 2002), and American Depository Receipts and cross-listed shares (Gagnon and Karolyi 2010) in equities; on-the-run versus off-the-run U.S. Treasuries (Krishnamurthy 2002), swap spreads (Duarte et al. 2007), mortgages (Gabaix et al. 2007), and Treasury-inflation-protected securities (Fleckenstein et al. 2014) in fixed income; the CDS-bond basis (Garleanu and Pedersen (2011), Bai and Collin-Dufresne (2018)) in credit; and covered interest rate parity (Fong et al. (2010), Du et al. (2018)) in foreign exchange. While these examples are only a subset from a large literature, they serve to highlight how law of one price deviations occur across asset classes and are subject to intense scrutiny given the model-free nature in which they challenge the conventional wisdom. In some cases, the deviations are explained by frictions such as transaction costs, financing costs, or convergence risks. In other cases, the deviations are found to significantly predict returns, posing a challenge for standard asset pricing models.

Within the limits-to-arbitrage literature, this paper most closely relates to studies of equity volatility puzzles and anomalies. These studies can be broadly grouped into option pricing and VIX futures papers. Beginning with the option pricing papers, two stylized facts since the crash of 1987 are that: (i) index options are expensive, with delta-hedged straddles earning a significant risk premium known as the variance risk premium, and (ii) out-of-the-money put options are particularly expensive, exhibiting a smirk or higher level of implied volatility relative to at-the-money options (Bates (2000), Coval and Shumway (2001), Jurek and Stafford (2015)). In response to these irregularities that are not explained by the Black-Scholes-Merton model, the literature developed more sophisticated option pricing models that could account for some of the new patterns by allowing for stochastic volatility and jumps in the underlying process (Heston (1993), Duffie et al. (2000), Madan et al. (1998)). But even with these improvements, the no-arbitrage models still struggle to fit and explain some of the empirical properties of option prices (Bates 2003). Another strand of the literature investigates the importance of demand-pressure and equilibrium effects (Grossman and Zhou (1996), Bollen and Whaley (2004), Garleanu et al. (2009)). When dealers absorb demand shocks from end-users or portfolio insurers, demand pressure can impact option prices and the implied volatility surface, resolving some of the well-known option pricing puzzles. An important conclusion from these studies is that demand pressure for one option impacts the prices of other options that have related, unhedgeable features.

Another puzzle from the equity volatility literature is that implied volatility shocks earn a much smaller risk premium than realized volatility shocks (Dew-Becker et al. (2017), Andries et al. (2015)). This result poses a challenge to consumption-based asset pricing models with Epstein-Zin preferences where investors would be willing to pay a premium for both types of volatility shocks. In the data, the Sharpe ratios from receiving fixed in variance swap forwards and from selling VIX futures are much higher for shorter maturities, particularly for one-month variance swaps that are exposed to realized variance risk. Cheng (2018) builds on this result by studying the conditional price of implied volatility risk or the VIX premium. The VIX premium is defined as the difference between the VIX futures price and a statistical forecast of the VIX index at the maturity of the futures contract. This is similar to the variance risk premium which is often defined as the VIX index, or a one-month variance swap rate, minus a forecast of realized variance. The VIX premium and variance risk premium reflect the risk premiums for being exposed to implied volatility and realized volatility shocks respectively.

Cheng (2018) finds that the VIX premium declines when measures of systematic risk increase. For example, the VIX premium declines when stock market returns are negative or when realized or implied volatility increases. This result is surprising because the returns from selling VIX futures are highly correlated with stock market returns. To the extent that the equity and VIX futures markets are integrated by a common stochastic discount factor, one would expect the VIX premium to increase along with the equity risk premium when measures of systematic risk increase. In particular, Martin (2017) argues that the SVIX provides a lower bound for the equity risk premium, but the VIX premium is decreasing in measures of systematic risk like the SVIX, not increasing.

To partially resolve this puzzle, Cheng (2018) documents that dealer hedging demand for VIX futures declines when risk increases. One interpretation is that hedgers take profit on long positions when risk increases, leading to the anomalous response of the VIX premium to risk. Dong (2016) presents related work and argues that dealer positions in VIX futures reflect hedging demand from VIX ETP issuance. Dong (2016)'s reduced form analysis and theoretical model indicate that VIX ETP demand impacts the underlying VIX futures price. Mixon and Onur (2019) provide a related study of demand pressure using a detailed regulatory dataset that includes dealer net positions by contract at a daily frequency, in contrast to the public data used in other studies which includes weekly net position data that is aggregated across contracts. Mixon and Onur (2019) also find evidence of demand pressure with demand increasing the average level and slope of the VIX futures curve. Similar to Garleanu et al. (2009), Mixon and Onur (2019) find that demand for one VIX futures contract spills over to impact the prices of other contracts. However, the estimated demand effects are quantitatively small and typically within the spread between the no-arbitrage upper bound and lower bound for VIX futures.

This paper makes several contributions to the limits-to-arbitrage and equity volatility literatures. First, it provides a detailed, model-free measurement of no-arbitrage deviations across the VIX futures and S&P 500 index options markets. The no-arbitrage deviation measure is the difference between the price of a VIX futures contract and the price of a synthetic variance swap forward rate estimated from index option prices. Existing studies either estimate a risk premium for VIX futures that relies on time-series data and a parametric statistical forecasting model as in (Cheng 2018) or they compute a deviation measure that relies on synthetic variance swap rates and a convexity adjustment from VIX options (Dong (2016), Park (2019)). This paper avoids using VIX options to focus instead on the synthetic variance swap forward rates that are estimated from the larger and more liquid S&P 500 index options market and because VIX options are only available starting in 2006.

Another advantage of the deviation measure in this paper is that it allows for a direct and straightforward measurement of arbitrage opportunities. When the deviation measure is positive, there is a static arbitrage opportunity because the price of a VIX futures contract is above its variance swap forward rate, its corresponding non-parametric upper bound. This paper provides new evidence on the frequency of no-arbitrage upper bound and lower bound violations. The lower bound is estimated from the deviation measure and a convexity adjustment from a term-structure model. In addition, the paper shows that the deviation measure significantly predicts the returns of hedged VIX futures. This adds another predictor to the VIX futures literature to complement the VIX premium of Cheng (2018) and slope factor of Johnson (2017) which also predict VIX futures returns. The predictability in this paper is strongest when VIX futures are hedged with variance swap forwards, a relative value result that is new to the literature. Compared to existing studies, the size of the return predictability and Sharpe ratio in the trading strategy are large when VIX futures are hedged with variance swap forwards.

Finally, this paper provides new evidence on the importance of the risk and demand channels in driving the no-arbitrage deviations over time. In a VAR that accounts for the interaction of risk and demand factors, the response to a risk shock is 3-4 times larger in magnitude than the response to a demand shock over short horizons. The results indicate that demand by itself is not sufficient to explain the negative response of the no-arbitrage deviation measure to risk. This observation coupled with the no-arbitrage violations documented in the paper brings back into focus the puzzle as to why the prices of VIX futures decline relative to variance swap forward rates when risk increases.

3 Equity Volatility Markets

3.1 History and Products

Index options complete markets and expand the set of contingent claims that investors can trade by allowing for the construction of Arrow-Debreu securities on the state of the stock market over different horizons (Ross (1976b), Breeden and Litzenberger (1978)). In practice, however, the option portfolios that span different payoffs can be complicated to construct and involve trading a significant number of options, resulting in large implementation costs. While narrower in focus, the variance swap and VIX futures markets allow investors to trade the level of realized and implied volatility directly, without requiring a need for static or dynamic option trading strategies.

The S&P 500 index options started trading on the Chicago Board Options Exchange (CBOE) in 1983 but it was not until the late 1990s that the variance swap market started to gain traction, perhaps encouraged by the high levels of index-option implied volatility that followed the Asian financial and LTCM crises (Carr and Lee 2009). The motivation for trading variance swaps stemmed from investor demand to obtain exposure to the variance risk premium and pure volatility risk. Although these payoffs can be approximated by a delta-hedged straddle, they cannot be spanned with a limited number of options. Consider a delta-hedged straddle. When the underlying moves away from the strike price, the exposure to realized and implied volatility changes. In Black-Scholes parlance, the gamma and vega of a delta-hedged straddle change with the moneyness of the option. In contrast, a variance swap exchanges a fixed payment at maturity for a floating payment equal to the realized volatility of the underlying asset over the life of the swap, thus providing direct exposure to realized variance risk over different horizons.

The dawn of the variance swap market coincided with a number of theoretical developments and influential papers that showed how to replicate the payoff of a variance swap with a static portfolio of options and a dynamic trading strategy in the underlying (Demeterfi et al. (1999), Carr and Madan (1999), Bakshi and Madan (2000), and Britten-Jones and Neuberger (2000)). This no-arbitrage replication argument is model-free in the sense that it only requires the ability to trade a continuum of European options with different strike prices, no jumps in the underlying asset, constant interest rates, and continuous trading of the underlying. Seeing the generality of this approach and the growth in the variance swap market, in 2003, the CBOE revised its definition of the VIX index to follow the no-arbitrage formula for pricing variance swaps with index option prices (Carr and Wu (2006), CBOE (2019)).

The CBOE then introduced trading in VIX futures and options in 2004 and 2006 respectively. The payoff to a VIX futures contract is the difference between the futures price and a special opening quotation of the VIX. From its definition, the VIX index upon which VIX futures are based is equal to the square root of a one-month synthetic variance swap rate. The swap rate is "synthetic" because it is computed using the price of a portfolio of S&P 500 index options. This relationship binds together VIX futures prices and index option prices up to a convexity adjustment. The convexity adjustment is the definition of the VIX as the square-root of the price of the index option portfolio, not just the price of the portfolio itself. The square root convexity adjustment expresses the VIX index in volatility units that are familiar to option traders. The VIX is defined in the same unit as the Black-Scholes-Merton implied volatility parameter.

This paper relies on a detailed dataset of synthetic variance swap rates for the S&P 500 index that are derived in Van Tassel (2019) to study the no-arbitrage relationship between the index options and VIX futures markets. The synthetic variance swap rates are computed from the theoretical price for a variance swap following Carr and Wu (2009). The implementation is similar to the computation of the VIX index but excludes the strike truncation and discretization that the CBOE uses. The paper then converts the synthetic variance swap rates into variance swap forward rates for comparison to VIX futures prices. For example, a two-month variance swap can be decomposed into a one-month variance swap plus a one-month forward one-month variance swap. The ability to perform this decomposition is essential for the analysis in the paper. Variance swap forwards provide a non-parametric upper bound for the prices of VIX futures contracts (Carr and Wu 2006). This result serves as the basis for the main no-arbitrage deviation measure introduced in the paper.

In recent work, Martin (2017) has extended the generality of variance swap pricing. By using simple returns instead of log returns to compute the realized variance payoff for the floating leg of a variance swap, Martin (2017) shows that the assumption of a continuous underlying can be relaxed. Why does this paper use synthetic variance swap rates computed from the traditional formula rather than simple variance swaps instead? While simple variance swaps have the advantage that they require no third-order approximation for jumps, they are not easily decomposed into forward rates and they do not serve as the basis for the VIX index. Thus, they do not have as close of a relationship to the pricing of VIX futures as do synthetic variance swap rates and therefore are less well-suited for studying no-arbitrage relationships across the index options and VIX futures markets.

3.2 Market Size and Participants

To understand why no-arbitrage violations may occur across equity volatility markets, it is helpful to provide a background on the institutional details of these markets. The relative size of the different markets, types of participants, investor positioning, and linkages across markets are all potential determinants of no-arbitrage violations.

The S&P 500 index options and VIX futures markets are large and liquid exchange-based markets. Figure 1 provides a brief summary of these markets including the growth in their open interest over time and a breakdown of investor positioning in VIX futures. The top left plot illustrates the size of the index options market relative to the VIX futures market over time. Both markets have exhibited significant growth over the past decade. S&P 500 index options had an average open interest of $3.4 billion in Black-Scholes vega in 2018. VIX futures had an average open interest of 462 thousand contracts in 2018, equal to $462 million of "vega" or gains and losses for a one-point change in VIX futures prices given the contract multiplier of $1000. While the VIX futures market experienced a nearly 10-fold increase in open interest over the past decade, the index options market is still 7-times larger in 2018 as measured by vega. To the extent that larger market size corresponds to greater liquidity and more informative prices, the index option market provides a valuable signal about the fair value of VIX futures prices. The no-arbitrage deviation measure introduced later in the paper is motivated by this idea, using index option prices to bound VIX futures prices.

Focusing on the VIX futures market, the top right plot shows that post-crisis growth coincides with the rise of volatility-exchange-traded products (VIX ETPs). VIX ETPs issued by banks and broker-dealers allow retail investors to gain exposure to implied volatility without needing to trade in the index options or VIX futures markets directly. For example, consider the VXX, one of the first ETPs introduced to the market in early 2009. The VXX ETN tracks the S&P 500 VIX Short-Term Futures Index, rolling from the front-month contract to the second-month contract to provide investors with a weighted average futures maturity of one-month. Instead of needing to perform the roll themselves, investors can simply purchase and hold the exchange-traded note. Over time, similar products were introduced that track different VIX futures indices with long or short exposure, potentially building in embedded leverage. The top right plot of figure 1 illustrates that the growth of VIX futures has coincided with the growth of VIX ETPs.

One explanation for the simultaneous growth in the VIX futures and ETP markets is that dealers issue VIX ETPs to satisfy retail demand and then hedge in the underlying VIX futures market. Dong (2016) suggests this hypothesis and shows that short-term ETP net vega is highly correlated with dealer positions from the CFTC's Commitment of Traders (CoT) Report.[1] The bottom left plot replicates and extends this result. Dealer positions closely match the magnitude of VIX short-term ETP net vega and the two demand variables track each other with a correlation of around 60% from 2010 to 2018. This result suggests that retail demand may be an important determinant for the growth in the VIX futures market. If retail demand to buy volatility causes dealers to hedge in the VIX futures market, this could affect VIX futures prices and drive them above the no-arbitrage bounds implied by the index option market.

If dealer positions in the VIX futures market are just a veil for retail demand, who absorbs the retail demand shocks? The bottom right plot attempts to address this question by plotting net positions for different trader types including dealers, leveraged funds, asset managers, and other reportable traders from the CFTC's CoT Report. Focusing on the post-crisis period, the largest net position by magnitude belong to dealers and leveraged funds. Dealer and leveraged funds net positions exhibit a correlation of -91% from 2010 to 2018. One interpretation of this result is that leveraged funds absorb retail demand shocks that are passed through by dealers from the volatility ETP market.

With this context in mind, the paper next describes the no-arbitrage bounds on VIX futures prices and introduces a deviation measure which tracks arbitrage opportunities over

[1] The VIX ETP demand variable is equal to the leverage-weighted market capitalization of short-term ETPs net of short interest. Following Dong (2016), the formula is $D_{ETP} = \sum_{i \in ST}(Shrout_i - ShortInt_i) \cdot P_i \cdot M_i$ where $Shrout_i$ is shares outstanding, $ShortInt_i$ is short interest, P_i is price, and M_i is the direction and leverage multiplier. I compute D_{ETP} using Bloomberg data for the VXX, VIIX, VIXY, UVXY, TVIX, XIV, SVXY, IVOP, XXV, and VXXB ETPs. The multipliers are equal to $M = [1\ 1\ 1\ 2\ 2\ \text{-}1\ \text{-}1\ \text{-}1\ \text{-}1\ 1]$. I adjust the multipliers for UVXY and SVXY to 1.5 and -.5 after 2/28/18. To express this demand variable as VIX Short-Term ETP Net Vega in \$ million of vega in the plot, I divide D_{ETP} by the VIX index. Similarly, the total vega number is $\sum_{i \in ST} Shrout_i \cdot P_i \cdot |M_i| / VIX$.

time and across contracts. The paper finds that there are arbitrage opportunities across the VIX futures and index option markets. The deviation measure widens when risk increases and predicts VIX futures returns. The paper then returns to the demand variables described above to investigate the relationship between the no-arbitrage deviation measure, risk, and demand pressure.

4 No-Arbitrage Deviation Measure

4.1 VIX Futures Bounds

The upper bound for the price of a VIX futures contract is the maturity-matched one-month variance swap forward rate expressed in volatility units. The derivation of the upper bound follows from the definition of the VIX index and an application of Jensen's inequality (Carr and Wu 2006). The VIX is defined as the square-root of a one-month variance swap rate. By taking the square-root, the definition expresses the VIX in volatility units that are familiar to traders who use the Black-Scholes-Merton model. At the same time, the square-root introduces a wedge between the price of VIX futures and their corresponding variance swap forward rates.

To derive the upper bound, let the price of a VIX futures contract be the risk-neutral $\mathbb{Q}$ expected value of the VIX index at maturity,

$$Fut_{t,T} = E_t^{\mathbb{Q}}[VIX_T]. \tag{1}$$

As noted above, the VIX is defined as the square-root of a one-month variance swap rate $VIX_t \equiv \sqrt{VS_{t,t+1}}$.[2] Let variance swaps be modeled as the risk-neutral expected value of realized variance from the trade date until the maturity of the swap,

$$VS_{t,T} = E_t^{\mathbb{Q}}\left[RV_{t,T}\right]. \tag{2}$$

Similarly, let variance swap forward rates be the risk-neutral expected value of realized variance between two forward starting dates,

$$Fwd_{t,T_1,T_2} = E_t^{\mathbb{Q}}\left[RV_{T_1,T_2}\right]. \tag{3}$$

[2]In practice the VIX index is expressed in annualized volatility units. The paper omits the annualization in this derivation for notional simplicity. Empirically the paper compares VIX futures prices to variance swap forward rates in annualized units, using calendar time to perform the annualization as in the CBOE construction of the VIX (CBOE 2019).

It follows that the prices of VIX futures are bounded above by variance swap forward rates,

$$
\begin{aligned}
Fut_{t,T} &= E_t^{\mathbb{Q}}[VIX_T] \\
&= E_t^{\mathbb{Q}}\left[\sqrt{VS_{T,T+1}}\right] \\
&\leq \sqrt{E_t^{\mathbb{Q}}[VS_{T,T+1}]} \\
&= \sqrt{E_t^{\mathbb{Q}}\left[E_T^{\mathbb{Q}}[RV_{T,T+1}]\right]} \\
&= \sqrt{Fwd_{t,T,T+1}}.
\end{aligned} \tag{4}
$$

The one-month variance swap forward rate from time T to time $T+1$ expressed in volatility units is $\sqrt{Fwd_{t,T,T+1}}$.

If VIX futures violate the upper bound so that $Fut_{t,T} > \sqrt{Fwd_{t,T,T+1}}$, there is a static arbitrage opportunity . To see this, note that $VS_{T,T+1} = f(VIX_T) = VIX_T^2$ is a convex function. This motivates the trade:

Arbitrage Trade for Upper Bound Violation

Time	t	T
Sell $f'(Fut_{t,T})$ VIX futures contracts	0	$f'(Fut_{t,T})(Fut_{t,T} - VIX_T)$
Pay fixed in VS forward	0	$VS_{T,T+1} - Fwd_{t,T,T+1}$
Total	0	$VS_{T,T+1} - f(Fut_{t,T}) - f'(Fut_{t,T})(VIX_T - Fut_{t,T})$ $+f(Fut_{t,T}) - Fwd_{t,T,T+1}$

The first term in the payoff at time T is non-negative because f is convex. The second term is positive because $f(Fut_{t,T}) = Fut_{t,T}^2 > Fwd_{t,T,T+1}$ by assumption. The trade is an arbitrage because it requires no investment at time t and guarantees a positive payoff at time T.

The prices of VIX futures contracts are also bounded below by volatility swap rates. The derivation of the lower bound proceeds along similar lines. In particular, the price of a VIX futures contract satisfies,

$$
\begin{aligned}
Fut_{t,T} &= E_t^{\mathbb{Q}}\left[\sqrt{E_T^{\mathbb{Q}}[RV_{T,T+1}]}\right] \\
&\geq E_t^{\mathbb{Q}}\left[E_T^{\mathbb{Q}}\left[\sqrt{RV_{T,T+1}}\right]\right] \\
&= E_t^{\mathbb{Q}}\left[\sqrt{RV_{T,T+1}}\right] \\
&\equiv Fvol_{t,T,T+1},
\end{aligned} \tag{5}
$$

where the last line defines a volatility swap forward rate. In the discussion below the paper will focus on the upper bound because the upper bound can be estimated non-parametrically directly from observed variance swap rates. To compute the lower bound the paper will rely on a dynamic term-structure model to estimate the difference between the upper and lower bounds following the approach in Van Tassel (2019).

4.2 No-Arbitrage Deviation Measure

The ability to bound VIX futures prices with variance swap forward rates motivates defining the no-arbitrage deviation measure,

$$\text{Deviation}_{t,n} \equiv \text{Fut}_{t,n} - \text{Fwd}_{t,n}. \tag{6}$$

This deviation measure is the difference between the price of the n-th futures contract $\text{Fut}_{t,n}$ expiring at time τ_n and the corresponding one-month variance swap forward rate $\text{Fwd}_{t,n}$ expressed in volatility units. The deviation measure is defined in levels rather than as a percentage because the absolute size of the no-arbitrage violations matter. Abstracting from transaction costs, arbitrageurs should trade against a large upper bound violation, regardless of whether the VIX is at 20 or 50. Note also the slight abuse of notation compared to the previous section. In the deviation measure, $\text{Fwd}_{t,n}$ denotes the variance swap forward rate in volatility units to be directly comparable to VIX futures prices.

The deviation measure is attractive because of its straightforward interpretation and for its ability to be estimated non-parametrically with minimal assumptions. When the deviation measure is positive, VIX futures prices are above their upper bound which creates a static arbitrage opportunity or law of one price violation. When the deviation measure is negative, VIX futures are cheap relative to variance swap forwards to the extent that the bound is relatively tight. Large negative deviations may also violate the lower bound, but this is harder to measure because the prices of volatility swaps are harder to observe. Compared to variance swaps, volatility swaps are traded infrequently and are difficult to replicate from index option prices. Therefore, the paper uses the upper bound from variance swap forward rates to define the deviation measure.

Computing the deviation measure requires the prices of VIX futures and of variance swap forwards. The VIX futures price is directly observable from settlement prices or intraday trade-and-quote data. The variance swap forward price can be computed directly from a variance swap curve. The paper computes the variance swap forward curve assuming flat forward rates between observed variance swap maturities.[3] Figures 2 and 3 provide two

[3]For example, let VS_{t,T_1} and VS_{t,T_2} be variance swap rates at maturities T_1 and T_2. The forward rates between T_1 and T_2 are assumed to be constant and equal to Fwd_{t,T_1,T_2} in annualized variance units where Fwd_{t,T_1,T_2} satisfies $VS_{t,T_1}(T_1 - t) + Fwd_{t,T_1,T_2}(T_2 - T_1) = VS_{t,T_2}(T_2 - t)$. Let $F(\tau)$ denote the corresponding forward curve, i.e. $F(t, \tau) = Fwd_{t,T_1,T_2}\ \forall \tau \in (T_1 - t, T_2 - t)$ and a similar definition for other maturities τ. The upper bound for the n-month VIX futures contract expiring at time τ_n is then equal to

examples of this computation on two dates with large no-arbitrage deviations. Figure 2 plots the variance swap curve, instantaneous forward curve, and one-month forward rates for the VIX futures maturity dates. Figure 3 compares the VIX futures prices to their corresponding variance swap forward rates. Table 1 provides the corresponding data, including the synthetic variance swap curve from which forward rates are estimated and the VIX futures prices. The example on February 27, 2012 illustrates a case when the prices of VIX futures are above the upper bound resulting in a static arbitrage opportunity. The size of the violation is large as the deviation measure is more than 1% for several of the contracts. The example on August 10, 2011 illustrates the opposite case when the prices of VIX futures are well below the upper bound, often by as much as 3% to 5%. Using an estimate from a term-structure model for the difference between the upper and lower bounds for the different contracts on this date, it appears that VIX futures prices are below the lower bound. This indicates an arbitrage opportunity if arbitrageurs are able to transact at volatility swap forward rates near the lower bound estimates.

The paper computes the deviation measure with synthetic variance swap rates from Van Tassel (2019) and synchronized VIX futures prices. The synthetic variance swap rates are constructed from index option prices from OptionMetrics data for maturities of $\tau = \{1, 2, 3, 6, 9, 12, 15, 18, 24\}$ months. To be synchronous with the option quotes, the VIX futures prices are settlement prices from 4:15pm ET prior to March 3, 2008 and then mid-prices at 4pm ET from Thomson Reuters Tick History when available. This change in timing matches the OptionMetrics quotes which switched to 4pm in March 2008. As a robustness check, the Appendix computes the deviation measure in a number of alternative ways and finds results that are broadly the same. The baseline deviation measure is 97% to 100% correlated with alternative measures that make a convexity adjustment to reduce the bias in the upper bound, either using a term-structure model or regression. In addition, the baseline measure is around 90% correlated with deviation measures computed from Bloomberg and CBOE volatility index data and 70% to 80% correlated with deviation measures computed from OTC variance swap quotes. The baseline measure has the advantage that it is available over the full sample period from March 2004 to December 2018, whereas the alternative datasets are only available for part of the sample period.

$\mathrm{Fwd}_{t,n} = \sqrt{\int_{\tau_n}^{\tau_{n+1}} F(t,s)ds}$. Similar qualitative results are obtained if the flat forward rates are smoothed with a moving average filter like a lowess smoother (unreported in the paper).

4.3 Deviation Measure Time-Series

The deviation measure varies significantly over time. Figure 4 plots the average deviation and average absolute deviation for the front six contracts from March 2004 to December 2018. During the first part of the sample period the average deviation exhibits positive values on a significant fraction of dates indicating the presence of arbitrage opportunities. Since 2012 the average deviation declines from around -.50% to -1.00% and is positive less often. The plot also highlights how the deviation measure responds to changes in risk. Around negative events that coincide with stock market declines like the financial crisis, equity flash crash, and S&P downgrade of U.S. debt, the deviation measure decreases while the absolute deviation increases in magnitude.

Table 2 presents summary statistics of the deviation measure for the front six contracts from 2007 to 2018 when a balanced panel is available. The last column averages the statistics across contracts. The results indicate that the deviation measure is negative on average with a bias that increases from around -.25% for the front contract to around -1% for the longer-dated contracts. The negative bias is statistically significant as indicated by the t-statistics and consistent with the definition of the deviation measure which equals the VIX futures price minus an upper bound. In terms of variability, the deviation measure has a standard deviation of around .80% across contracts and exhibits negative skewness and excess kurtosis, particularly for the first and second contract. The deviation measure is also positively autocorrelated across one-day, one-week, and one-month horizons. The results are consistent with the time-series plot, indicating that there are large and persistent law of one price deviations across the VIX futures and index options markets.

The summary statistics also highlight differences across contracts. Panel B reports the correlation matrix of the deviation measure. While the pairwise correlation is positive across contracts, it tends to decline as the contract maturities get further apart and is lowest for the front contract. The average pairwise correlation for the other contracts is around 50%. Panel C reports how the average deviation varies across contracts over time. The results indicate that the decline in the average deviation from the time-series plot is driven by the longer-dated contracts rather than the front contract. Combined with the previous results, the summary statistics indicate that the deviation measure exhibits substantial variability both over time and within the cross-section.

4.4 The Frequency of Law of One Price Violations

Table 3 reports the frequency of law of one price violations. Panel A.I reports the fraction of days when the deviation measure exceeds a positive threshold $\text{Deviation}_{t,n} \geq \tau$ where $\tau \in$

$\{0, .25, .50\}$ over the full 2004 to 2018 sample period. Positive values of the deviation measure indicate the presence of arbitrage opportunities in which the prices of VIX futures are above their upper bound from variance swap forward rates. The table indicates that almost 20% of contract-date observations correspond to upper bound violations. The violations are most frequent for the front and second contracts which exhibit positive values for $\text{Deviation}_{t,n}$ on 35% and 24% of days respectively. The results for the lower bound violations in Panel B.I is similar. Nearly 20% of contract-date observations correspond to a lower-bound violation with these violations occurring more frequently for longer-dated contracts. Taken together, the results indicate that VIX futures exhibit frequent law of one price violations relative to their no-arbitrage bounds.

Are these results driven by the early years in the sample, small upper bound violations, or longer-dated less liquid contracts? The answer is no. Panels A.I and B.I show that upper bound and lower bound violations occur for all contracts. Panels A.II and B.II report the number of upper bound and lower bound violations from 2010 to 2018. For the front contract, the number of upper bound violations remains around 35% in the post-crisis period. The average number of upper bound violations declines from 19% to 12% of contract-date observations, but this remains a substantial fraction of violations. Moreover, the number of lower bound violations increases from 22% to 30%, with large increases for the longer-dated contracts. This rules out the hypothesis that the arbitrage violations are solely driven by the early years in the sample; these violations do not go away after liquidity improved or after traders learned of the no-arbitrage relationships. To illustrate this point graphically, Figure 5 reports the deviation measure for the front contract and the VIX futures price and variance swap forward rate. The top plot shows that the deviation measure indicates consistent law of one price violations across the entire sample period. The bottom plot shows that the VIX futures price and variance swap forward rate track each other closely throughout the entire sample period. Thus, it does not appear that the violations are driven by the early years in the sample. Instead, they are pervasive.

The no-arbitrage violations are also large in magnitude. The table indicates that around 10% (5%) of contract-date observations have upper bound violations greater than .25% (.50%) over the full sample period. For the lower bound there are around 14% (8%) of contract-date observations with violations lower than -.25% (-.50%). The change in the 2010-2018 period for the large violations is similar to the change for the baseline violations around $\tau = 0$. These results also address concern about measurement noise. In particular, estimation of variance swap forward rates introduces some noise in the deviation measure which is a potential concern for measuring the frequency of arbitrage opportunities. The fact that a significant fraction of law of one price violations are of large magnitude assuages this concern. In addition, the time-series plots show a clear relationship between the estimated deviation measure and risk events. If measurement noise dominated the results, these plots would appear much noisier and the deviation measure would not predict the returns of VIX futures. The next section investigates return predictability and finds that the deviation measure is highly significant at predicting returns, consistent with the idea that it identifies no-arbitrage violations between the prices of VIX futures and variance swap rates.

5 VIX Futures Return Predictability

5.1 VIX Futures and Variance Swap Forward Returns

To study return predictability, the paper defines the excess return from selling the n-th VIX futures contract over horizon h as,

$$R_{t+h,n}^{VIX} = Fut_{t,n} - Fut_{t+h,n}. \tag{7}$$

Similarly, the paper defines the excess return from receiving fixed in a one-month variance swap forward with a maturity matched to the n-th futures contract as,

$$R_{t+h,n}^{VSF} = Fwd_{t,n} - Fwd_{t+h,n}, \tag{8}$$

where the forward price $F_{t,n}$ is expressed in variance units. The Appendix shows that similar qualitative results hold for log- and percentage-returns.[4] While percentage-returns may be preferred in other settings, such as for stocks where prices are not stationary and the percentage-return is the return on investment from buying one share, this is not the case for VIX futures. The return defined above $R_{t+h,n}^{VIX}$ is the return an investor would receive from selling one contract, abstracting from the multiplier and margin for simplicity. In contrast, generating percentage-returns for VIX futures requires a more complicated trading strategy that adjusts the position size each day to account for the level of the futures price, a difference that would entail additional transaction costs in practice. The trading strategy discussed below will define returns that are similar to $R_{t+h,n}^{VIX}$ taking the multiplier, margin, and transaction costs into account.

Table 4 reports summary statistics for VIX futures and variance swap forwards prices and returns. The table includes the front six contracts from 2007 to 2018 when a balanced panel is available and reports the average statistic across contracts in the last column. For

[4]The log- and percentage-returns are defined as $Fut_{t+h,n}/Fut_{t,n} - 1$ and $\log(Fut_{t+h,n}/Fut_{t,n})$.

prices, the table indicates that the unconditional term-structure is upward sloping. The standard deviation is decreasing in the contract number reflecting the mean-reversion of implied volatility. The prices are positively skewed and exhibit excess kurtosis. For returns, the average return tends to be positive but is not statistically significant. The exception is the fifth and sixth contract for variance swap forwards where the average return is actually negative. These results are consistent with Dew-Becker et al. (2017) and Andries et al. (2015) who find a large risk premium from being exposed to realized volatility, but a smaller risk premium for being exposed to implied volatility. The returns also exhibit negative skewness and excess kurtosis which is common for volatility selling strategies.

5.2 Return Predictability

Table 5 reports the main return predictability regressions. The regression is a two-step procedure. First, VIX futures are regressed onto variance swap forward returns to obtain a hedge ratio for each contract β_n. The variance swap forward returns are highly significant for VIX futures returns, exhibiting an average explanatory power of around 70% to 80% as reported in the Appendix. The hedged returns are then regressed onto the deviation measure to see whether the no-arbitrage deviation predicts VIX futures returns that are hedged with variance swap forward returns. The idea is that, when the deviation measure is high, VIX futures are expensive relative to variance swap forwards, so the returns from selling VIX futures and paying fixed in variance swap forward rates should be high. Similarly, when deviation is low, VIX futures are inexpensive relative to variance swap forward rates so the returns from selling VIX futures should also be low. The table confirms this hypothesis with the highly significant and positive point estimates on the deviation measure across specifications.

Panel A reports results for the full sample from 2004 to 2018 at a weekly horizon $h = 5$. For ease of interpretation, the dependent and independent variables are standardized. For example, in Panel A.I, a one-standard deviation increase in the deviation measure for the second contract predicts a .23 standard deviation higher return with an $R^2 = 5.1\%$. A one-standard deviation increase in the deviation measure for the sixth contract predicts a .38 standard deviation higher return with an $R^2 = 14.5\%$. The average R^2 and t-statistic across contracts are 8.3% and 5.7 respectively. Panel A.II shows that these results are robust to including the VIX, realized variance over the past 21 days (RV), the CRSP value-weighted stock market return over the past week ($RMRF$), and volume for the n-th contract normalized by open interest (VLM). If anything, including these additional variables increases the strength of the deviation measure as a predictor and also increases the in-sample explanatory power.

Are these results specific to the sample period, investment horizon, or deviation measure specification? The answer is no. The deviation measure robustly predicts the returns of VIX futures hedged with variance swap forward rates across the specifications explored in the paper and Appendix. For example, Panel B runs the regressions for a post-crisis sample period from 2010 to 2018 and finds similar results. If anything, the predictability in the post-crisis period increases for the second contract but remains stable for the other contracts. The average R^2 and t-statistic for the deviation measure in Panel B.II are 12% and 7.52. The Appendix reports additional results showing that the predictability holds over daily and monthly horizons, for lagged and bias-adjusted versions of the deviation measure, and when the deviation measure is computed with synthetic variance swap rates from alternative datasets like Bloomberg data or the CBOE volatility indices.

The evidence suggests that the deviation measure significantly predicts the returns of VIX futures hedged with variance swap forwards. As noted before, these results provide reassurance that the deviation measure is picking up a no-arbitrage deviation and does not merely reflect measurement noise from estimating variance swap forward rates. But if the deviation measure is picking up mispricing, is the VIX futures or index option market driving the predictability? Since the analysis is relative, it is challenging to provide a sharp answer to this question. One hypothesis is that the larger and more established index options market would serve as a fair-value measure for VIX futures. Alternatively, given the high liquidity of VIX futures in the latter part of the sample and the fact that VIX futures contracts are traded directly, rather than being derived from option portfolios across maturities, one might hypothesize that VIX futures would be fairly priced with variance swap forward rates driving the predictability. One way to test these hypotheses is through the predictability of the deviation measure when the returns of VIX futures and variance swap forwards are hedged with stock market returns. If the deviation measure is picking up mispricing for either VIX futures or variance swap forwards, then it should remain significant at predicting the returns in these markets after hedging with returns from another market, namely the stock market.

Table 6 reports the return predictability results for VIX futures hedged with CRSP value-weighted returns. As before, VIX futures returns are first regressed onto stock market returns and then the hedged returns are regressed onto the deviation measure. The explanatory power of the stock market is somewhat weaker for VIX futures returns than using variance swap forward returns for hedging as in the previous table (see the Appendix). Despite this, the results indicate that the deviation measure is still significant at predicting VIX futures returns hedged with the stock market. While the predictive power is somewhat lower, in

Panel A, the point estimates are all positive and significant, except for the third contract which is only positive. Similar to the previous results, Panel A.II shows that the results are relatively stable after including additional predictors like the VIX and realized variance to proxy for the variance risk premium. During the post-crisis period, Panel B indicates that the point estimates remain significant or become more significant in the case of the third contract. Interpreting these results, while hedging VIX futures with variance swap forwards from the previous table may be closer to a textbook arbitrage, the results from Table 6 indicate that the deviation measure still exhibits some predictability for VIX futures even when they are hedged with the stock market.

Table 7 then performs the analogous regressions for variance swap forwards hedged with CRSP value-weighted returns. In this case, the expected sign on the deviation measure is negative to the extent that the deviation measure picks up mispricing in variance swap forwards relative to VIX futures. In Panel A, there is some evidence of predictability for variance swap forwards matched to contracts three to six, but no evidence for contracts one and two which have the wrong sign. In Panel B.II, deviation predicts hedged variance swap forward returns with maturities matched to the fifth and sixth contract, but not for contracts one through four. The average R^2 and t-statistic in Panel B.II across contracts are 2.2% and -1.68 which contrasts 4% and 3.4 in Table 6 when VIX futures are hedged with the stock market. While there is some evidence of predictability for variance swap forward rates, compared to the previous table, the results suggest that the deviation measure is picking up larger amounts of predictability for VIX futures than for variance swap forward rates.

5.3 Relative Value Trading Strategy

The results from the previous section indicate that the no-arbitrage deviation measure significantly predicts the returns of VIX futures hedged with variance swap forwards across various in-sample specifications. How robust is this predictability over time and how can its magnitude be interpreted? Does accounting for financing or transaction costs eliminate the positive CAPM alpha?

One way to investigate these questions is with a trading strategy. To that end, the paper computes the returns from a basic strategy that sells (buys) VIX futures contracts when the deviation measure exceeds a high (low) threshold value and hedges with either variance swap forwards or stock market returns. To normalize the deviation measure within contract and over time, the strategy converts the deviation measure into a rolling z-score that is computed using one-year of lagged data,

$$Z_{t,n} \equiv \frac{\text{Deviation}_{t,n} - \mu_{t,n}}{\sigma_{t,n}}. \tag{9}$$

Similarly, the hedge ratio for each contract $\beta_{t,n}$ is computed from rolling regressions with one-year of lagged data. The regressions are analogous to the first step in the return predictability regression from the previous section, except that they use rolling instead of full sample data. This approach accounts for any time-variation in the hedge ratios and ensures that the hedge ratios are in the investor's information set.

Returns for the baseline trading strategy for a short position in the n-th VIX futures contract are defined as,

$$R_{t+h,n} \equiv \frac{1000 \cdot (Fut_{t,n} - Fut_{t+h,n} - \beta^{hedge}_{t,n} \cdot R^{hedge}_{t+h,n}) - Margin_{t,n} \cdot Rf_t \cdot \frac{h}{365}}{Margin_{t,n}}. \tag{10}$$

This definition uses maximal leverage to margin following Garleanu and Pedersen (2011). The payoff is the change in the VIX futures price minus the hedging return times the contract multiplier. The margin is the initial margin for the n-th VIX futures contract. The financing cost is the margin multiplied by the risk-free rate times the actual number of days h over 365 days in a year. The proxy for the risk-free rate is the three-month U.S. Treasury bill rate. The return for a long position is defined analogously.

One aspect of the return definition is that it abstracts from the financing cost for the hedging return. On one hand, this assumption may be conservative. If an exchange or bilateral counterparty offered portfolio margin, the capital requirement for the hedged trade would be lower than the margin for a naked position in VIX futures because of the reduced risk due to the hedge. On the other hand, if an investor was subject to a leverage constraint or if margin was required for both the VIX futures trade and hedge, more capital might be required. The paper focuses on the margin for VIX futures contracts because that data is available historically. It is less clear, for example, what the margin would be on a bilateral, over-the-counter variance swap forward trade in the pre-crisis or post-crisis era. While some of the results are sensitive to this margin assumption, most are not. For example, the CAPM α-to-margin is determined in part by the maximum leverage, which is directly related to the margin assumption. In contrast, doubling or tripling the margin would have no impact on the Sharpe ratio of the strategy.

The return predictability regressions indicate that the deviation measure significantly predicts the returns of VIX futures hedged with variance swap forwards across the front six contracts. Motivated by this, the paper focuses on a baseline trading strategy that hedges with variance swap forwards, has a one-week horizon $h = 5$, and uses a threshold value of $\tau = .50$ z-scores to trade the front six contracts. If there are multiple contracts to trade on the same day, the strategy forms an equal-weighted portfolio. As a robustness check, the paper also explores alternative strategies that vary the threshold value, number of contracts traded, and hedging instrument.

Figure 6 plots the performance of the relative value strategy in comparison to the stock market. The returns are normalized to 10% annualized volatility for comparison. From 2004 to 2018 the VIX futures trading strategy earns a Sharpe Ratio (SR) of 3.0 versus .5 for CRSP value-weighted returns. The plot reports the cumulative sum of overlapping weekly returns for both series. The VIX futures strategy performs well across market environments with minimal drawdowns compared to the stock market. Even during the financial crisis, the VIX futures strategy continues to earn positive returns.

Table 8 builds on these results by reporting summary statistics for the trading strategy returns compared to stock market and volatility factor returns. Panel A reports summary statistics for the relative value trading strategy across specifications and time periods. Panel B reports summary statistics for stock market and volatility factor returns. The first column in Panel A corresponds to the performance of the baseline strategy from Figure 6 that trades the front six contracts with a threshold of $\tau = .50$ and hedges with variance swap forwards (VSF). The strategy returns are adjusted to have 1.39% weekly volatility for $\sqrt{52}\times$ 1.39% = 10% annualized volatility as in the time-series plot. The weekly SR of .42 annualizes to $\sqrt{52}\times$ .42 = 3.0 matching Figure 6. Beyond this metric, the strategy delivers positively skewed, fat-tailed returns that are only negative 27% of the time. The maximum drawdown of 6.7% contrasts a maximum drawdown of around 50% for the stock market over the same period of time (Panel B, column 1). To obtain 10% annualized volatility, the returns are de-levered and multiplied by 11.6%. The unlevered mean return is 4.99%. This is close to the alpha-to-margin of 4.90% per week, indicating that the maximally leveraged strategy delivers large returns that are largely unexplained by the market factor.

The subsequent columns in Table 8 vary aspects of the trading strategy or sample period to examine the robustness of the results. The second column (2) in Panel A uses stock market returns ($RMRF$) as a hedge instead of variance swap forward returns. Similar to the return predictability regressions, the performance declines somewhat but the strategy still obtains an annualized SR of 1.43 and CAPM alpha-to-margin of 3.6%. The returns remain positively skewed and continue to exhibit a small maximum drawdown of 10%.

The third column (3) in Panel A adds transaction costs to the strategy from column (2) that hedges with the stock market. The returns with transaction costs incorporate the bid-ask spread from closing quote data for VIX futures. For example, when selling VIX futures, the strategy sells at the bid and buys the contract back at the ask a week later. The strategy does not trade if the bid-ask spread exceeds .25 volatility units, which is five times the typical bid-ask spread in recent years. This constraint is included to avoid trading during illiquid periods of time when it may be costly or difficult to execute the strategy. As will be shown below, the strategy with transaction costs performs better when trading the front two contracts instead of the front six contracts. The front two contracts are more liquid and have lower transaction costs on average. Using the same threshold of $\tau = .50$ as before, the strategy with transaction costs and a stock market hedge earns an annualized SR of .73 with a CAPM alpha-to-margin of 2.6% per week. The returns continue to be positively skewed and are only negative 32% of the time. The next columns (4) to (6) show that similar results hold in a post-crisis sample from 2010 to 2018. The SRs and CAPM alpha-to-margin are slightly lower, but still large. The maximum drawdown and percentage of negative returns are similar to the full sample estimates.

Panel B presents summary statistics for stock market and volatility returns for comparison. Column (1) reports results for CRSP value-weighted returns (RMRF). Columns (2) and (3) report results for receiving fixed in one-month variance swaps and selling the front-month VIX futures contracts. The horizon is one-week and the returns are computed over the same period of time as for the trading strategy returns in Panel A. The annualized SRs for the stock market, variance swap, and VIX futures returns are .51, 1.08, and .65. While the relative value strategy earns a higher SR than the stock market and the strategy of unconditionally selling front month VIX futures, the SR from receiving fixed in one-month variance swaps is similar to the relative value strategy when hedging with stock market returns. Unlike the relative value strategy, the stock market and volatility selling returns are negatively skewed. The stock market returns and returns from unconditionally selling the front-month VIX futures contracts also exhibit large drawdowns during the financial crisis.

The summary statistics indicate that the trading strategy earns excess returns relative to the CAPM as measured by the large alpha-to-margin estimates. How significant is this the outperformance and is it robust to other factor models? Table 9 answers this question by reporting alpha-to-margin estimates for different specifications including the CAPM, a four-factor model with the Fama-French three factors and Carhart momentum factor (FFC4), and a six-factor model (FFCV6) that adds realized and implied volatility factor returns from one-month variance swap forwards (VS1) and front-month VIX futures contracts (VX1). Columns 1-3 report results for the full sample period and columns 4-6 report results for a post-crisis sample from 2010 to 2018.

The baseline strategy that hedges with variance swap forwards earns an alpha-to-margin of 4% to 5% across specifications and sample periods. The insignificant point estimates on the factors and low R2s indicate that the variance swap forward hedge has removed almost all of the systematic risk in the trading strategy returns. The large alpha and low explanatory power of the factor models is consistent with the idea of trading against an arbitrage spread. Panels B and C find similar qualitative results for the other trading strategy specifications. The alpha estimates remain significant and the explanatory power of the factor models remains low across specifications. While the volatility factors are more significant in Panels B and C than in Panel A, the highest R2s are only 10% to 15% and the negative loadings on the volatility factors generally increase the alpha estimates. Overall the alpha estimates appear robust to different factor model specifications and time periods.

The top plot in Figure 7 reports the performance of the three main specifications that hedge with variance swap forward rates, the stock market, and the stock market with transaction costs. The relative value strategy hedged with variance swap forwards or stock market returns outperforms the stock market over nearly the entire sample period. The strategy with transaction costs underperforms until the financial crisis and then starts to outperform the stock market. The higher values for the relative value strategies at the end of the sample relative to the stock market are consistent with their higher SRs from the summary statistics table. Note also that the stock market is a difficult benchmark to outperform during this sample period. Alternative equity risk factors such as value (HML), small-minus-big (SMB), and momentum (MOM) generally underperformed the stock market, earning annualized SRs of .01, .08, and .12, that also underperform the relative value strategy.

Is the outperformance of the relative value strategies driven by the trading threshold or number of contracts traded? The bottom plot in Figure 7 investigates this question. There are two takeaways for the baseline strategy that hedges with variance swap forward rates. First, as the number of contracts traded increases, the SR increases. This result is consistent with the return predictability regressions. Since the deviation measure predicts returns across all contracts, trading more contracts provides the strategy with more relative value opportunities to exploit, improving performance. For example, the SRs with a threshold of $\tau = .50$ when trading up to the front 1, 2, 4, and 6 contracts are .92, 1.58, 2.63, and 3.00.

The second observation is that the SRs tend to peak at an intermediate trading threshold τ and then decline. For large thresholds the strategy trades less frequently and only against large deviations, passing up valuable trading opportunities. This pattern is slightly different, however, when including transaction costs. In this case, it is still desirable to trade against large opportunities, but there is also a benefit to trading less frequently to lower transaction costs. This interpretation echoes findings from the literature on optimal portfolio choice with transaction costs (Davis and Norman 1990; Gârleanu and Pedersen 2013). When incorporating transaction costs, the SRs do not decline as rapidly with the trading threshold and in some cases it is better to use a larger threshold. The SRs are also higher when only trading the front 1 or 1-2 contracts, reflecting the lower transaction costs of these contracts relative to the longer-dated contracts. The Appendix shows that similar results hold in the 2010 to 2018 post-crisis period.

In summary, the results indicate that trading against the no-arbitrage deviation measure earns significant returns with minimal exposure to traditional risk factors. The alpha-to-margin estimates and high SRs are robust across a range of specifications and sample periods. The strongest results correspond to the strategy that hedges with variance swap forwards. This strategy is most similar to the textbook arbitrage trades between VIX futures and variance swap forwards or volatility swap forwards described earlier in the paper.

6 Discussion

6.1 No-Arbitrage Deviation versus Risk and Demand Factors

What drives the no-arbitrage violations and deviation measure over time? Two channels identified by the limits-to-arbitrage literature are risk and demand. If arbitrageurs are risk averse or have limited capital, demand shocks can push prices away from fundamental values. The resulting effect of limits-to-arbitrage and demand shocks can be anomalies such as return predictability and no-arbitrage violations.

The paper estimates a vector autoregression (VAR) to study how the no-arbitrage deviation measure relates to risk and demand factors, $y_t = [DEV_t \, VIX_t \, DNP_t]$. The variables in the VAR are the average deviation measure across the front six contracts (DEV), the CBOE Volatility index (VIX), and the dealer net position (DNP) in VIX futures from the CFTC's Commitment of Traders Report (CoT). The DEV variable focuses on the average deviation to keep the size of the VAR small. The VIX and DNP variables serve as proxies for risk and demand. Figure 1 motivates using DNP as a demand variable by illustrating the high correlation between DNP and VIX ETP demand, a proxy for retail demand. The advantage of DEV relative to VIX ETP demand is that it only depends on quantities, not on prices. The DNP variable in the VAR is the dealer net position as a fraction of open interest. This normalization bounds DNP between 0 and 1, removing the time trend in net position size that reflects the growing VIX futures market over the sample period. The Appendix shows that similar qualitative results hold using other proxies for risk and demand such as stock market returns and realized variance for risk and VIX ETP demand and the delta of VIX options traded by retail customers for demand.

The sample period is 2010 to 2018 using weekly observations on CoT release dates when the DNP variable is reported. The sample period is motivated by several observations. First, the start date matches prior studies investigating the relationship between the pricing of VIX futures and demand (Cheng 2018). During this period there are no breaks in the reporting of the DNP variable.[5] In addition, the sample corresponds to a post-crisis period when the VIX ETP market is growing. Dong (2016) argues that the introduction of VIX ETP trading represents a structural break in the VIX futures market, with ETPs introducing new channels for demand to impact futures prices. Finally, the sample period is beneficial because a balanced panel is available to compute the average deviation measure across the front six futures contracts.

Figure 8 presents the time-series relationship between the average deviation and the risk and demand variables that are used in the VAR. The top plot shows that the deviation measure decreases when the VIX increases. Since the deviation measure tracks the difference in prices for nearly identical claims across two markets, it is surprising that the no-arbitrage deviation responds to risk, as risk should play a similar role in both markets. One hypothesis for what drives the result is that increases in risk may prevent traders from being able to engage in arbitrage trades that would drive prices back to fundamental values, perhaps as a result of binding margin or value-at-risk constraints. Alternatively, hedgers may take profit on long positions when risk increases, leading to demand shocks that are correlated with increases in risk and traders temporary inability to exploit arbitrage opportunities due to financing constraints. In addition to the correlation with risk, the bottom plot shows that the deviation measure is also highly correlated with the DNP demand variable. This correlation may be driven by systematic demand shocks as discussed above or by idiosyncratic shocks such as mechanical roll effects from VIX ETPs. To more precisely identify how the no-arbitrage deviation measure relates to risk and demand and to better understand the lead-lag relationships, the paper estimates a VAR and studies its associated impulse response functions.

6.2 VAR Impulse Response Functions

The paper estimates a trivariate VAR for the variables $y_t = [DEV_t \; VIX_t \; DNP_t]$ at a weekly frequency from 2010 to 2018. Figure 9 reports impulse response functions (IRFs) from the VAR. The IRFs are from a Cholesky decomposition with the variables ordered as: VIX, DNP, DEV. The optimal lag length is selected by the SBIC criterion. The Appendix reports the IRFs with different orderings as a robustness check and finds qualitatively similar results.

[5] There is a gap in CoT report for VIX futures from December 2008 to June 2009 when open interest was low and the position breakdown by trader type was not reported.

The top row in Figure 9 reports the IRFs for DEV in response to VIX and DNP shocks. The top left plot shows that a one-standard deviation increase in VIX corresponds to a .25 standard deviation decrease in DEV that mean reverts after four to six weeks. The top right plot shows that a one standard deviation increase in DNP corresponds to a .05 to .10 standard deviation increase in DEV that mean reverts over a longer horizon. These responses are large in the sense that they represent several bid-ask spreads in VIX futures and are the same order of magnitude as the coefficients on the deviation measure in the return predictability regressions. The impact of the VIX shock is about 3-4 times larger in magnitude than the impact of a demand shock over short horizons. Despite the tight relationship between the deviation measure and demand in the time-series plot, the VAR indicates that the risk shock is more significant over shorter horizons. For longer horizons the demand shock remains more significant and has a slightly larger magnitude than the VIX shock. Taken together, the results indicate that both the risk and demand channels have an impact on the no-arbitrage deviation measure over different horizons.

The middle row in Figure 9 reveals how the DNP variable reacts to VIX and DEV shocks. The middle left plot shows that a one standard deviation increase in the VIX corresponds to a .10 decrease in DNP that persists for one to three months. The middle right plot shows that a one standard deviation increase in DEV corresponds to an increase in DNP by .10 standard deviations that peaks after one to two months and then persists over a longer period of time. These results have a mixed interpretation. On one hand, DNP decreases when risk increases. This result is consistent with dealers acting as hedgers that take profit on long positions when risk increases. On the other hand, DNP increases when the deviation measure increases. This result suggests that dealers also act as momentum traders, increasing their long position in VIX futures when the prices of VIX futures increase relative to variance swap forward rates. To the extent that the DNP variable is a veil for retail demand, the results suggest that retail traders use VIX ETPs to hedge volatility risk and chase momentum.

The bottom row Figure 9 reports how the VIX responds to DNP and DEV shocks. The responses are largely insignificant. This result provides a reassuring placebo test. One would not expect changes in DNP or DEV to impact the VIX unless changes in these variables led to arbitrage trading that moved index option prices and thus the VIX. For example, if DEV increases, arbitrage traders might sell VIX futures and pay fixed in variance swap forwards. Implementing the variance swap forward trade synthetically could put downward pressure on the VIX if traders pay fixed in long-dated synthetic variance swaps by buying long-dated options and receive fixed in short-dated synthetic variance swap rates by selling short-dated options. The IRF in the bottom right plot provides some evidence that there is a short-term negative effect of DEV shocks on the VIX over short horizons. Over longer horizons, however, the response becomes insignificant and the point estimate is close to zero. Given the large size of the index options market relative to the VIX futures market, the insignificant responses that are consistent with the null hypothesis of no effect on the VIX seem most plausible.

6.3 Panel Regressions

Another way to study the deviation measure and its relationship with risk and demand is by exploiting a panel-regression approach. While the VAR accounts for the persistence of the different variables and their joint interactions, it can be difficult to interpret large-scale VARs. This observation motivates estimating the VAR with only three variables, one of which is the average deviation measure across the front six contracts. In a panel-regression, the deviation measure for different contracts can be studied directly along with fixed effects to study within contract variation and control variables to account for the economic environment.

Table 10 reports panel regressions that find similar relationships between the deviation measure and the risk and demand variables despite the different identification approach from the VAR. The regression specification is,

$$\text{Deviation}_{t+h,n} = \beta \Delta x_{t+h} + \rho \text{Deviation}_{t,n} + \delta \text{Controls}_t + \text{FEs} + \epsilon_{t+h,n}. \tag{11}$$

The horizon is one-week $h = 5$ to be similar to the VAR analysis. The regressions include overlapping observations from daily data using the front six contracts. The explanatory variable Δx_{t+h} and control variables are standardized but the deviation measure is not. The summary statistics for the deviation measure in Table 2 show that the deviation measure has an average standard deviation of about .80%. This makes the results roughly comparable to the IRFs from the VAR where the variables are standardized. The first three specifications 1-3 in Table 10 show how the deviation measure responds to changes in the explanatory variable controlling for the persistence of the deviation measure. The next three specifications 4-6 repeat this analysis adding control variables and fixed effects. The control variables include the time-to-maturity, initial margin, and open interest for the n-th contract on date t as well as the lagged VIX index. Fixed effects are included for the contract, calendar year, and for the contract being the first, second, etc. contract to maturity.

Panels A reports the results for the risk variables from 2007 to 2018. Panel B finds similar results from 2010 to 2018, showing that the results are not driven by the financial crisis. As in the VAR analysis, the results indicate that the deviation measure is decreasing in risk. When stock market returns are negative or volatility increases, the deviation measure tends to decline. The magnitude of the point estimates is similar to the IRFs. For example, in the panel regressions, a one-standard deviation increase in the VIX is associated with a decline in the deviation measure of around -.15 in Panel A or -.18 in Panel B. In the VAR, the IRF of the deviation measure to a VIX shock has a point estimate of around -.10 to -.25 in the first few weeks after a shock. Despite the different approaches, the VAR and panel regression reveal a similar relationship between the deviation measure and VIX index. Moreover, the panel regressions highlight how the relationship between the deviation measure and risk is not specific to the VIX index, but also holds for stock market returns and for the realized variance of the S&P 500 index.

Panel C reports the results for the demand variables from 2010 to 2018. The post-crisis sample period is motivated by data availability. As in the VAR analysis, the deviation measure is increasing in demand pressure. A one-standard deviation increase in the dealer net position corresponds to a .05 standard-deviation increase in the deviation measure. This response is similar to the IRF in which the deviation measure increases by around .05 to .10 in response to a DNP shock in the weeks following the shock. Beyond the DNP variable, the regressions show that the deviation measure is also increasing in VIX ETP net demand and in the delta of VIX options traded by retail customers. Similar to the DNP variable, the ETP and option demand variables are normalized by open interest and then standardized. Compared to the risk variables in Panel B, the point estimates and explanatory power of the demand variables is slightly lower in Panel C. Over short horizons, changes in risk are associated with larger changes in the deviation measure than changes in demand, similar to the VAR analysis.

6.4 Other Channels

Beyond demand pressure, the limits-to-arbitrage literature highlights a number of additional channels that can lead to no-arbitrage violations and return predictability. For example, changes in margin requirements may lead to binding financial constraints, resulting in margin spirals if arbitrageurs are forced to unwind positions (Brunnermeier and Pedersen 2008). End-of-month and end-of-quarter dates can result in funding pressure that is associated with heightened arbitrage deviations in the foreign exchange market (Du et al. 2018). Salient events like FOMC announcements may result in return predictability (Lucca and Moench 2015). How does the no-arbitrage deviation measure studied in this paper relate to these events?

Figure 10 investigates this question by reporting event study plots to illustrate how the deviation measure and VIX index respond to various events. The top row reports results for margin changes. There are 41 increases and 33 decreases in the initial margin for the front month VIX futures contract from 2004 to 2018. In the two weeks leading up to a margin increase, the VIX index tends to increase by around four points. During this time, the deviation measure decreases by around .20 units. Given the small number of events, this decline is not statistically significant. After the margin increase, the VIX tends to decrease and the deviation measure tends to increase which is a decline in magnitude. From the plot, there does not appear to be a very strong response of the deviation measure to margin increases or decreases, aside from the typical reaction to changes in the VIX. The subsequent plots highlight similar null results for FOMC announcements, non-farm payroll announcements, month-end dates, and quarter-end dates. There may be a small decrease in the deviation measure on the day following FOMC announcements, but otherwise the deviation measure seems largely unrelated to these events. The Appendix provides related analysis for the risk and demand variables by highlighting how the deviation measure responds to the largest changes in risk and demand over the sample period. While the magnitudes are different, the qualitative results are similar with changes in risk and demand being associated with the largest changes in the deviation measure.

7 Conclusion

This paper documents systematic law of one price deviations across the VIX futures and S&P 500 index options markets. The prices of VIX futures violate estimates of their no-arbitrage upper and lower bounds, indicating the presence of arbitrage opportunities. A no-arbitrage deviation measure equal to the difference between the price of a VIX futures contract and its corresponding synthetic variance swap forward rate is found to significantly predicts returns. A relative value trading strategy that exploits the deviation measure earns a significant Sharpe ratio with minimal exposure to traditional risk factors.

These results are surprising because the no-arbitrage relationships investigated in the paper are well understood in the academic literature and should be known by option traders. Even if it is difficult to trade VIX futures against synthetic variance swap forwards in practice, there is still significant predictability when using the deviation measure to predict VIX futures returns hedged with the stock market, and a sizeable Sharpe ratio and alpha-to-margin estimate persist even after taking VIX futures transaction costs into account. Thus, the results cannot be explained by a mere appeal to implementation challenges or transaction costs.

Instead, the paper finds evidence that the no-arbitrage deviations are related to systematic risk and demand pressure. Other channels like margin changes, economic announcements, and month-end effects are less significant in explaining the variation in the deviations over time. To the extent that dealer positions reflect retail demand, the results are consistent with retail traders using VIX ETPs to hedge volatility risk and chase momentum, driving the prices of VIX futures away from fundamental values implied by the index option market.

An implication of these results is that investors and policymakers should be cautious when interpreting signals from equity volatility markets. Large no-arbitrage deviations indicate that the VIX futures and index option markets are sending different messages about future risks.

References

Adrian, T. and H. S. Shin (2013). Procyclical leverage and value-at-risk. *The Review of Financial Studies 27*(2), 373–403.

Andries, M., T. Eisenbach, M. Schmalz, and Y. Wang (2015). The term structure of the price of variance risk. *Working Paper*.

Bai, J. and P. Collin-Dufresne (2018). The cds-bond basis. *Financial Management*.

Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics 55*(2), 205–238.

Bates, D. (2000). Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics 94*, 181–238.

Bates, D. S. (2003). Empirical option pricing: A retrospection. *Journal of Econometrics 116*(1-2), 387–404.

Bollen, N. P. and R. E. Whaley (2004). Does net buying pressure affect the shape of implied volatility functions? *The Journal of Finance 59*(2), 711–753.

Breeden, D. T. and R. H. Litzenberger (1978). Prices of state contingent claims implicit in option prices. *Journal of Business 51*(4), 621–651.

Britten-Jones, M. and A. Neuberger (2000). Option prices, implied price processes, and stochastic volatility. *The Journal of Finance 55*(2), 839–866.

Brunnermeier, M. K. and L. H. Pedersen (2008). Market liquidity and funding liquidity. *The Review of Financial Studies 22*(6), 2201–2238.

Carr, P. and R. Lee (2009). Volatility derivatives. *Annu. Rev. Financ. Econ. 1*(1), 319–339.

Carr, P. and D. Madan (1999). Introducing the covariance swap. *Risk*, 47–51.

Carr, P. and L. Wu (2006). A tale of two indices. *The Journal of Derivatives 13*(3), 13–29.

Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies 22*(3), 1311–1341.

CBOE (2019). The cboe volatility index. *https://www.cboe.com/micro/vix/vixwhite.pdf*.

Cheng, I.-H. (2018). The vix premium. *The Review of Financial Studies 32*(1), 180–227.

Coval, J. D. and T. Shumway (2001). Expected option returns. *The journal of Finance 56*(3), 983–1009.

Davis, M. H. and A. R. Norman (1990). Portfolio selection with transaction costs. *Mathematics of operations research 15*(4), 676–713.

Demeterfi, K., E. Derman, M. Kamal, and J. Zou (1999). A guide to volatility and variance swaps. *The Journal of Derivatives 6*(4), 9–32.

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2017). The price of variance risk. *Journal of Financial Economics 123*(2), 225–250.

Dong, X. S. (2016). Price impact of ETP demand on underliers. *Working Paper*.

Du, W., A. Tepper, and A. Verdelhan (2018). Deviations from covered interest rate parity. *The Journal of Finance 73*(3), 915–957.

Duarte, J., F. A. Longstaff, and F. Yu (2007). Risk and return in fixed-income arbitrage: Nickels in front of a steamroller? *Review of Financial Studies 20*(3), 769–811.

Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica 68*(6), 1343–1376.

Fleckenstein, M., F. A. Longstaff, and H. Lustig (2014). The tips-treasury bond puzzle. *The Journal of Finance 69*(5), 2151–2197.

Fong, W.-M., G. Valente, and J. K. Fung (2010). Covered interest arbitrage profits: The role of liquidity and credit risk. *Journal of Banking & Finance 34*(5), 1098–1107.

Gabaix, X., A. Krishnamurthy, and O. Vigneron (2007). Limits of arbitrage: theory and evidence from the mortgage-backed securities market. *The Journal of Finance 62*(2), 557–595.

Gagnon, L. and G. A. Karolyi (2010). Multi-market trading and arbitrage. *Journal of Financial Economics 97*(1), 53–80.

Garleanu, N. and L. H. Pedersen (2011). Margin-based asset pricing and deviations from the law of one price. *The Review of Financial Studies 24*(6), 1980–2022.

Gârleanu, N. and L. H. Pedersen (2013). Dynamic trading with predictable returns and transaction costs. *The Journal of Finance 68*(6), 2309–2340.

Garleanu, N., L. H. Pedersen, and A. M. Poteshman (2009). Demand-based option pricing. *Review of Financial Studies 22*(10), 4259–4299.

Gromb, D. and D. Vayanos (2010). Limits of arbitrage. *Annu. Rev. Financ. Econ. 2*(1), 251–275.

Grossman, S. J. and Z. Zhou (1996). Equilibrium analysis of portfolio insurance. *The Journal of Finance 51*(4), 1379–1403.

Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of financial studies 6*(2), 327–343.

Johnson, T. L. (2017). Risk premia and the vix term structure. *Journal of Financial and Quantitative Analysis 52*(6), 2461–2490.

Jurek, J. W. and E. Stafford (2015). The cost of capital for alternative investments. *The Journal of Finance 70*(5), 2185–2226.

Krishnamurthy, A. (2002). The bond/old-bond spread. *Journal of Financial Economics 66*(2-3), 463–506.

Lamont, O. A. and R. H. Thaler (2003). Anomalies: The law of one price in financial markets. *Journal of Economic Perspectives 17*(4), 191–202.

Lee, C. M., A. Shleifer, and R. H. Thaler (1991). Investor sentiment and the closed-end fund puzzle. *The Journal of Finance 46*(1), 75–109.

Liu, J. and F. A. Longstaff (2003). Losing money on arbitrage: Optimal dynamic portfolio choice in markets with arbitrage opportunities. *The Review of Financial Studies 17*(3), 611–641.

Longstaff, F. A., P. Santa-Clara, and E. S. Schwartz (2001). The relative valuation of caps and swaptions: Theory and empirical evidence. *The Journal of Finance 56*(6), 2067–2109.

Lucca, D. O. and E. Moench (2015). The pre-fomc announcement drift. *The Journal of Finance 70*(1), 329–371.

Madan, D. B., P. P. Carr, and E. C. Chang (1998). The variance gamma process and option pricing. *Review of Finance 2*(1), 79–105.

Martin, I. (2017). What is the expected return on the market? *The Quarterly Journal of Economics 132*(1), 367–433.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science 4*, 141–183.

Mitchell, M., T. Pulvino, and E. Stafford (2002). Limited arbitrage in equity markets. *The Journal of Finance 57*(2), 551–584.

Mixon, S. and E. Onur (2019). Derivatives pricing when supply and demand matter: Evidence from the term structure of vix futures. *Journal of Futures Markets 39*(9), 1035–1055.

Park, Y.-H. (2019). Variance disparity and market frictions. *Working Paper*.

Ross, S. A. (1976a). The arbitrage theory of capital asset pricing. *Journal of Economic Theory 13*(3), 341–360.

Ross, S. A. (1976b). Options and efficiency. *Quarterly Journal of Economics 90*(1), 75–89.

Shleifer, A. (1986). Do demand curves for stocks slope down? *The Journal of Finance 41*(3), 579–590.

Shleifer, A., J. B. DeLong, L. Summers, and R. Waldmann (1990). Noise trader risk in financial markets. *Journal of Political Economy 98*(4), 703–738.

Van Tassel, P. (2019). Volatility term premia. *Working Paper*.

Table 1: Computing the Deviation Measure on Example Dates

This table illustrates two examples of how the paper computes the no-arbitrage deviation measure each day. First, the paper obtains the variance swap curve in annualized variance units. Second, the paper computes the variance swap forward curve assuming flat forward rates between the observed maturities $Fwd_{t,T_1,T_2} = (VS_{t,T_2} \cdot T_2 - VS_{t,T_1} \cdot T_1)/(T_2 - T_1)$. Third, the paper computes the difference between the VIX futures price and the one-month variance swap forward rate for the corresponding maturity. The deviation measure $\text{Deviation}_{t,n}$ for the n-th contract is the difference between the VIX futures price $\text{Fut}_{t,n}$ and the one-month variance swap forward rate $\text{Fwd}_{t,n}$.

Example: Computation of Deviation Measure on 27Feb12

Step 1: Variance swap rates (variance units)									
Maturity	1	2	3	6	9	12	15	18	24
$VS_{t,T}$	3.39	3.90	4.55	5.83	6.44	6.88	7.18	7.52	8.16
Step 2: Variance swap forward rates (variance units)									
Maturity	0-1	1-2	2-3	3-6	6-9	9-12	12-15	15-18	18-24
Fwd_{t,T_1,T_2}	3.39	4.40	5.87	7.11	7.66	8.21	8.37	9.24	10.07
Step 3: Deviation measure $\text{Deviation}_{t,n} = \text{Fut}_{t,n} - \text{Fwd}_{t,n}$ (volatility units)									
Maturity	21Mar	18Apr	16May	20Jun	18Jul	22Aug	19Sep	17Oct	21Nov
Contract	1	2	3	4	5	6	7	8	9
$\text{Fut}_{t,n}$	21.40	24.38	25.87	27.02	28.17	28.78	29.65	29.75	29.75
$\text{Fwd}_{t,n}$	20.35	23.23	25.67	26.67	26.67	27.48	27.67	27.67	28.47
$\text{Deviation}_{t,n}$	1.05	1.14	0.20	0.36	1.51	1.30	1.98	2.08	1.28

Example: Computation of Deviation Measure on 10Aug11

Step 1: Variancee swap rates (variance units)									
Maturity	1	2	3	6	9	12	15	18	24
$VS_{t,T}$	21.00	15.89	14.14	11.94	11.11	10.72	10.49	10.28	10.03
Step 2: Variance swap forward rates (variance units)									
Maturity	0-1	1-2	2-3	3-6	6-9	9-12	12-15	15-18	18-24
Fwd_{t,T_1,T_2}	21.00	10.77	10.65	9.73	9.45	9.56	9.58	9.24	9.28
Step 3: Deviation measure $\text{Deviation}_{t,n} = \text{Fut}_{t,n} - \text{Fwd}_{t,n}$ (volatility units)									
Maturity	17Aug	21Sep	19Oct	16Nov	21Dec	18Jan	15Feb	21Mar	
Contract	1	2	3	4	5	6	7	8	
$\text{Fut}_{t,n}$	36.25	28.98	28.00	27.25	26.10	27.33	27.80	27.80	
$\text{Fwd}_{t,n}$	43.23	32.75	32.23	31.20	31.20	31.07	30.74	30.74	
$\text{Deviation}_{t,n}$	-6.98	-3.78	-4.23	-3.95	-5.10	-3.74	-2.94	-2.94	

Table 2: Summary Statistics for Deviation Measure

This table reports summary statistics for the law of one price deviations between VIX futures and variance swap forwards. The deviation measure is defined as $\text{Deviation}_{t,n} = \text{Fut}_{t,n} - \text{Fwd}_{t,n}$ in annualized volatility units. Panel A reports summary statistics for the deviation variable. Panel B reports the correlation of the deviations across contracts. Panel C reports the average deviation by year for each contract. The sample period is 2007 to 2018 to obtain a balanced panel across contracts.

	$\text{Deviation}_{t,n}$ Summary Statistics						
Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)	Avg.
Panel A: Summary Statistics							
Mean	-0.25	-0.47	-1.02	-1.26	-0.97	-1.08	-0.84
Standard Deviation	0.82	0.74	0.77	0.78	0.88	0.93	0.82
t-statistic	-5.06	-7.34	-13.37	-15.38	-10.17	-11.07	-10.40
Skewness	-3.68	-1.72	0.32	-0.01	-0.83	-0.78	-1.12
Kurtosis	31.49	14.07	5.50	3.13	5.11	5.80	10.85
Minimum	-8.87	-7.70	-4.89	-4.27	-5.86	-7.15	-6.46
25th-Percentile	-0.46	-0.78	-1.50	-1.75	-1.40	-1.60	-1.25
Median	-0.15	-0.41	-1.04	-1.26	-0.89	-1.03	-0.80
75th-Percentile	0.13	-0.08	-0.52	-0.75	-0.40	-0.48	-0.35
Maximum	4.26	4.42	4.22	2.16	1.58	1.44	3.01
Autocorrelation 1-day	0.56	0.81	0.87	0.90	0.92	0.90	0.83
Autocorrelation 5-day	0.33	0.65	0.73	0.80	0.84	0.78	0.69
Autocorrelation 21-day	0.17	0.36	0.52	0.61	0.65	0.64	0.49
Panel B: Correlation Matrix							
$\text{Deviation}_{t,1}$	1.00	0.38	0.06	0.06	0.20	0.19	0.18
$\text{Deviation}_{t,2}$	0.38	1.00	0.57	0.53	0.62	0.61	0.54
$\text{Deviation}_{t,3}$	0.06	0.57	1.00	0.72	0.53	0.55	0.49
$\text{Deviation}_{t,4}$	0.06	0.53	0.72	1.00	0.74	0.54	0.52
$\text{Deviation}_{t,5}$	0.20	0.62	0.53	0.74	1.00	0.83	0.58
$\text{Deviation}_{t,6}$	0.19	0.61	0.55	0.54	0.83	1.00	0.54
Panel C: Average Deviation by Year							
2007	-0.33	-0.26	-0.33	-0.39	-0.31	-0.31	-0.32
2008	-0.69	-0.57	-0.32	-0.66	-1.22	-1.41	-0.81
2009	-0.10	0.11	-0.47	-0.84	-1.12	-1.26	-0.61
2010	-0.44	-0.09	-0.48	-0.62	-0.25	-0.45	-0.39
2011	-0.38	-0.74	-1.38	-1.77	-1.33	-1.35	-1.16
2012	-0.15	-0.29	-1.04	-1.26	-0.23	-0.12	-0.52
2013	-0.09	-0.25	-0.82	-0.95	-0.39	-0.64	-0.53
2014	-0.12	-0.50	-1.26	-1.56	-1.12	-1.31	-0.98
2015	-0.27	-0.90	-1.66	-1.94	-1.64	-1.76	-1.36
2016	-0.14	-0.66	-1.36	-1.58	-1.18	-1.36	-1.05
2017	-0.02	-0.69	-1.55	-1.84	-1.22	-1.32	-1.11
2018	-0.21	-0.85	-1.55	-1.73	-1.57	-1.62	-1.25

Table 3: The Frequency of the Law of One Price Violations

This table reports the frequency of law of one price violations for VIX futures contracts from their no-arbitrage bounds. Panel A reports the frequency of upper bound violations over the full sample from 2004 to 2018 (A.I) and during a post-crisis sample 2010 to 2018 (A.II). An upper bound violation occurs when the deviation measure is greater than zero, $\text{Deviation}_{t,n} > 0$. The table also reports the frequency of violations that are greater than thresholds of .25% and .50%. Panel B reports the analogous frequency of lower bound violations. The lower bound is estimated for each contract-day as the upper bound minus the difference between the upper bound and lower bound from a term-structure model for that contract-day pair. The results indicate that around 20% of contract-date observations exhibit an upper bound or lower bound violations on average.

Frequency of Law of One Price Violations

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)	Avg.
Panel A.I: VIX Futures - Upper Bound > Threshold 2004-2018							
Threshold = 0	0.35	0.24	0.13	0.14	0.14	0.14	0.19
Threshold = .25	0.17	0.13	0.07	0.08	0.09	0.09	0.11
Threshold = .50	0.07	0.06	0.03	0.05	0.06	0.06	0.05
Panel A.II: VIX Futures - Upper Bound > Threshold 2010-2018							
Threshold = 0	0.36	0.15	0.03	0.02	0.09	0.09	0.12
Threshold = .25	0.15	0.07	0.02	0.01	0.05	0.06	0.06
Threshold = .50	0.06	0.03	0.00	0.00	0.03	0.04	0.03
Panel B.I: VIX Futures - Lower Bound < Threshold 2004-2018							
Threshold = 0	0.11	0.12	0.31	0.35	0.19	0.21	0.22
Threshold = -.25	0.06	0.06	0.21	0.25	0.12	0.13	0.14
Threshold = -.50	0.03	0.03	0.12	0.15	0.08	0.08	0.08
Panel B.II: VIX Futures - Lower Bound < Threshold 2010-2018							
Threshold = 0	0.08	0.16	0.49	0.55	0.25	0.28	0.30
Threshold = -.25	0.04	0.08	0.33	0.40	0.15	0.17	0.20
Threshold = -.50	0.03	0.04	0.19	0.25	0.10	0.10	0.12

Table 4: Summary Statistics for VIX Futures and Variance Swap Forwards

Panel A reports summary statistics for VIX futures prices for the front six contracts and for variance swap one-month forward rates for the corresponding maturities. Panel B reports summary statistics for VIX futures and variance swap forward one-week excess returns ($h = 5$). The sample period is 2007 to 2018 for a balanced panel.

VIX Futures and VS Forwards Summary Statistics

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)	Average
Panel A.I: VIX Futures Prices $Fut_{t,n}$							
Mean	20.07	20.81	21.30	21.65	21.94	22.20	21.33
Standard Deviation	8.57	7.66	7.06	6.60	6.28	6.02	7.03
Skewness	2.17	1.82	1.59	1.38	1.21	1.10	1.54
Kurtosis	9.04	7.21	6.20	5.10	4.33	3.90	5.96
Median	17.54	18.58	19.27	19.67	20.08	20.27	19.23
Panel A.II: VS Forward Rates $Fwd_{t,n}$ (annualized volatility units)							
Mean	20.32	21.28	22.32	22.91	22.91	23.28	22.17
Standard Deviation	8.80	7.74	6.91	6.48	6.47	6.21	7.10
Skewness	2.21	1.90	1.62	1.41	1.41	1.32	1.65
Kurtosis	9.50	7.66	6.28	5.21	5.18	4.86	6.45
Median	17.82	19.28	20.64	21.27	21.30	21.62	20.32
Panel B.I: VIX Futures Returns (percent) $R^{VIX}_{t+h,n} = Fut_{t,n} - Fut_{t+h,n}$							
Mean	0.17	0.15	0.09	0.06	0.06	0.05	0.10
Standard Deviation	2.42	1.99	1.57	1.34	1.19	1.10	1.60
Sharpe Ratio	0.07	0.08	0.06	0.05	0.05	0.05	0.06
t-statistic	1.81	1.89	1.48	1.19	1.24	1.23	1.47
Skewness	-1.56	-1.20	-1.03	-0.86	-0.80	-0.74	-1.03
Kurtosis	12.77	10.53	9.42	8.21	8.29	7.35	9.43
Median	0.27	0.30	0.20	0.15	0.10	0.10	0.19
Panel B.I: VS Forward Returns (basis points) $R^{VSF}_{t+h,n} = Fwd_{t,n} - Fwd_{t+h,n}$							
Mean	0.05	0.50	1.01	0.80	-0.38	-0.09	0.31
Standard Deviation	18.98	12.37	9.54	8.21	6.75	6.10	10.33
Sharpe Ratio	0.00	0.04	0.11	0.10	-0.06	-0.02	0.03
t-statistic	0.07	0.99	2.44	2.23	-1.54	-0.45	0.62
Skewness	-3.84	-2.82	-2.84	-2.81	-0.30	-0.09	-2.12
Kurtosis	67.11	45.54	30.56	28.92	15.93	12.59	33.44
Median	1.17	1.07	1.63	1.18	-0.32	0.11	0.81

Table 5: Deviation Measure Predicts the Returns of VIX Futures Hedged with Variance Swap Forwards

This table reports return predictability regressions for hedged VIX futures returns over a weekly horizon (h=5). The columns report results for each of the front six contracts. The first step regresses VIX futures returns onto variance swap forward returns to estimate the hedge ratio $\hat{\beta}_n$. The second step regresses the hedged return onto the deviation measure. The deviation measure significantly predicts hedged returns across contracts and sample periods, and is robust to the presence of other predictors like the VIX and realized variance (RV) which proxy for the variance risk premium. The variables in the second step regression are z-scored for ease of interpretation. Panel A (B) reports results for the full (post-crisis) sample.

Return Predictability Regression: $R^{VIX}_{t+h,n} - \hat{\beta}_n R^{VSF}_{t+h,n} = \gamma_n' x_{t,n} + \epsilon_{t+h,n}$

Panel A.I: Full sample from 2004 to 2018

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.19***	0.23***	0.27***	0.25***	0.36***	0.38***
	(0.03)	(0.05)	(0.05)	(0.03)	(0.07)	(0.08)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.038	0.051	0.075	0.063	0.129	0.145

Panel A.II: Full sample from 2004 to 2018 with controls

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.27***	0.31***	0.33***	0.32***	0.37***	0.39***
	(0.04)	(0.04)	(0.05)	(0.04)	(0.05)	(0.07)
VIX	0.29***	-0.00	0.01	-0.03	0.07	0.23**
	(0.11)	(0.15)	(0.12)	(0.13)	(0.11)	(0.11)
RV	-0.13	0.09	0.14	0.20	-0.16	-0.23**
	(0.11)	(0.19)	(0.16)	(0.16)	(0.12)	(0.09)
RMRF	-0.01	-0.09**	-0.02	-0.11***	-0.02	0.03
	(0.04)	(0.04)	(0.05)	(0.04)	(0.05)	(0.05)
VLM	-0.01	0.13***	0.16***	0.11***	0.17***	0.16***
	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)	(0.03)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.065	0.076	0.117	0.111	0.175	0.175

Panel B.I: Post-crisis sample from 2010 to 2018

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.23***	0.34***	0.28***	0.26***	0.40***	0.38***
	(0.04)	(0.03)	(0.04)	(0.04)	(0.05)	(0.05)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.052	0.112	0.078	0.068	0.158	0.143

Panel B.II: Post-crisis sample from 2010 to 2018 with controls

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.29***	0.39***	0.32***	0.30***	0.40***	0.37***
	(0.04)	(0.05)	(0.04)	(0.05)	(0.05)	(0.05)
VIX	0.26***	0.11	-0.00	-0.12	-0.00	0.07
	(0.09)	(0.10)	(0.11)	(0.10)	(0.10)	(0.10)
RV	-0.09	-0.04	0.06	0.10	-0.03	-0.04
	(0.08)	(0.08)	(0.09)	(0.08)	(0.09)	(0.09)
RMRF	0.04	0.00	-0.02	-0.03	0.02	0.07
	(0.05)	(0.04)	(0.05)	(0.05)	(0.05)	(0.05)
VLM	-0.02	0.15**	0.20***	0.09**	0.05	0.08**
	(0.04)	(0.06)	(0.04)	(0.04)	(0.04)	(0.04)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.079	0.136	0.118	0.077	0.160	0.150

Newey-West SEs with 15 lags in parentheses, * p<.10, ** p<.05, *** p<.01

Table 6: Deviation Measure Predicts the Returns of VIX Futures Hedged with Stock Market Returns

This table reports return predictability regressions for VIX futures returns that are hedged with stock market returns over a weekly horizon (h=5). The regressions are analogous to Table 5. When hedging with stock market returns, the deviation measure remains significant at predicting VIX futures returns, particularly during the post-crisis period. These results provide evidence that the deviation measure is identifying mispricing of VIX futures relative to variance swap forwards.

Return Predictability Regression: $R^{VIX}_{t+h,n} - \beta_n RMRF_{t+h,n} = \gamma_n' x_{t,n} + \epsilon_{t+h,n}$

Panel A.I: Full sample from 2004 to 2018						
Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.24***	0.28***	0.04	0.07**	0.17***	0.16***
	(0.07)	(0.07)	(0.04)	(0.03)	(0.06)	(0.06)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.055	0.079	0.001	0.005	0.029	0.026

Panel A.II: Full sample from 2004 to 2018 with controls						
Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.28***	0.28***	0.05	0.06*	0.13***	0.12***
	(0.06)	(0.06)	(0.04)	(0.04)	(0.04)	(0.04)
VIX	0.43***	0.23	0.18	0.13	0.06	0.07
	(0.12)	(0.17)	(0.17)	(0.16)	(0.16)	(0.16)
RV	-0.36***	-0.30	-0.36*	-0.30	-0.17	-0.18
	(0.12)	(0.19)	(0.20)	(0.19)	(0.17)	(0.16)
RMRF	0.03	-0.01	-0.00	0.01	-0.00	-0.02
	(0.07)	(0.07)	(0.07)	(0.07)	(0.08)	(0.07)
VLM	-0.02	0.09***	0.02	0.01	0.04	0.06*
	(0.04)	(0.03)	(0.03)	(0.04)	(0.04)	(0.03)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.075	0.104	0.041	0.037	0.043	0.044

Panel B.I: Post-crisis sample from 2010 to 2018						
Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.14**	0.23***	0.16***	0.16***	0.18***	0.15***
	(0.06)	(0.06)	(0.05)	(0.05)	(0.06)	(0.06)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.019	0.053	0.025	0.025	0.031	0.022

Panel B.II: Post-crisis sample from 2010 to 2018 with controls						
Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.22***	0.27***	0.14***	0.13***	0.15***	0.12**
	(0.05)	(0.06)	(0.04)	(0.05)	(0.05)	(0.05)
VIX	0.30***	0.15	0.06	0.02	-0.01	-0.01
	(0.10)	(0.12)	(0.12)	(0.12)	(0.11)	(0.11)
RV	-0.12	-0.08	-0.11	-0.09	-0.05	-0.06
	(0.10)	(0.11)	(0.12)	(0.12)	(0.11)	(0.11)
RMRF	0.01	-0.02	0.00	0.01	0.01	0.03
	(0.06)	(0.05)	(0.05)	(0.05)	(0.05)	(0.06)
VLM	-0.03	0.05	-0.05	-0.05	-0.03	0.01
	(0.05)	(0.05)	(0.04)	(0.05)	(0.05)	(0.04)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.054	0.063	0.029	0.032	0.033	0.027

Newey-West SEs with 15 lags in parentheses, * p<.10, ** p<.05, *** p<.01

Table 7: Deviation is Less Significant for Variance Swap Forward Returns Hedged with the Stock Market, Especially in Post-Crisis Period

This table reports return predictability regressions for variance swap forwards that are hedged with stock market returns over a weekly horizon (h=5). The regressions are analogous to Table 5. In comparison to VIX futures, the deviation measure is less significant at predicting hedged variance swap forward returns, particularly during the post-crisis period where there is lack of significance for some contracts. These results provide less evidence that the deviation measure is identifying mispricing of variance swap forwards relative to VIX futures. That said, the deviation measure does predict hedged variance swap forward returns for most contracts and it tends to have the expected, negative sign, indicating that the evidence is somewhat mixed.

Return Predictability Regression: $R^{VSF}_{t+h,n} - \beta_n RMRF_{t+h,n} = \gamma_n' x_{t,n} + \epsilon_{t+h,n}$

Panel A.I: Full sample from 2004 to 2018

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.05	0.13	-0.15***	-0.11***	-0.11**	-0.16***
	(0.09)	(0.09)	(0.05)	(0.03)	(0.06)	(0.05)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.002	0.016	0.022	0.012	0.012	0.025

Panel A.II: Full sample from 2004 to 2018 with controls

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	0.03	0.07	-0.17***	-0.16***	-0.15***	-0.21***
	(0.08)	(0.06)	(0.05)	(0.04)	(0.05)	(0.06)
VIX	0.34**	0.33	0.29	0.26	0.11	-0.01
	(0.15)	(0.23)	(0.20)	(0.20)	(0.17)	(0.15)
RV	-0.45**	-0.46	-0.57**	-0.55**	-0.15	-0.10
	(0.23)	(0.30)	(0.28)	(0.27)	(0.23)	(0.17)
RMRF	0.04	0.04	0.00	0.09	0.01	-0.06
	(0.06)	(0.08)	(0.06)	(0.08)	(0.10)	(0.08)
VLM	0.02	0.02	-0.08***	-0.05**	-0.09***	-0.06*
	(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.032	0.048	0.123	0.115	0.019	0.038

Panel B.I: Post-crisis sample from 2010 to 2018

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	-0.12***	0.00	-0.04	-0.02	-0.12**	-0.16***
	(0.04)	(0.06)	(0.05)	(0.06)	(0.05)	(0.05)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.014	-0.000	0.001	0.000	0.015	0.026

Panel B.II: Post-crisis sample from 2010 to 2018 with controls

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Deviation	-0.05	0.01	-0.08*	-0.07	-0.14***	-0.18***
	(0.05)	(0.06)	(0.04)	(0.06)	(0.05)	(0.05)
VIX	0.16*	0.11	0.11	0.14	0.03	-0.03
	(0.10)	(0.12)	(0.13)	(0.13)	(0.12)	(0.11)
RV	-0.09	-0.08	-0.18	-0.20*	-0.05	-0.07
	(0.10)	(0.11)	(0.12)	(0.12)	(0.10)	(0.10)
RMRF	-0.06	-0.03	0.01	0.03	-0.02	-0.04
	(0.05)	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
VLM	-0.00	-0.04	-0.18***	-0.11**	-0.06	-0.04
	(0.04)	(0.05)	(0.04)	(0.05)	(0.04)	(0.04)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.025	0.005	0.037	0.016	0.017	0.036

Newey-West SEs with 15 lags in parentheses, * p<.10, ** p<.05, *** p<.01

Table 8: Summary Statistics for Relative Value Trading Strategies

This table reports summary statistics for weekly returns from the relative value trading strategies based on the deviation measure in comparison to stock market and volatility factor returns. Panel A reports statistics for the relative value strategy using VIX futures returns hedged with variance swap forward returns (1), stock market returns (2), and stock market returns including transaction costs (3). Columns (4-6) report results for a post-crisis sample from 2010 to 2018. Panel B reports analogous statistics for stock market returns (RMRF), receiving fixed in one-month variance swap forwards (VS1), and selling the front month VIX futures contract (VX1).

Weekly Return Summary Statistics

Panel A: VIX futures trading strategy levered for 10% annualized volatility

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Hedge	VSF	RMRF	RMRF	VSF	RMRF	RMRF
Transaction Costs	No	No	Yes	No	No	Yes
Post-Crisis	No	No	No	Yes	Yes	Yes
Mean	0.58	0.27	0.14	0.54	0.21	0.16
Standard Deviation	1.39	1.39	1.39	1.39	1.39	1.39
Sharpe Ratio	0.42	0.20	0.10	0.39	0.15	0.12
t-statistic	14.54	6.29	3.26	10.69	3.81	2.98
Skewness	0.71	1.41	2.09	1.25	1.61	2.03
Kurtosis	18.44	13.07	20.62	26.12	17.32	23.82
Minimum	-11.58	-7.04	-8.50	-9.64	-7.27	-9.67
25th-Percentile	-0.06	-0.39	-0.27	-0.08	-0.40	-0.23
Median	0.51	0.16	0.00	0.51	0.13	0.00
75th-Percentile	1.15	0.84	0.50	1.12	0.75	0.54
Maximum	16.48	13.91	15.17	18.02	14.36	17.27
Negative Percent	26.66	39.99	32.64	27.29	41.94	31.30
Sortino Ratio	0.87	0.42	0.20	0.79	0.30	0.23
Maximum Drawdown	6.72	9.67	16.68	6.30	9.96	8.17
Leverage (percentage)	11.62	7.87	5.78	12.71	8.12	6.57
CAPM α-to-margin	4.90	3.61	2.64	4.12	2.74	2.75
Observations	3698	3698	3698	2246	2246	2246

Panel B: Stock market, variance swap, and VIX futures factor returns

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Return	RMRF	VS1	VX1	RMRF	VS1	VX1
Post-Crisis	No	No	No	Yes	Yes	Yes
Mean	0.17	0.03	0.19	0.24	0.03	0.23
Standard Deviation	2.37	0.21	2.20	2.08	0.14	2.14
Sharpe Ratio	0.07	0.15	0.09	0.11	0.20	0.11
t-statistic	2.21	4.19	2.54	2.90	5.40	2.68
Skewness	-0.92	-0.97	-1.71	-0.80	-1.98	-1.29
Kurtosis	11.74	68.08	15.17	6.57	30.99	12.50
Minimum	-19.81	-3.03	-21.04	-14.47	-1.43	-19.70
25th-Percentile	-0.86	-0.00	-0.40	-0.71	-0.00	-0.45
Median	0.32	0.03	0.21	0.38	0.03	0.25
75th-Percentile	1.44	0.07	1.15	1.45	0.06	1.25
Maximum	18.26	2.84	13.02	9.00	0.98	10.35
Negative Percent	41.24	26.04	33.59	39.31	26.40	33.53
Sortino Ratio	0.11	0.21	0.12	0.18	0.29	0.16
Maximum Drawdown	50.76	5.74	45.03	16.38	1.18	16.69
Observations	3698	3698	3698	2246	2246	2246

Table 9: Alpha-to-Margin Estimates for the Relative Value Trading Strategies

This table reports weekly factor return regressions for the relative value trading strategies using the maximum leverage to required initial margin for VIX futures. The alpha-to-margin estimates are large and significant across hedging portfolios, factor model specifications, and sample periods.

Factor Model Specification	CAPM (1)	FFC4 (2)	FFCV6 (3)	CAPM (4)	FFC4 (5)	FFCV6 (6)
Panel A: Variance Swap Hedge						
Alpha	4.90***	4.92***	4.83***	4.12***	4.16***	4.02***
	(0.35)	(0.35)	(0.39)	(0.43)	(0.44)	(0.59)
RMRF	0.37	0.28	0.01	0.44	0.32	-0.09
	(0.32)	(0.38)	(0.35)	(0.42)	(0.47)	(0.43)
HML		-0.06	-0.04		0.02	-0.04
		(0.26)	(0.25)		(0.39)	(0.38)
SMB		0.32	0.36		0.55	0.60
		(0.36)	(0.35)		(0.47)	(0.47)
MOM		-0.14	-0.11		-0.16	-0.22
		(0.19)	(0.19)		(0.30)	(0.29)
VS1			4.06			9.42
			(4.52)			(16.40)
VX1			0.04			-0.07
			(0.44)			(0.82)
Adjusted R^2	0.005	0.006	0.008	0.006	0.009	0.013
Panel B: Stock Market Hedge						
Alpha	3.61***	3.65***	4.00***	2.74***	2.85***	3.23***
	(0.54)	(0.55)	(0.55)	(0.72)	(0.72)	(0.72)
RMRF	-0.74**	-0.92**	1.21**	-0.78	-0.89	2.12***
	(0.31)	(0.39)	(0.50)	(0.56)	(0.62)	(0.61)
HML		-0.45	-0.50		-1.60**	-1.28***
		(0.46)	(0.40)		(0.63)	(0.49)
SMB		0.60	0.31		0.91	0.64
		(0.56)	(0.51)		(0.78)	(0.66)
MOM		-0.55*	-0.54*		-1.59***	-1.38***
		(0.33)	(0.30)		(0.54)	(0.46)
VS1			-7.93*			-14.90
			(4.58)			(9.92)
VX1			-2.38***			-2.96***
			(0.71)			(0.93)
Adjusted R^2	0.010	0.014	0.069	0.009	0.034	0.119
Panel C: Stock Market Hedge with Transaction Costs						
Alpha	2.64***	2.72***	3.24***	2.75***	2.93***	3.55***
	(0.74)	(0.75)	(0.73)	(0.87)	(0.87)	(0.81)
RMRF	-0.98**	-1.40***	1.98***	-1.20*	-1.38*	2.97***
	(0.39)	(0.52)	(0.73)	(0.68)	(0.76)	(0.81)
HML		-0.28	-0.35		-1.93**	-1.45**
		(0.70)	(0.59)		(0.79)	(0.60)
SMB		1.48**	1.02		1.35	0.96
		(0.74)	(0.64)		(0.96)	(0.80)
MOM		-0.79*	-0.76**		-2.24***	-1.92***
		(0.42)	(0.37)		(0.62)	(0.50)
VS1			-10.48			-27.37**
			(7.88)			(11.13)
VX1			-3.95***			-3.91***
			(1.13)			(1.40)
Adjusted R^2	0.009	0.016	0.093	0.013	0.045	0.155
Post-Crisis	No	No	No	Yes	Yes	Yes
Observations	3698	3698	3698	2246	2246	2246

Newey-West SEs with 15 lags in parentheses, * p<.10, ** p<.05, *** p<.01

Table 10: The No-Arbitrage Deviation Measure is Decreasing in Risk and Increasing in Demand Pressure for VIX Futures

This table reports a panel regression of the deviation measure onto changes in different risk and demand variables over a one-week horizon $h = 5$. Specifications 1-3 control for the lagged deviation. Specifications 4-6 add controls variables and fixed effects. The control variables include the contract-specific time-to-maturity, open interest, and initial margin and the lagged VIX index to proxy for the economic environment. The fixed effects include the contract, contract number, and calendar year. The explanatory variables are z-scored for ease of interpretation. Across the different variables proxying for risk and demand and regression specifications, the results indicate that the deviation measure is decreasing in risk and increasing in demand.

$$\text{Deviation}_{t+h,n} = \beta\Delta x_{t+h} + \rho\text{Deviation}_{t,n} + \delta\text{Controls}_t + \text{FEs} + \epsilon_{t+h,n}$$

Panel A: Risk Factors from 2007 to 2018

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Explanatory Variable	RMRF	RV	VIX	RMRF	RV	VIX
Δx_{t+h}	0.10***	-0.05***	-0.14***	0.11***	-0.25***	-0.15***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.05)	(0.02)
$\text{Deviation}_{t,n}$	0.79***	0.79***	0.81***	0.60***	0.61***	0.60***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Observations	17689	17689	17689	17689	17689	17689
Adjusted R^2	0.655	0.646	0.667	0.694	0.695	0.707
Controls and FEs	No	No	No	Yes	Yes	Yes

Panel B: Risk Factors from 2010 to 2018

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Explanatory Variable	RMRF	RV	VIX	RMRF	RV	VIX
Δx_{t+h}	0.17***	-0.06**	-0.17***	0.17***	-0.33***	-0.18***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.05)	(0.02)
$\text{Deviation}_{t,n}$	0.85***	0.84***	0.86***	0.66***	0.64***	0.65***
	(0.02)	(0.02)	(0.01)	(0.03)	(0.03)	(0.03)
Observations	13260	13260	13260	13260	13260	13260
Adjusted R^2	0.749	0.724	0.756	0.789	0.782	0.794
Controls and FEs	No	No	No	Yes	Yes	Yes

Panel C: Demand Factors from 2010 to 2018

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Explanatory Variable	Dealer	ETP	Option	Dealer	ETP	Option
Δx_{t+h}	0.04***	0.02	0.05***	0.05***	0.03**	0.07***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
$\text{Deviation}_{t,n}$	0.85***	0.85***	0.85***	0.67***	0.67***	0.67***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Observations	13260	13260	13260	13260	13260	13260
Adjusted R^2	0.724	0.723	0.725	0.765	0.764	0.767
Controls and FEs	No	No	No	Yes	Yes	Yes

SEs double-clustered by date and contract, * p<.10, ** p<.05, *** p<.01

Figure 1: Equity Volatility Market Size and Dealer Position

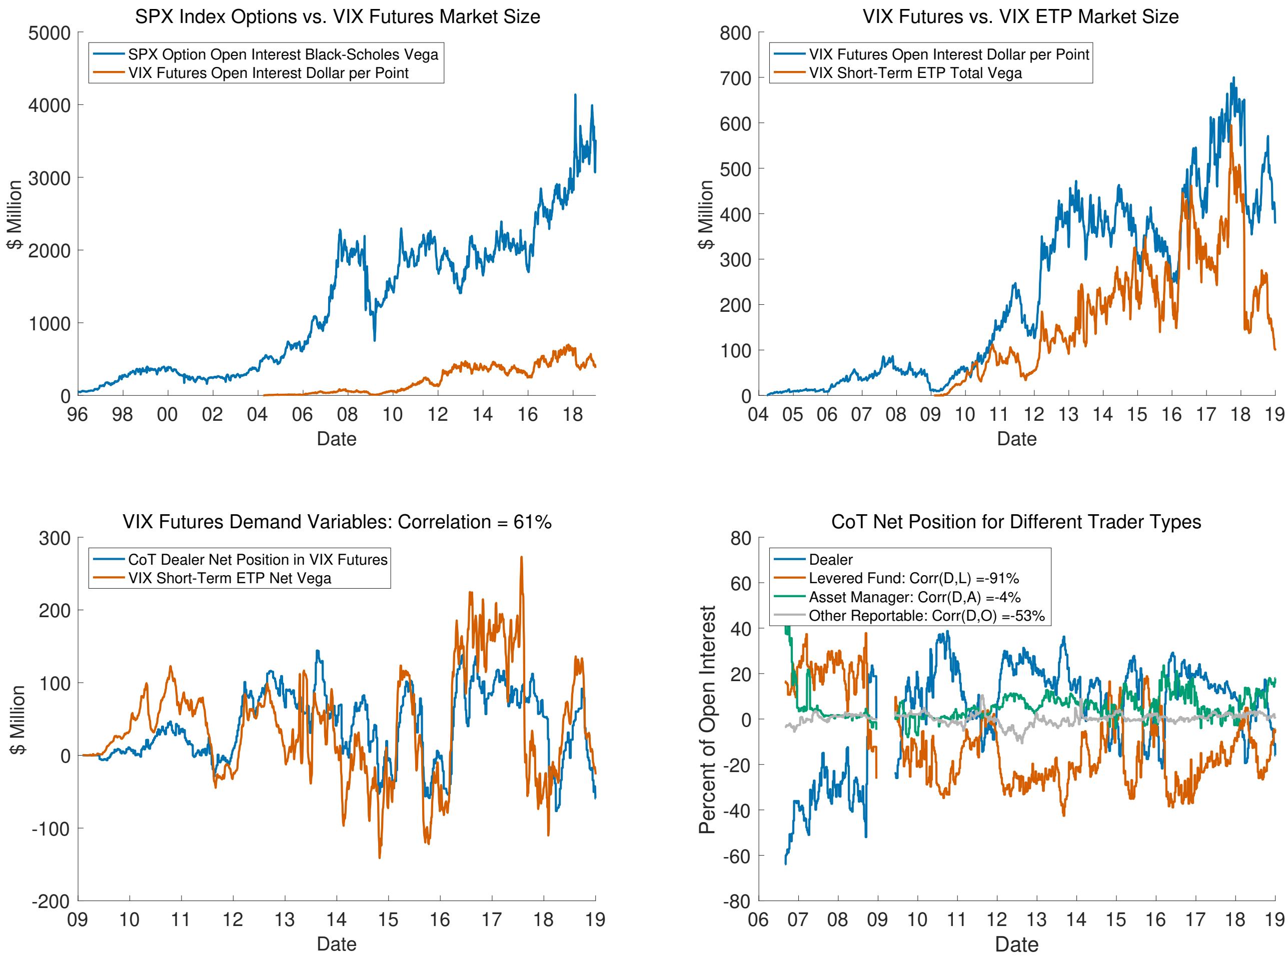

This figure provides an overview of equity volatility market size and investor positioning. The S&P 500 index options and VIX futures markets have experienced substantial growth over the last decade. The index options market has been much larger than the VIX futures market throughout the sample (top left). The growth in the VIX futures market has coincided with growth in the VIX ETP market starting in 2009 (top right). Focusing on the behavior of certain traders, dealer positions in VIX futures are highly related to demand for VIX ETPs (bottom left). Within the VIX futures market, dealers and leveraged funds generally take large and opposing positions (bottom right). The break in the time-series of CoT positions corresponds to a period in early 2009 when VIX futures open interest was low and the breakdown was not reported. One interpretation of the position data is that retail demand for ETPs is hedged by dealers against leveraged funds.

Figure 2: Variance Swap Forward Rate Estimation

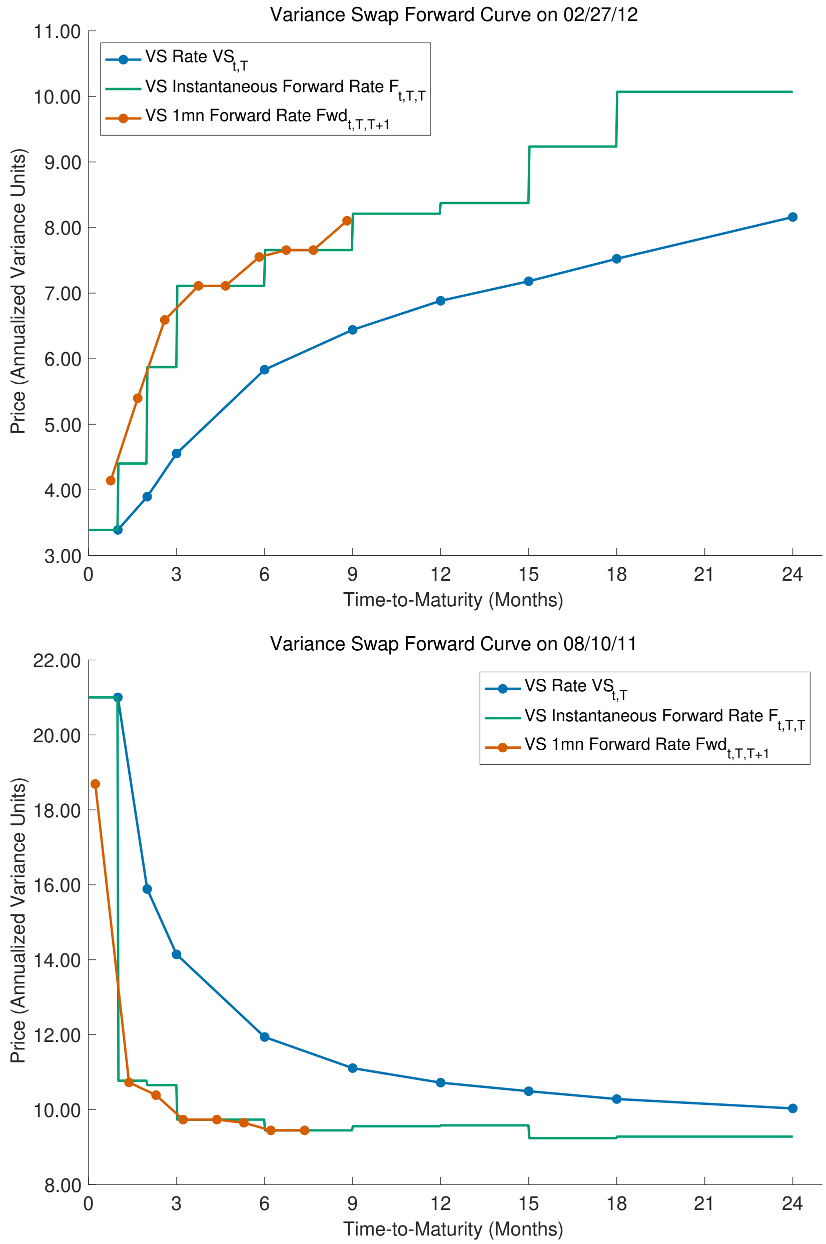

This figure illustrates the computation of one-month variance swap forward rates for the deviation measure. The variance swap curve in blue is decomposed into an instantaneous forward curve in green assuming flat forward rates between the observed maturities. The one-month forward rates at VIX futures maturity dates are an average of the green variance swap forward curve over the next month. Variance swap rates and forward rates in these plots are expressed in annualized variance units.

Figure 3: Examples of Law of One Price Violations

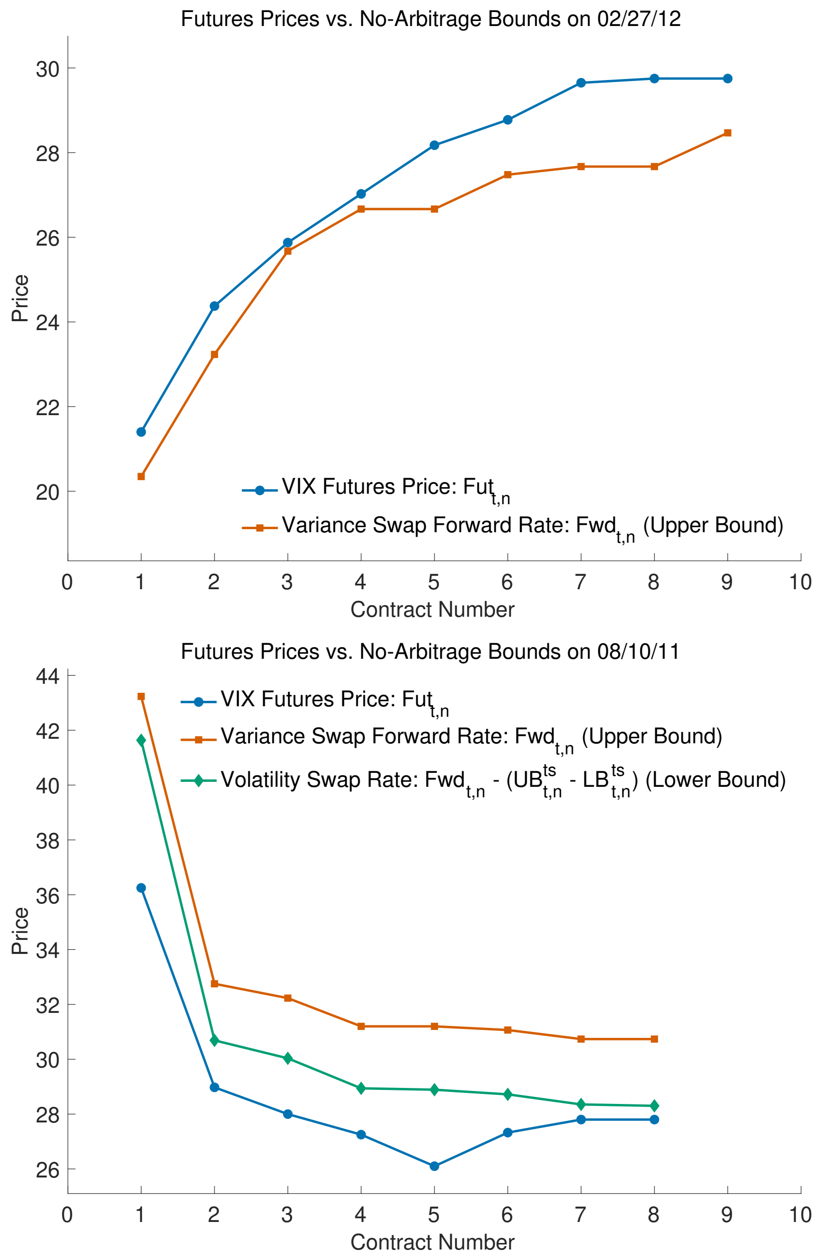

The deviation measure is the difference between the VIX futures price in blue and one-month variance swap forward rate in red both expressed in annualized volatility units. The top plot illustrates examples of static arbitrage opportunities in which the prices of VIX futures are above the upper bound. An arbitrageur could lock in a riskless profit with no capital outlay by selling the expensive VIX futures contracts with the hedge ratio from the paper, paying fixed in the variance swap forwards, and holding the position until maturity. The bottom plot illustrates an example when the prices of VIX futures are significantly below the upper bound and an estimate of the lower bound. In this case, an arbitrageur could lock in a profit by buying VIX futures and receiving fixed in volatility swap forwards. The lower bound is the upper bound in red minus the difference between an estimate of the upper bound and lower bound from a term-structure model.

Figure 4: Deviation of VIX Futures from the Law of One Price

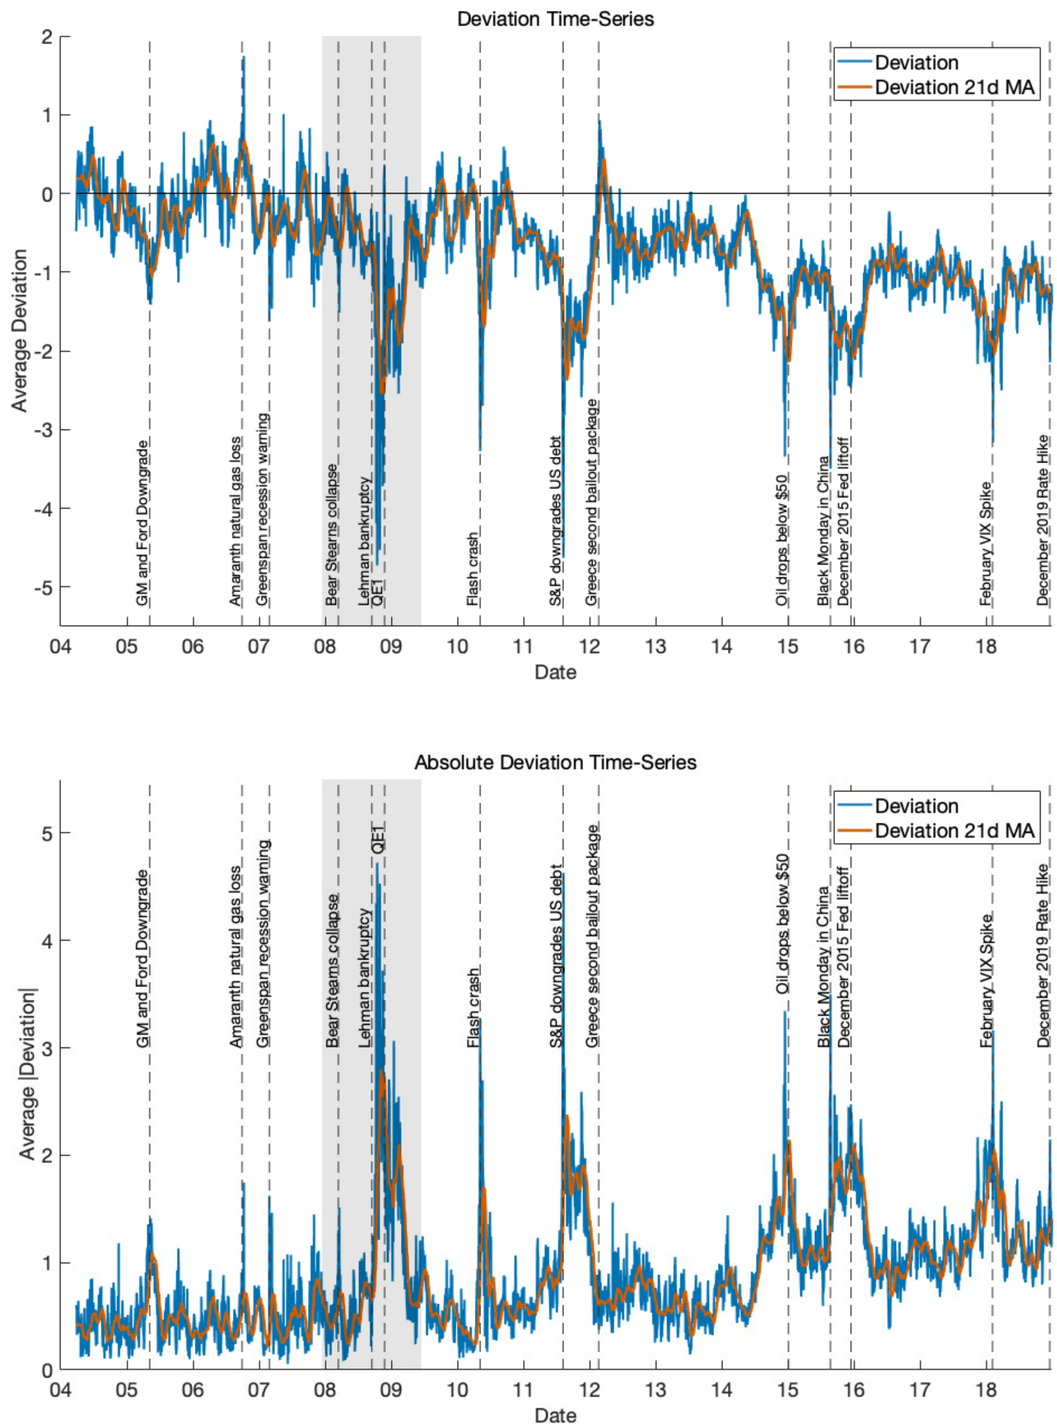

This figure plots the average deviation and average absolute deviation for the front six contracts from March 2004 to December 2018. The deviation measure for the n-th contract is the difference between the VIX futures price and its no-arbitrage upper bound, $\text{Deviation}_{t,n} = \text{Fut}_{t,n} - \text{Fwd}_{t,n}$. Positive values are law of one price violations. Gray shading indicates NBER recessions.

Figure 5: Deviation of the Front Contract from the Law of One Price

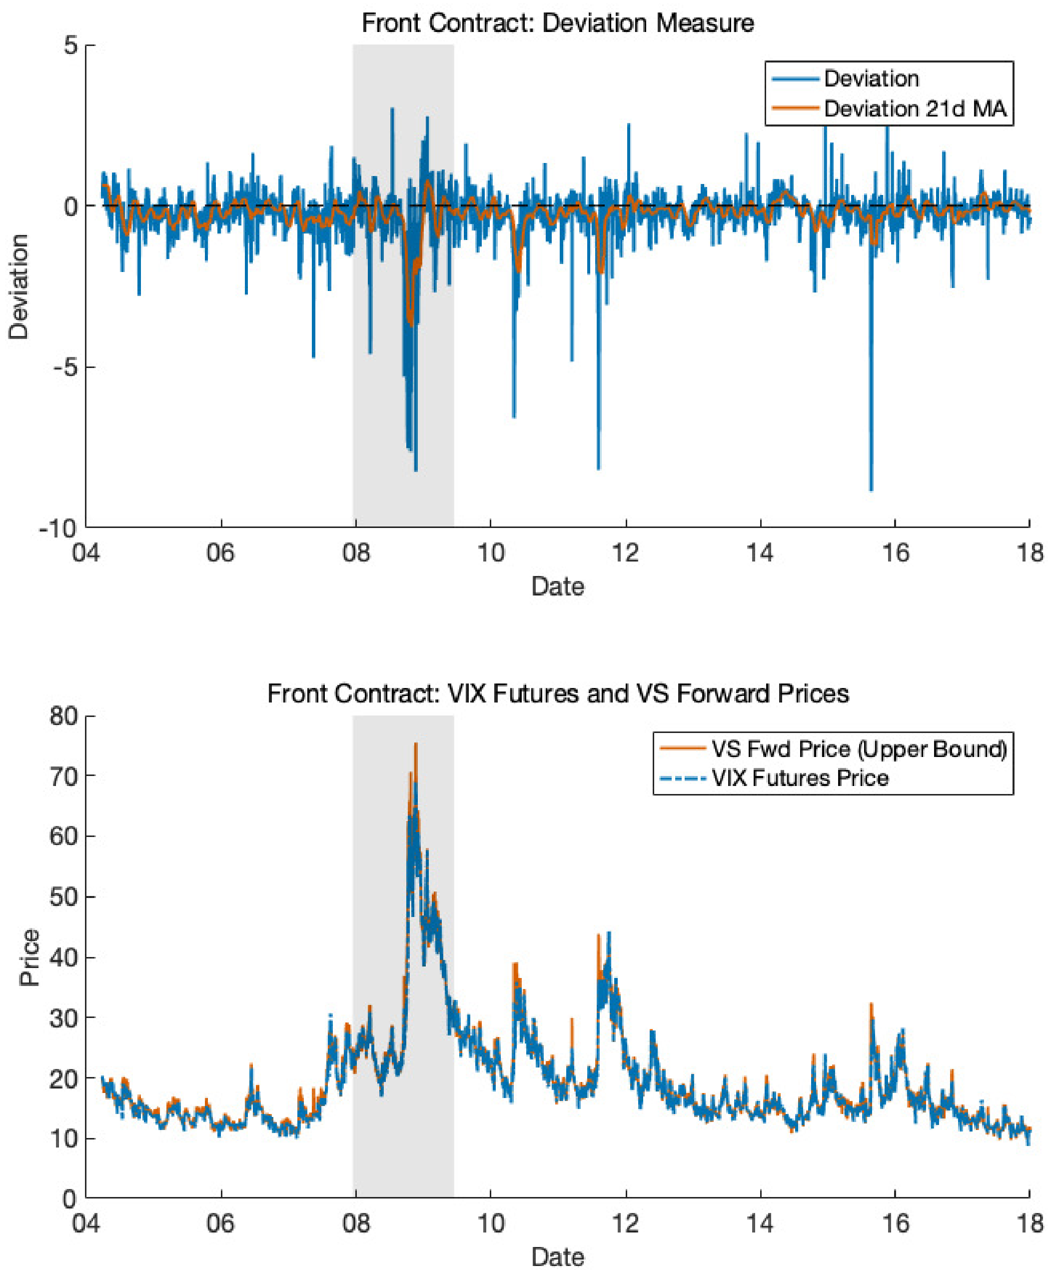

This figure plots the deviation measure for the front contract alongside the VIX futures and variance swap forward prices that are used to compute the deviation measure. While the futures price and forward price tend to track each other closely, there are periods with prolonged and significant law of one price deviations during the sample period.

Figure 6: Performance of Relative Value Strategy Based on Deviation Measure

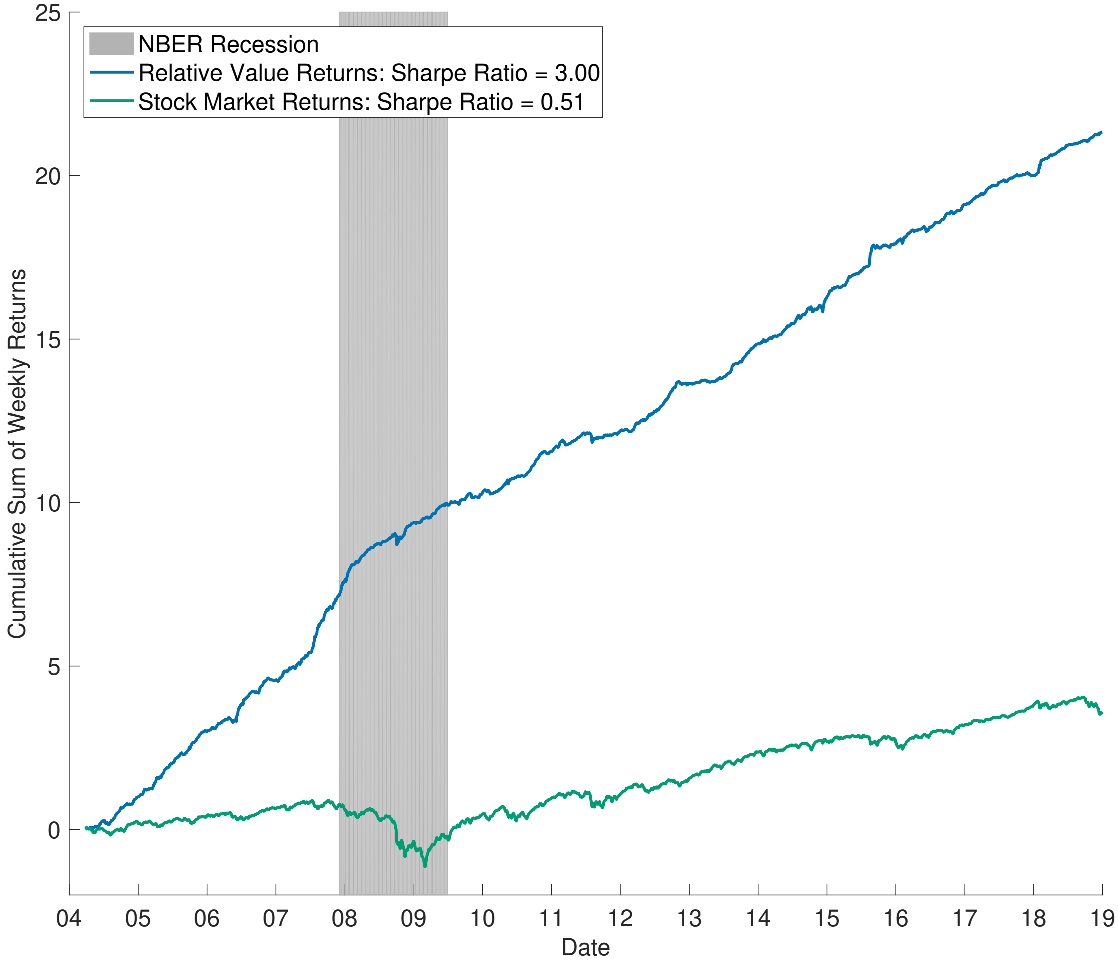

This figure plots the performance of the relative value trading strategy in VIX futures and variance swap forwards against stock market returns. The relative value strategy goes long (short) VIX futures when the deviation measure exceeds a low (high) threshold of $\tau = .50$ z-scores for the front six contracts and hedges with variance swap forwards. Each trade is held for a one-week horizon, with the strategy forming an equally-weighted portfolio when multiple contracts are traded on the same day. The plot reports the cumulative sum of weekly returns for the strategy and stock market which are normalized to 10% annualized volatility for comparison. The relative value strategy earns a large Sharpe ratio and exhibits low drawdowns compared to the stock market.

Figure 7: Robustness of Relative Value Trading Strategy

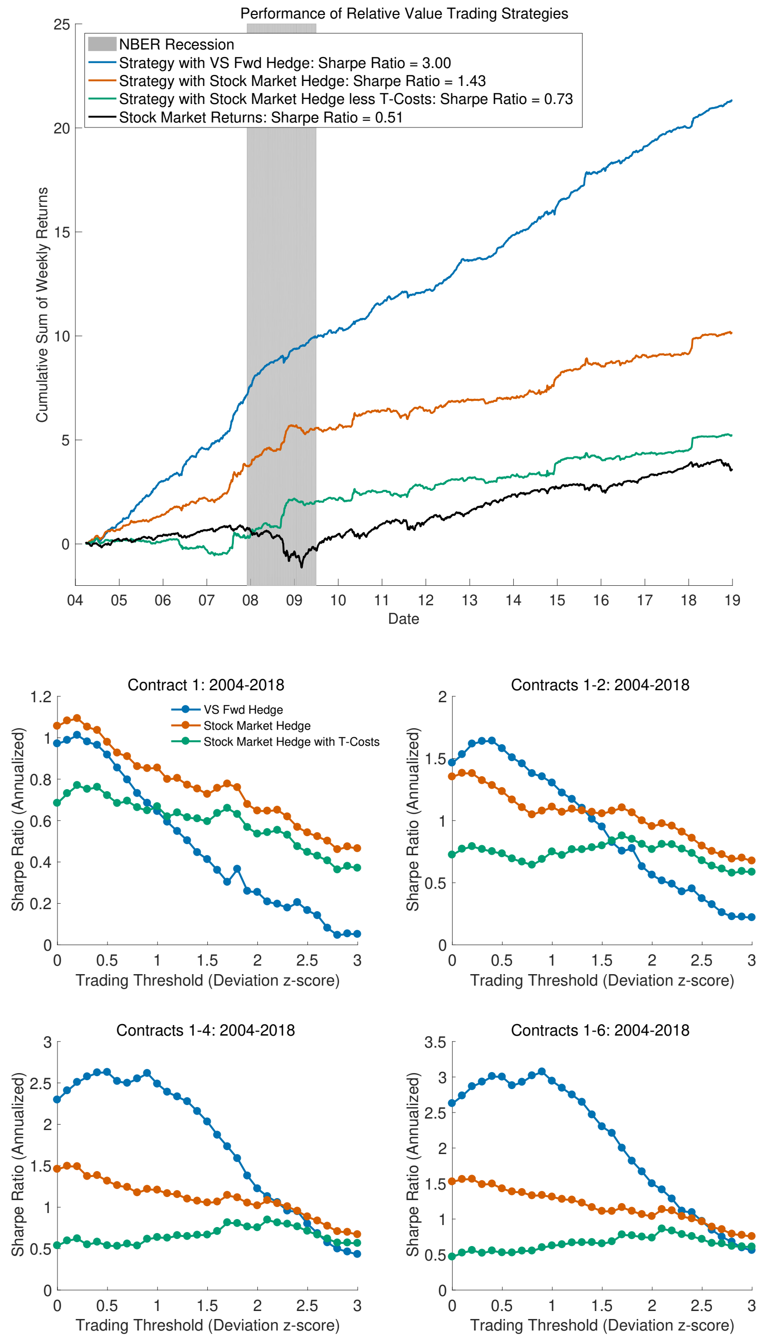

This figure illustrates the robustness of the relative value trading strategy to hedging with stock market returns, including transaction costs, trading different numbers of contracts, and varying the deviation-based trading threshold. The top plot reports the time-series of cumulative returns for the baseline strategies discussed in the paper in comparison to the stock market. The plot normalizes the annualized volatility of each series to 10% for comparison. The bottom plot reports the Sharpe ratios for the baseline strategies varying the number of contracts traded and threshold for trading. Across specifications, the relative value trading strategy earns a large Sharpe ratio and produces returns that are largely uncorrelated with traditional risk factors.

Figure 8: Time-Series of Deviation, VIX, and Dealer Net Position for VAR

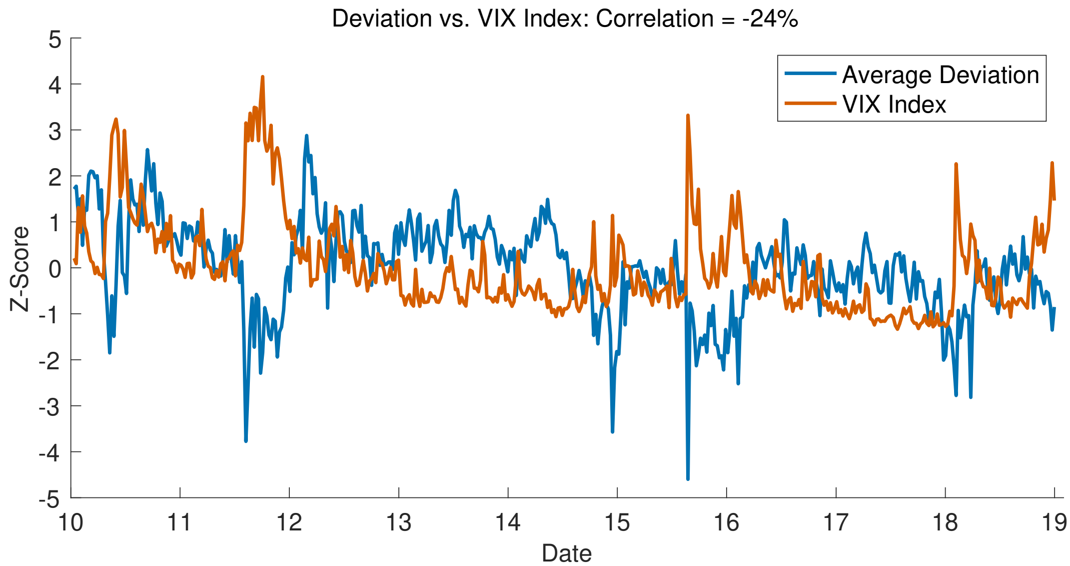

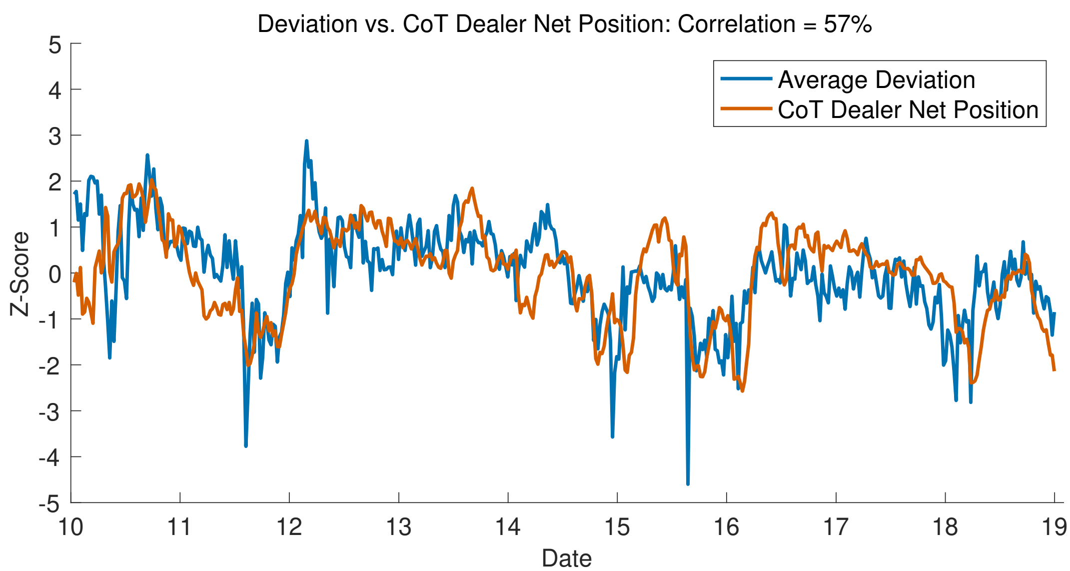

This figure plots the average deviation measure against the VIX index and dealer net position at a weekly frequency from 2010 to 2018 ($T = 469$). The VAR is estimated using these three variables, y_t =[Deviation$_t$ VIX$_t$ Dealer Net Position$_t$], which are standardized for comparison and ease of interpretation. The deviation measure is negatively correlated with the VIX index which serves as a proxy for risk (top plot). The deviation measure is positively correlated with the dealer net position which serves as a proxy for demand to buy VIX futures (bottom plot).

Figure 9: Impulse Response Functions of the No-Arbitrage Deviation Measure and Dealer Net Position to Shocks in the VIX index and Dealer Net Position

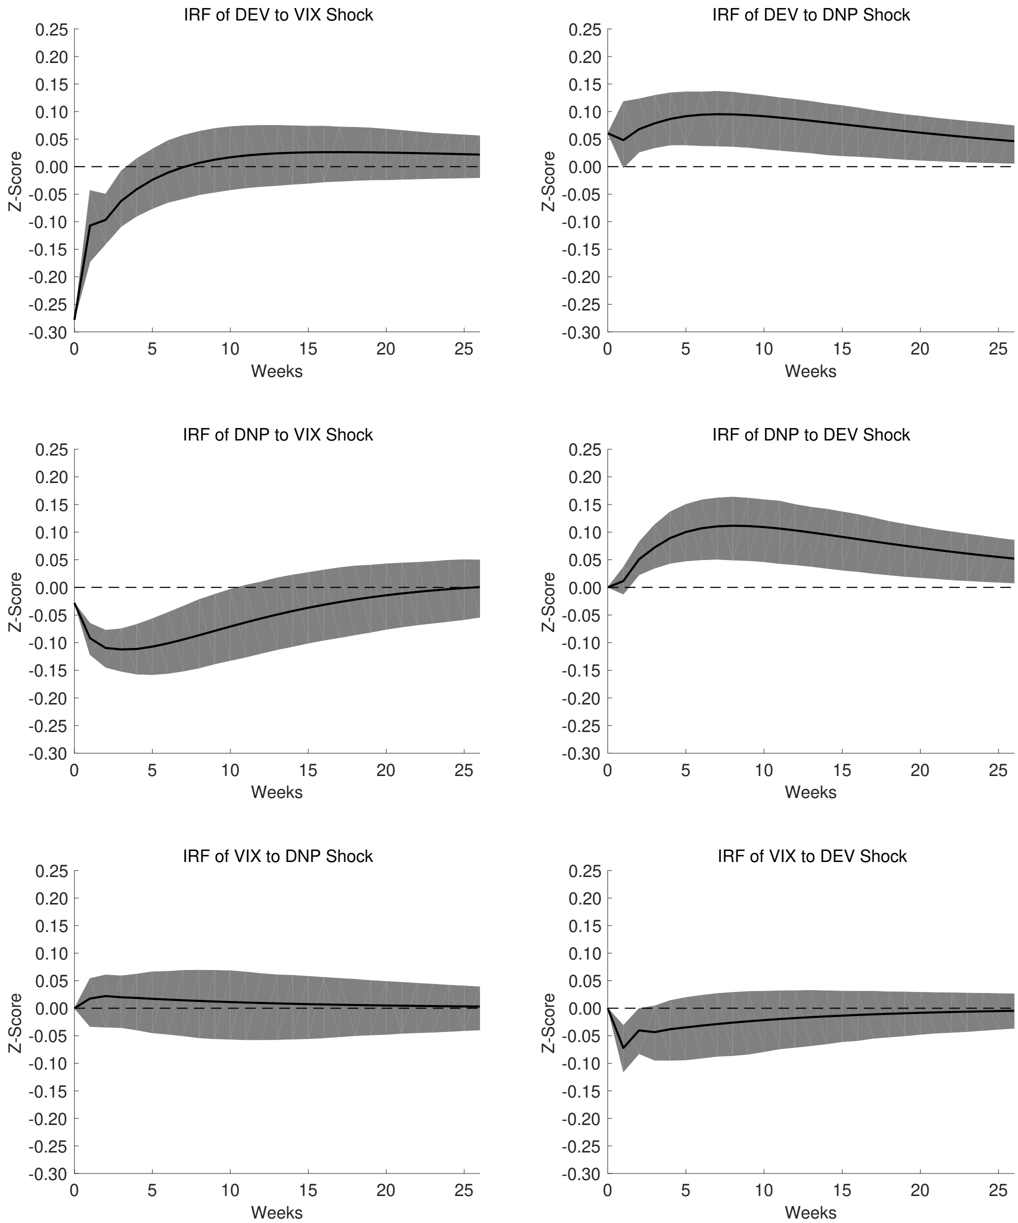

This figure reports the impulse response functions (IRFs) from the VAR. The IRFs are identified from a Cholesky decomposition with the ordering: VIX, DNP, DEV. The VAR is estimated using weekly data from 2010 to 2018 ($T = 469$). The IRFs in the first row show that DEV is decreasing (increasing) in VIX (DNP) shocks. The IRFs in the second row show that the DNP is decreasing (increasing) in VIX (DEV) shocks. The IRFs in the third row show that the VIX exhibits little response to DEV or DNP shocks, except for a short term-response to DEV shocks that becomes insignificant after a few weeks. The 95% pointwise confidence intervals in gray are block bootstrapped.

Figure 10: Event Study of Deviation Measure Around VIX Futures Margin Changes, FOMC days, Nonfarm Payrolls, Month-End, and Quarter-End Dates

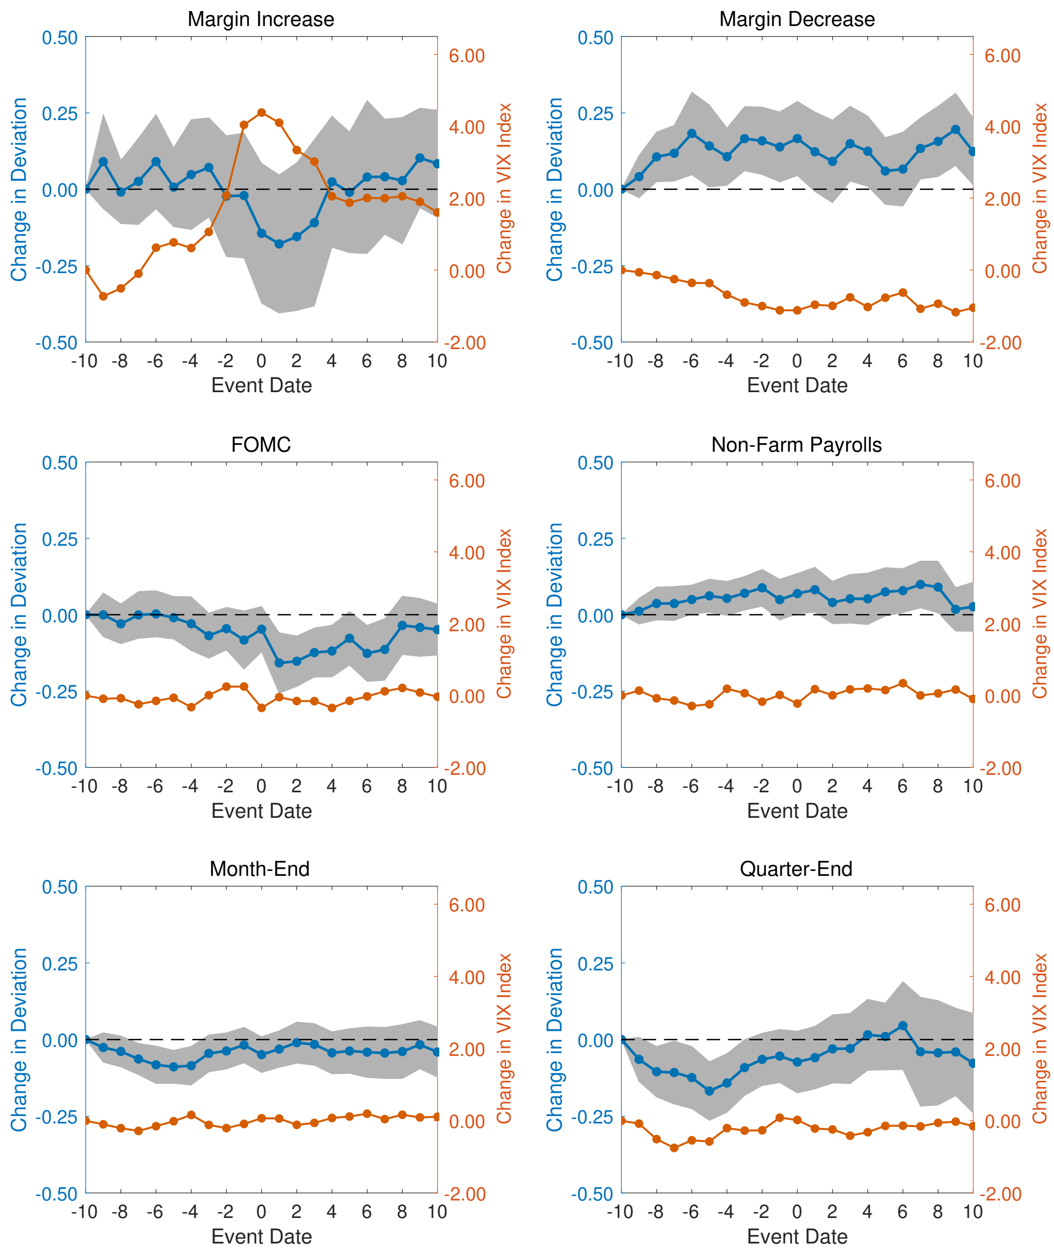

This figure plots the event time reaction of the average deviation measure and VIX index to changes in the initial margin for the front month VIX futures contract, FOMC announcements, non-farm payrolls announcements, end-of-month dates, and end-of-quarter dates. The figure reports the change in the VIX index in red and the change in the average deviation measure across the front six contracts in blue with a 95% pointwise confidence interval in gray for the ten days before and after the event.

Appendix:

"The Law of One Price in Equity Volatility Markets"

Peter Van Tassel

December 2019

A Appendix

Table A.1: Correlation of Alternative Deviation Measures

This table presents the correlation of the baseline deviation measure with alternative specifications. Panel A considers different ways of removing the bias from the upper bound, using a regression-based model (A.I) or term-structure model (A.II). Panel B considers alternative data sources. B.I uses VIX settlement prices that are not synchronized with SPX option quotes during the later years in the sample. B.II through B.V use alternative data sources for computing variance swap forward rates. The columns report the correlation for different contracts and for the average deviation across contracts. Overall, the alternative measures are highly correlated with the baseline measure.

Correlation of Baseline Deviation Measure with Alternative Measures

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)	Avg.
Panel A: Convexity Adjustments							
Panel A.I: Regression-based convexity adjustment (Mar04-Dec18)							
Correlation in levels	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Correlation in monthly changes	1.00	1.00	1.00	1.00	1.00	0.99	1.00
Correlation in weekly changes	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Panel A.II: Term-structure model convexity adjustment (Mar04-Dec18)							
Correlation in levels	0.99	0.97	0.97	0.97	0.97	0.97	0.96
Correlation in monthly changes	0.99	0.99	0.99	0.99	0.99	0.99	0.99
Correlation in weekly changes	0.99	0.99	0.99	0.99	0.99	1.00	0.98
Panel B: Alternative Data Sources							
Panel B.I: VIX settlement prices, not synchronized (Mar04-Dec18)							
Correlation in levels	0.96	0.96	0.98	0.99	0.99	0.99	0.98
Correlation in monthly changes	0.95	0.94	0.96	0.97	0.97	0.98	0.95
Correlation in weekly changes	0.94	0.90	0.93	0.94	0.95	0.97	0.89
Panel B.II: Bloomberg data for VS forward (Nov08-Dec18)							
Correlation in levels	0.66	0.69	0.71	0.77	0.85	0.89	0.87
Correlation in monthly changes	0.72	0.63	0.64	0.72	0.70	0.78	0.81
Correlation in weekly changes	0.69	0.48	0.66	0.66	0.67	0.73	0.72
Panel B.III: CBOE VIX, VIX3M, VIX6M for VS forward (Nov08-Dec18)							
Correlation in levels	0.68	0.61	0.85	0.85	0.88	0.86	0.93
Correlation in monthly changes	0.74	0.43	0.80	0.80	0.81	0.74	0.89
Correlation in weekly changes	0.71	0.15	0.80	0.69	0.70	0.58	0.80
Panel B.IV: Hedge Fund OTC quotes for VS forward (Mar04-Nov13)							
Correlation in levels	0.41	0.53	0.68	0.82	0.83	0.77	0.71
Correlation in monthly changes	0.46	0.41	0.51	0.66	0.65	0.66	0.53
Correlation in weekly changes	0.26	0.27	0.40	0.51	0.51	0.62	0.27
Panel B.V: Markit OTC quotes for VS forward (monthly Sep06-Dec15)							
Correlation in Levels	0.64	0.74	0.80	0.85	0.75	0.79	0.87
Correlation in monthly changes	0.67	0.72	0.74	0.82	0.81	0.82	0.88

Table A.2: Hedge Ratios for Return Predictability Regressions

This table reports the hedge ratios and explanatory power from the first stage in the return predictability regressions for Tables 5, 6, and 7. The hedge ratios are significant across contracts and sample periods. Compared to the stock market, variance swap forwards provide more explanatory power for VIX futures returns, particularly for the longer-dated contracts.

Hedge Ratios for Weekly Return Predictability Regressions

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: VIX Futures onto Variance Swap Forwards (Table 5)						
Panel A.I: Sample from 2004 to 2018						
β_n	10.36***	14.23***	14.48***	14.31***	14.98***	14.67***
	(1.50)	(1.16)	(1.14)	(1.14)	(0.78)	(0.87)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.574	0.770	0.757	0.757	0.715	0.664
Panel A.II: Sample from 2010 to 2018						
β_n	15.28***	21.04***	19.65***	18.53***	18.09***	16.64***
	(1.69)	(0.83)	(0.59)	(0.54)	(0.62)	(0.51)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.627	0.845	0.830	0.824	0.793	0.760
Panel B: VIX Futures onto Stock Market (Table 6)						
Panel B.I: Sample from 2004 to 2018						
β_n	0.73***	0.58***	0.45***	0.38***	0.34***	0.30***
	(0.04)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.574	0.566	0.551	0.535	0.519	0.496
Panel B.II: Sample from 2010 to 2018						
β_n	0.84***	0.70***	0.56***	0.47***	0.41***	0.37***
	(0.05)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.592	0.676	0.669	0.651	0.620	0.606
Panel C: Variance Swap Forwards onto Stock Market (Table 7)						
Panel C.I: Sample from 2004 to 2018						
β_n	0.05***	0.03***	0.02***	0.02***	0.02***	0.02***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.546	0.451	0.388	0.411	0.444	0.403
Panel C.II: Sample from 2010 to 2018						
β_n	0.05***	0.03***	0.02***	0.02***	0.02***	0.02***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.642	0.650	0.609	0.600	0.600	0.568

Newey-West SEs with $3 \cdot h$ lags in parentheses, * p<.10, ** p<.05, *** p<.01

Table A.3: Predicting VIX Futures Hedged Returns with LOOP Deviations: Robustness to Sample Period and Forecast Horizon

This table reports the predictability of the deviation measure for VIX futures returns hedged with variance swap forward returns over one-day, one-week, and one-month horizons across sample periods.

Return Predictability Regressions: $R^{VIX}_{t+h,n} - \beta_n R^{VS}_{t+h,n} = \gamma_n' x_{t,n} + \epsilon_{t+h,n}$

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A.I: Sample period 2004-2018, daily returns $h = 1$						
Deviation	0.21***	0.17***	0.21***	0.18***	0.27***	0.27***
	(0.04)	(0.03)	(0.04)	(0.02)	(0.04)	(0.05)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.044	0.030	0.042	0.031	0.073	0.072
Panel A.II: Sample period 2004-2018, weekly returns $h = 5$						
Deviation	0.19***	0.23***	0.27***	0.25***	0.36***	0.38***
	(0.03)	(0.05)	(0.05)	(0.03)	(0.07)	(0.08)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.038	0.051	0.075	0.063	0.129	0.145
Panel A.III: Sample period 2004-2018, monthly returns $h = 21$						
Deviation	0.26***	0.24***	0.29***	0.29***	0.41***	0.45***
	(0.09)	(0.06)	(0.05)	(0.05)	(0.08)	(0.07)
Observations	3344	3685	3685	3631	3194	3187
Adjusted R^2	0.068	0.059	0.084	0.084	0.167	0.198
Panel B.I: Sample period 2010-2018, daily returns $h = 1$						
Deviation	0.22***	0.19***	0.17***	0.17***	0.30***	0.26***
	(0.06)	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.046	0.037	0.030	0.027	0.092	0.068
Panel B.II: Sample period 2010-2018, weekly returns $h = 5$						
Deviation	0.23***	0.34***	0.28***	0.26***	0.40***	0.38***
	(0.04)	(0.03)	(0.04)	(0.04)	(0.05)	(0.05)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.052	0.112	0.078	0.068	0.158	0.143
Panel B.III: Sample period 2010-2018, monthly returns $h = 21$						
Deviation	0.14***	0.42***	0.44***	0.36***	0.51***	0.51***
	(0.05)	(0.04)	(0.07)	(0.07)	(0.06)	(0.06)
Observations	2019	2233	2233	2233	2233	2233
Adjusted R^2	0.019	0.174	0.196	0.133	0.255	0.256

Newey-West SEs with $3 \cdot h$ lags in parentheses, * p<.10, ** p<.05, *** p<.01

Table A.4: Predicting VIX Futures Hedged Returns with LOOP Deviations: Robustness to Deviation Measure

This table reports the predictability of the deviation measure for VIX futures returns hedged with variance swap forward returns over a weekly horizon (h=5) using different deviation measures. The sample period varies depending on when the different data sources are available.

Return Predictability Regressions: $R^{VIX}_{t+h,n} - \beta_n R^{VS}_{t+h,n} = \gamma_n' x_{t,n} + \epsilon_{t+h,n}$

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Deviation from paper ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation	0.19***	0.23***	0.27***	0.25***	0.36***	0.38***
	(0.03)	(0.05)	(0.05)	(0.03)	(0.07)	(0.08)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.038	0.051	0.075	0.063	0.129	0.145
Panel B: Deviation 1-day lag ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation 1d Lag	0.16***	0.13***	0.16***	0.12***	0.20***	0.20***
	(0.06)	(0.04)	(0.03)	(0.02)	(0.04)	(0.05)
Observations	3355	3682	3695	3636	3205	3037
Adjusted R^2	0.011	0.022	0.049	0.041	0.105	0.108
Panel C: Deviation 5-day moving average ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation 5d MA	0.17**	0.14***	0.16***	0.11***	0.20***	0.19***
	(0.07)	(0.05)	(0.03)	(0.02)	(0.04)	(0.05)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.012	0.025	0.050	0.035	0.106	0.097
Panel D: Regression-based convexity adjustment ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation - Regression Adj.	0.29***	0.20***	0.19***	0.15***	0.22***	0.23***
	(0.05)	(0.04)	(0.03)	(0.02)	(0.04)	(0.05)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.037	0.051	0.073	0.061	0.122	0.135
Panel E: Term-structure model convexity adjustment ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation - Term-Structure Adj.	0.29***	0.19***	0.17***	0.13***	0.23***	0.23***
	(0.05)	(0.04)	(0.03)	(0.02)	(0.05)	(0.05)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.036	0.046	0.061	0.047	0.134	0.134
Panel F: Bloomberg data ($h = 5$, Nov08-Dec18, daily overlapping)						
Deviation Bloomberg Data	0.34***	0.18***	0.12***	0.08***	0.26***	0.20***
	(0.08)	(0.05)	(0.03)	(0.03)	(0.05)	(0.05)
Observations	2267	2511	2511	2511	2511	2511
Adjusted R^2	0.047	0.037	0.027	0.017	0.170	0.107
Panel G: VIX, VIX3M, VIX6M Indices ($h = 5$, Nov08-Dec18, daily overlapping)						
Deviation VIX Indices	0.25***	0.14***	0.13***	0.12***	0.24***	0.17***
	(0.07)	(0.05)	(0.04)	(0.03)	(0.03)	(0.04)
Observations	2266	2510	2510	2510	2510	2510
Adjusted R^2	0.024	0.023	0.034	0.033	0.147	0.083
Panel H: Hedge Fund data ($h = 5$, Mar04-Nov13, daily overlapping)						
Deviation Hedge Fund Data	0.12	0.08	0.02	0.07***	0.14***	0.15***
	(0.08)	(0.05)	(0.06)	(0.03)	(0.04)	(0.05)
Observations	2209	2429	2429	2375	1938	1931
Adjusted R^2	0.007	0.008	0.001	0.013	0.041	0.045
Panel I: Markit data ($h = 21$, Sep06-Dec15, end-of-month data, non-overlapping)						
Deviation Markit Data	0.51**	0.11	0.06	0.11	0.28**	0.32**
	(0.23)	(0.27)	(0.11)	(0.09)	(0.14)	(0.14)
Observations	108	108	108	108	108	108
Adjusted R^2	0.019	-0.004	-0.006	0.009	0.076	0.117

Newey-West SEs with $3 \cdot h$ lags in parentheses, * p<.10, ** p<.05, *** p<.01

Table A.5: Predicting VIX Futures Hedged Returns with LOOP Deviations: Robustness to Return Definition - Percentage Returns

This table reports the predictability of the deviation measure for VIX futures returns hedged with variance swap forward returns over a weekly horizon (h=5) where returns are defined as the percentage return rather than the change in the futures price or change in the forward price. The deviation measure remains significant in these regressions. As noted in the paper, capturing percentage returns for VIX futures and variance swap forwards requires a dynamic trading strategy that may entail additional transaction costs in practice.

Predictability Regression: $R^{VIX}_{t+h,n} - \beta_n R^{VS}_{t+h,n} = \gamma_n' x_{t,n} + \epsilon_{t+h,n}$

Return Definitions: $R^{VIX}_{t+h,n} \equiv \frac{Fut_{t,n} - Fut_{t+h,n}}{Fut_{t,n}}, R^{VSF}_{t+h,n} \equiv \frac{F_{t,n} - F_{t+h,n}}{F_{t,n}}$

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A.I: Full sample from 2004 to 2018						
Deviation	0.19***	0.34***	0.37***	0.33***	0.32***	0.35***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.05)	(0.05)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.036	0.114	0.140	0.112	0.101	0.122
Panel A.II: Full sample from 2004 to 2018 with controls						
Deviation	0.23***	0.42***	0.43***	0.38***	0.30***	0.37***
	(0.04)	(0.04)	(0.03)	(0.04)	(0.05)	(0.05)
VIX	0.22**	0.14*	0.17*	0.10	0.01	0.20**
	(0.09)	(0.08)	(0.09)	(0.08)	(0.10)	(0.10)
RV	-0.17**	-0.05	0.00	0.06	-0.13	-0.19*
	(0.08)	(0.07)	(0.09)	(0.08)	(0.11)	(0.10)
RMRF	-0.02	-0.09***	-0.03	-0.09***	-0.02	0.01
	(0.02)	(0.03)	(0.04)	(0.03)	(0.04)	(0.04)
VLM	-0.04	0.14***	0.16***	0.06*	0.13***	0.17***
	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.04)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.045	0.150	0.187	0.150	0.139	0.154
Panel B.I: Post-crisis sample from 2010 to 2018						
Deviation	0.14***	0.29***	0.30***	0.31***	0.37***	0.37***
	(0.04)	(0.03)	(0.04)	(0.04)	(0.05)	(0.05)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.019	0.081	0.090	0.096	0.135	0.133
Panel B.II: Post-crisis sample from 2010 to 2018 with controls						
Deviation	0.21***	0.37***	0.35***	0.33***	0.33***	0.37***
	(0.03)	(0.05)	(0.04)	(0.05)	(0.05)	(0.05)
VIX	0.21***	0.21***	0.20**	0.04	-0.10	0.02
	(0.08)	(0.07)	(0.09)	(0.09)	(0.10)	(0.10)
RV	-0.08	-0.07	0.04	0.10	-0.06	-0.02
	(0.07)	(0.06)	(0.08)	(0.08)	(0.09)	(0.09)
RMRF	-0.03	-0.07**	-0.03	-0.07	-0.03	0.01
	(0.03)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)
VLM	-0.03	0.11**	0.17***	-0.00	-0.01	0.09**
	(0.04)	(0.05)	(0.03)	(0.05)	(0.04)	(0.04)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.038	0.127	0.173	0.120	0.154	0.140

Newey-West SEs with 15 lags in parentheses, * p<.10, ** p<.05, *** p<.01

Table A.6: Predicting VIX Futures Hedged Returns with LOOP Deviations: Robustness to Return Definition - Log Returns

This table reports the predictability of the deviation measure for VIX futures returns hedged with variance swap forward returns over a weekly horizon (h=5) where returns are defined as log-returns rather than the change in the futures price or change in the forward price. The deviation measure remains significant in these regressions. As noted in the paper, capturing log-returns for VIX futures and variance swap forwards requires a dynamic trading strategy that may entail additional transaction costs in practice.

Predictability Regression: $R^{VIX}_{t+h,n} - \beta_n R^{VS}_{t+h,n} = \gamma_n' x_{t,n} + \epsilon_{t+h,n}$

Return Definition: $R^{VIX}_{t+h,n} = \ln(Fut_{t,n}/Fut_{t+h,n}), R^{VSF}_{t+h,n} \equiv \ln(F_{t,n}/F_{t+h,n})$						
Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A.I: Full sample from 2004 to 2018						
Deviation	0.27***	0.42***	0.41***	0.35***	0.33***	0.35***
	(0.03)	(0.04)	(0.03)	(0.03)	(0.05)	(0.05)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.074	0.174	0.169	0.125	0.108	0.124
Panel A.II: Full sample from 2004 to 2018 with controls						
Deviation	0.28***	0.49***	0.46***	0.39***	0.30***	0.38***
	(0.03)	(0.04)	(0.03)	(0.04)	(0.05)	(0.05)
VIX	0.16*	0.13	0.16	0.11	-0.02	0.19*
	(0.09)	(0.09)	(0.10)	(0.09)	(0.10)	(0.11)
RV	-0.11	-0.02	-0.01	0.03	-0.13	-0.17*
	(0.08)	(0.09)	(0.10)	(0.08)	(0.10)	(0.10)
RMRF	0.00	-0.07**	-0.03	-0.07**	-0.01	0.02
	(0.03)	(0.03)	(0.04)	(0.03)	(0.04)	(0.04)
VLM	-0.11***	0.11**	0.16***	0.05	0.13***	0.17***
	(0.03)	(0.04)	(0.03)	(0.03)	(0.03)	(0.04)
Observations	3356	3698	3698	3644	3207	3200
Adjusted R^2	0.090	0.201	0.213	0.156	0.149	0.155
Panel B.I: Post-crisis sample from 2010 to 2018						
Deviation	0.23***	0.40***	0.35***	0.33***	0.38***	0.37***
	(0.04)	(0.03)	(0.04)	(0.04)	(0.05)	(0.05)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.055	0.160	0.119	0.112	0.145	0.137
Panel B.II: Post-crisis sample from 2010 to 2018 with controls						
Deviation	0.26***	0.48***	0.40***	0.36***	0.34***	0.37***
	(0.03)	(0.05)	(0.04)	(0.05)	(0.05)	(0.06)
VIX	0.15*	0.22**	0.19*	0.02	-0.13	0.01
	(0.08)	(0.11)	(0.10)	(0.09)	(0.10)	(0.10)
RV	-0.05	-0.06	0.07	0.12	-0.05	-0.01
	(0.08)	(0.08)	(0.09)	(0.09)	(0.10)	(0.09)
RMRF	-0.00	-0.05	-0.02	-0.07	-0.02	0.01
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)
VLM	-0.08**	0.07	0.19***	-0.00	-0.01	0.09**
	(0.04)	(0.07)	(0.04)	(0.05)	(0.04)	(0.04)
Observations	2030	2246	2246	2246	2246	2246
Adjusted R^2	0.072	0.197	0.210	0.133	0.169	0.144

Newey-West SEs with 15 lags in parentheses, * p<.10, ** p<.05, *** p<.01

Figure A.1: Examples of VIX Futures Law of One Price Deviations

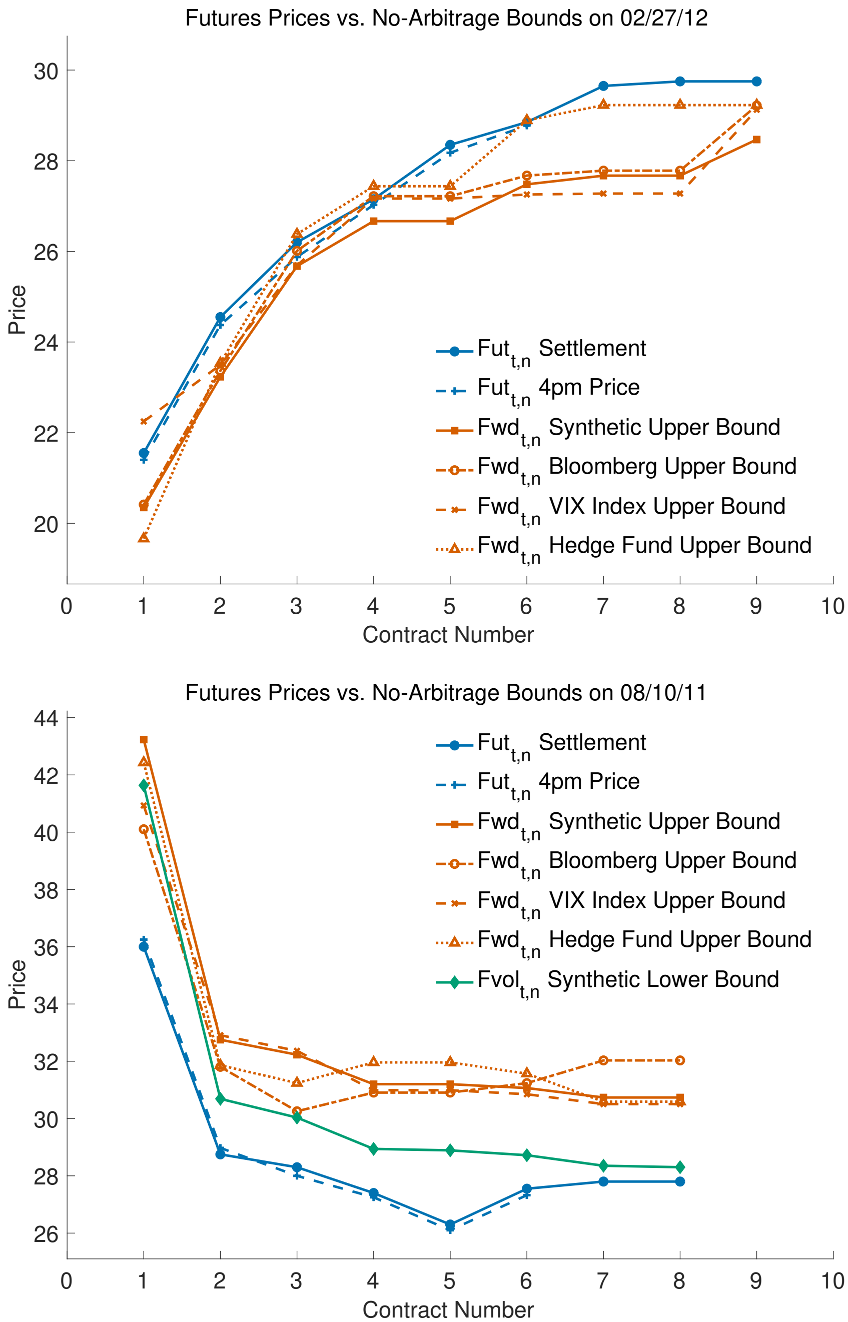

This figure compares VIX futures prices at 4pm and at 4:15pm (settlement) to variance swap forward rates estimated from a variety of alternative data sources on two different days. The results illustrate the robustness of Figure 3. Regardless of which of the four data sources are used to compute variance swap forward rates or which of the two VIX futures prices are used, the top plot features examples of static arbitrage opportunities in which the prices of VIX futures are above their non-parametric, no-arbitrage upper bounds. The bottom plot illustrates the opposite case where the prices of VIX futures are well below their upper bounds and are even below an estimate of the lower bounds.

Figure A.2: Deviation Measure Computed from Alternative Datasets

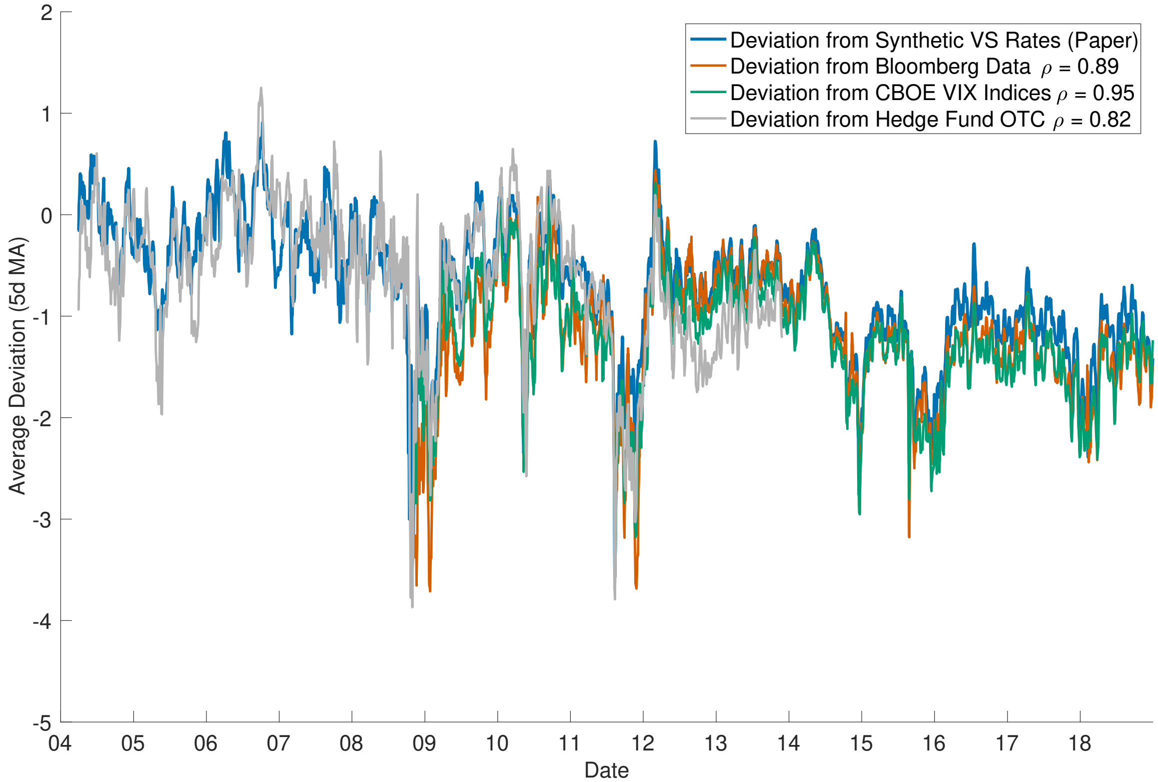

This figure plots the main deviation measure from the paper against alternative deviation measures that are computed by estimating variance swap forward rates from different datasets. For each deviation measure, the figure reports the average deviation across the front six contracts as a five-day moving average. The deviation measure from the paper is highly correlated with the alternative deviation measures. The correlation is reported in the legend over the sample period when both the main deviation measure and alternative measure are available. As before, positive values are law of one price violations. Gray shading indicates NBER recessions.

Figure A.3: Deviation Measure By VIX Futures Contract

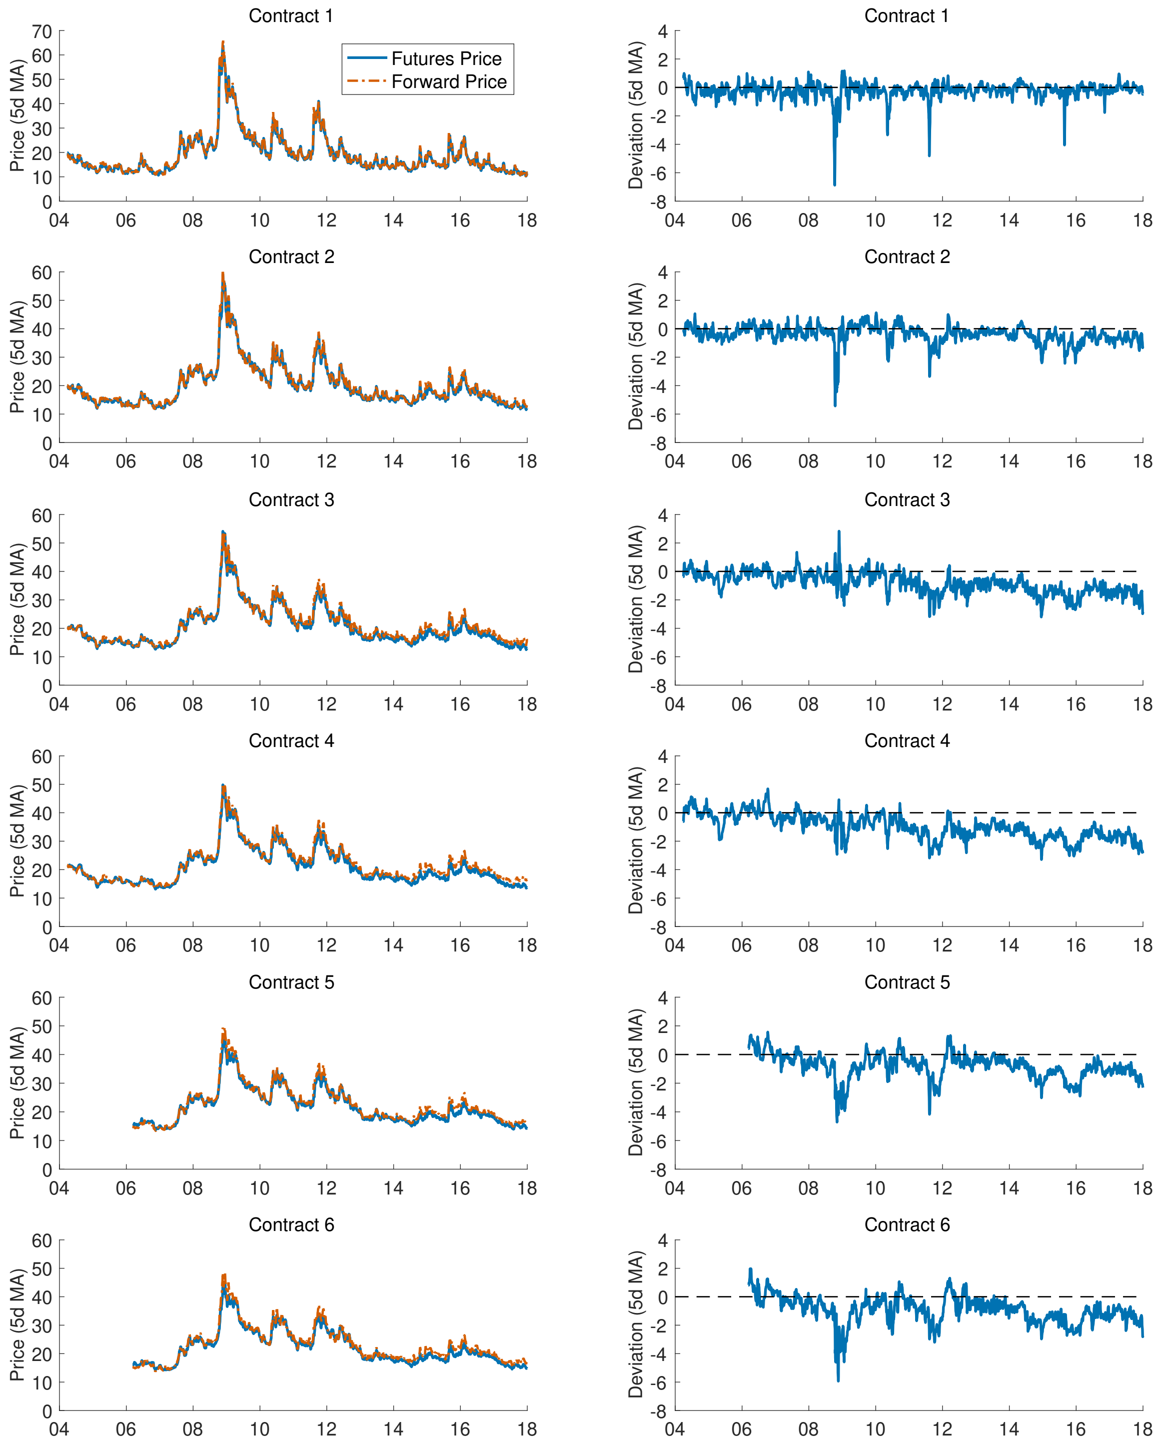

This figure plots the VIX futures price and variance swap forward rate (upper bound) on the left and deviation measure, or difference between the futures price and forward price, on the right for each contract. The prices and deviation measure for the longer-dated contracts become available later in the sample, motivating the 2007-2018 and 2010-2018 sample periods that are used in some of the regression and summary statistics analysis to provide a balanced panel across contracts.

Figure A.4: Distribution of VIX Futures Prices Relative to the No-Arbitrage Upper Bound from 2007 to 2018

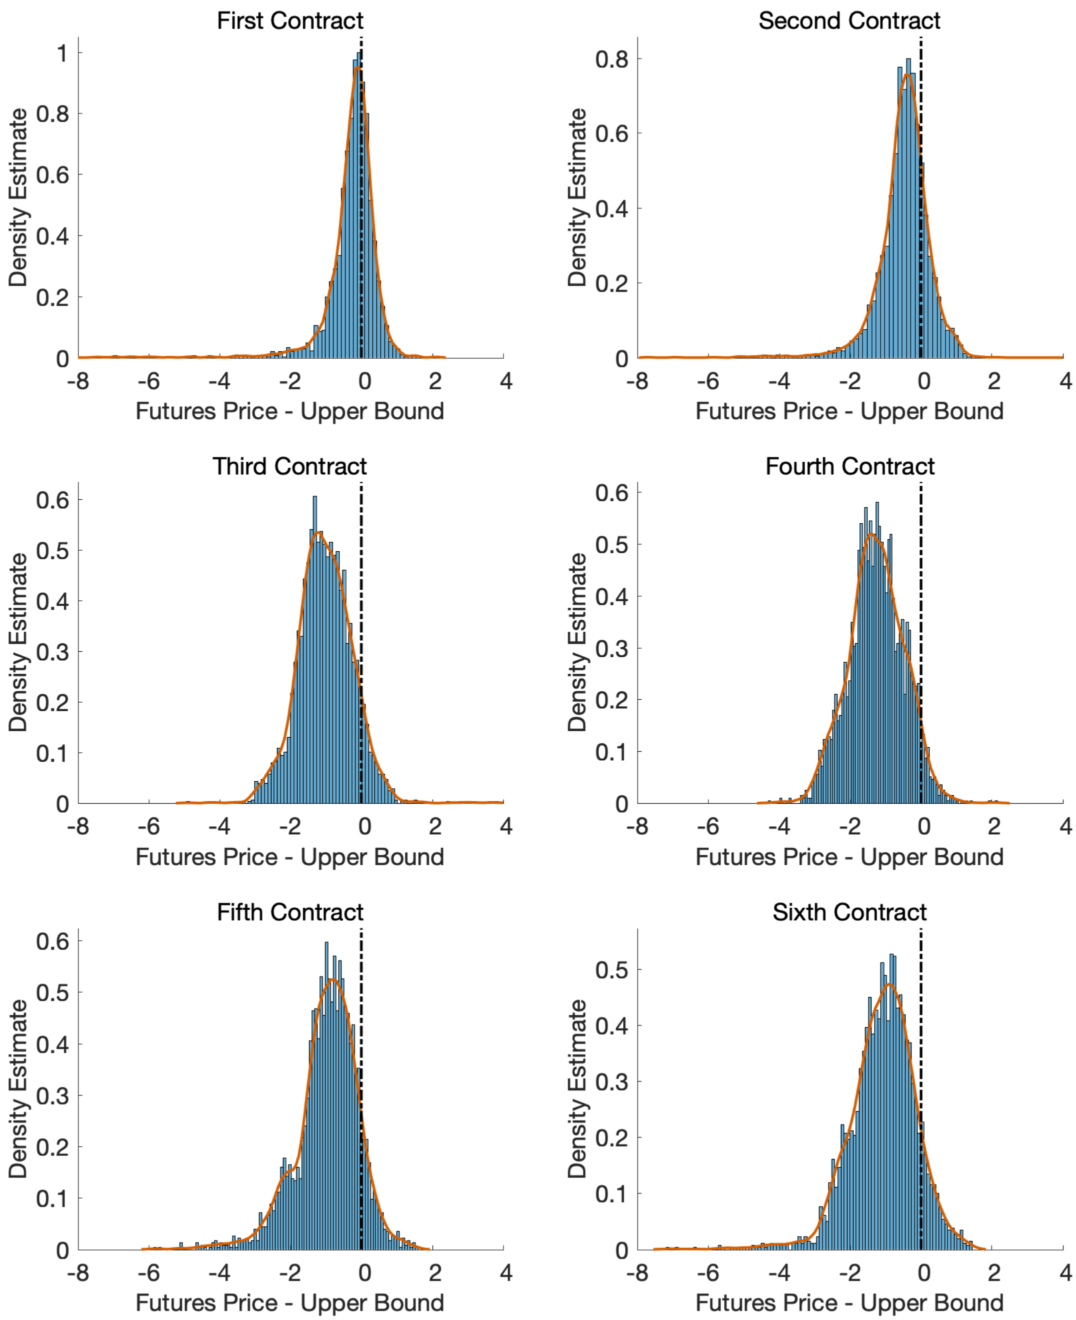

This figure plots histograms and kernel density estimates of the law of one price deviation measure by VIX futures contract from 2007 to 2018. The distribution is negatively skewed for the front contract and second contract. The histograms indicate the presence of law of one price violations from the probability mass for the deviation measure being greater than zero which corresponds to VIX futures prices being greater than the upper bound. The histograms also reveal how the deviation measures exhibit large negative values that may also represent law of one price violations to the extent that VIX futures prices go below volatility swap forward rates. The lower bound violations cannot be measured directly from these histograms, but are reported in the next figure.

Figure A.5: Distribution of VIX Futures Prices Relative to the No-Arbitrage Lower Bound from 2007 to 2018

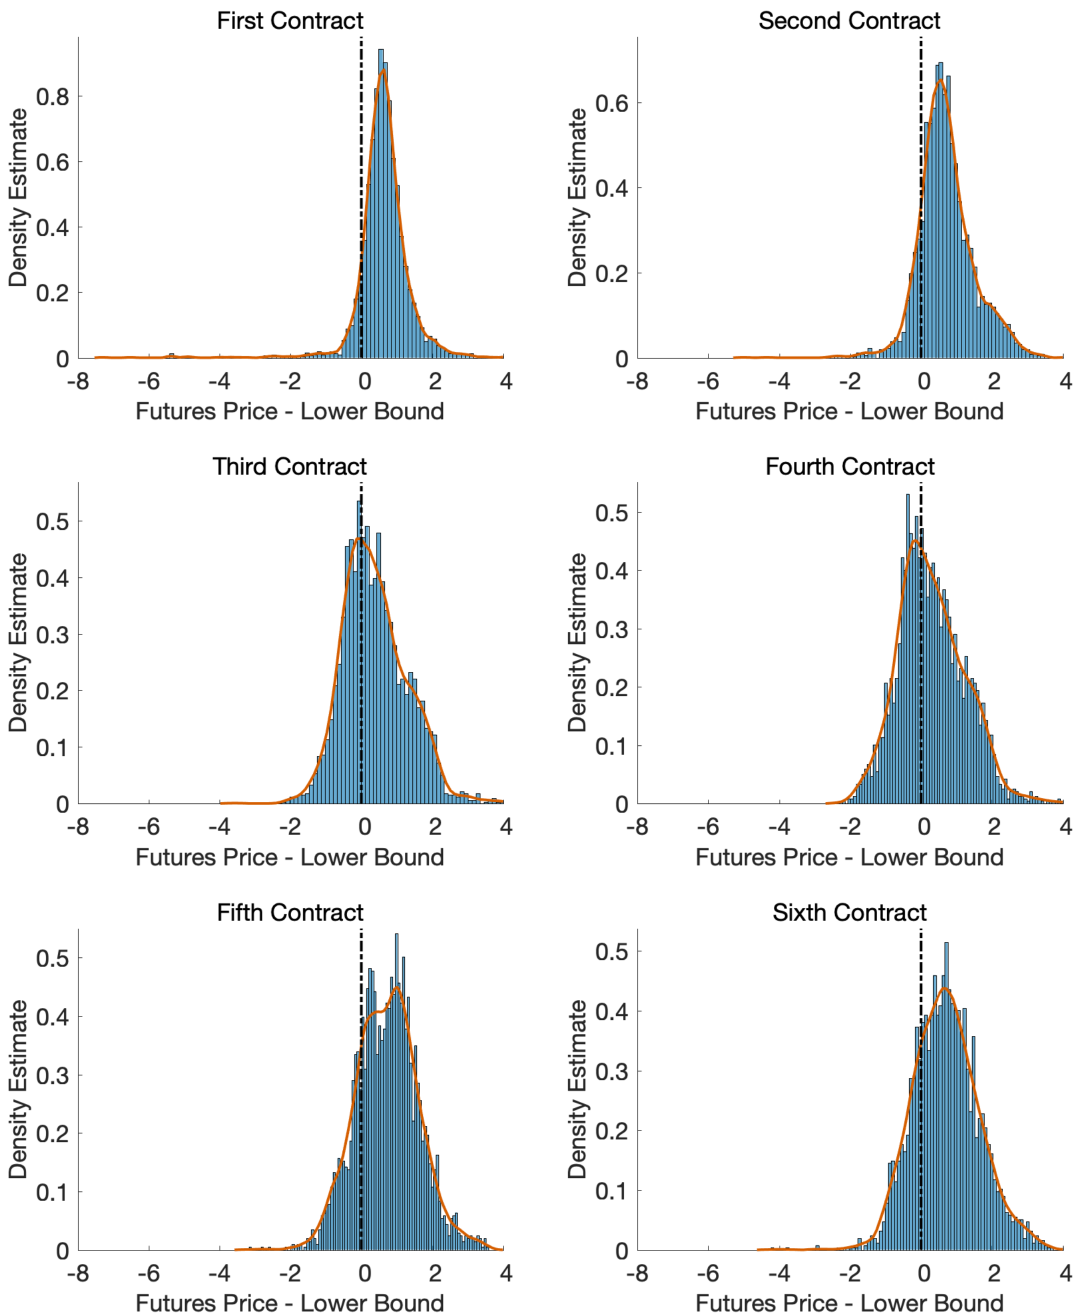

This figure plots histograms and kernel density estimates of VIX futures prices relative to their law of one price lower bound from 2007 to 2018. The lower bound is computed as $Fwd_{t,n} - (UB_{t,n} - LB_{t,n})$ where $UB_{t,n}$ is the variance swap forward rate and $LB_{t,n}$ is the volatility swap forward rate estimated from a no-arbitrage term-structure model on day t for the n-th contract following the approach in Van Tassel (2019). In this case, the histograms indicate the presence of law of one price violations from the probability mass below zero which corresponds to cases in which the VIX futures price is below the lower bound. Lower bound violations are more pronounced for the longer-dated contracts consistent with Table 3.

Figure A.6: Impulse Response Functions of Deviation and Dealer Position to VIX and Dealer Position Shocks Across Trivariate VAR Specifications

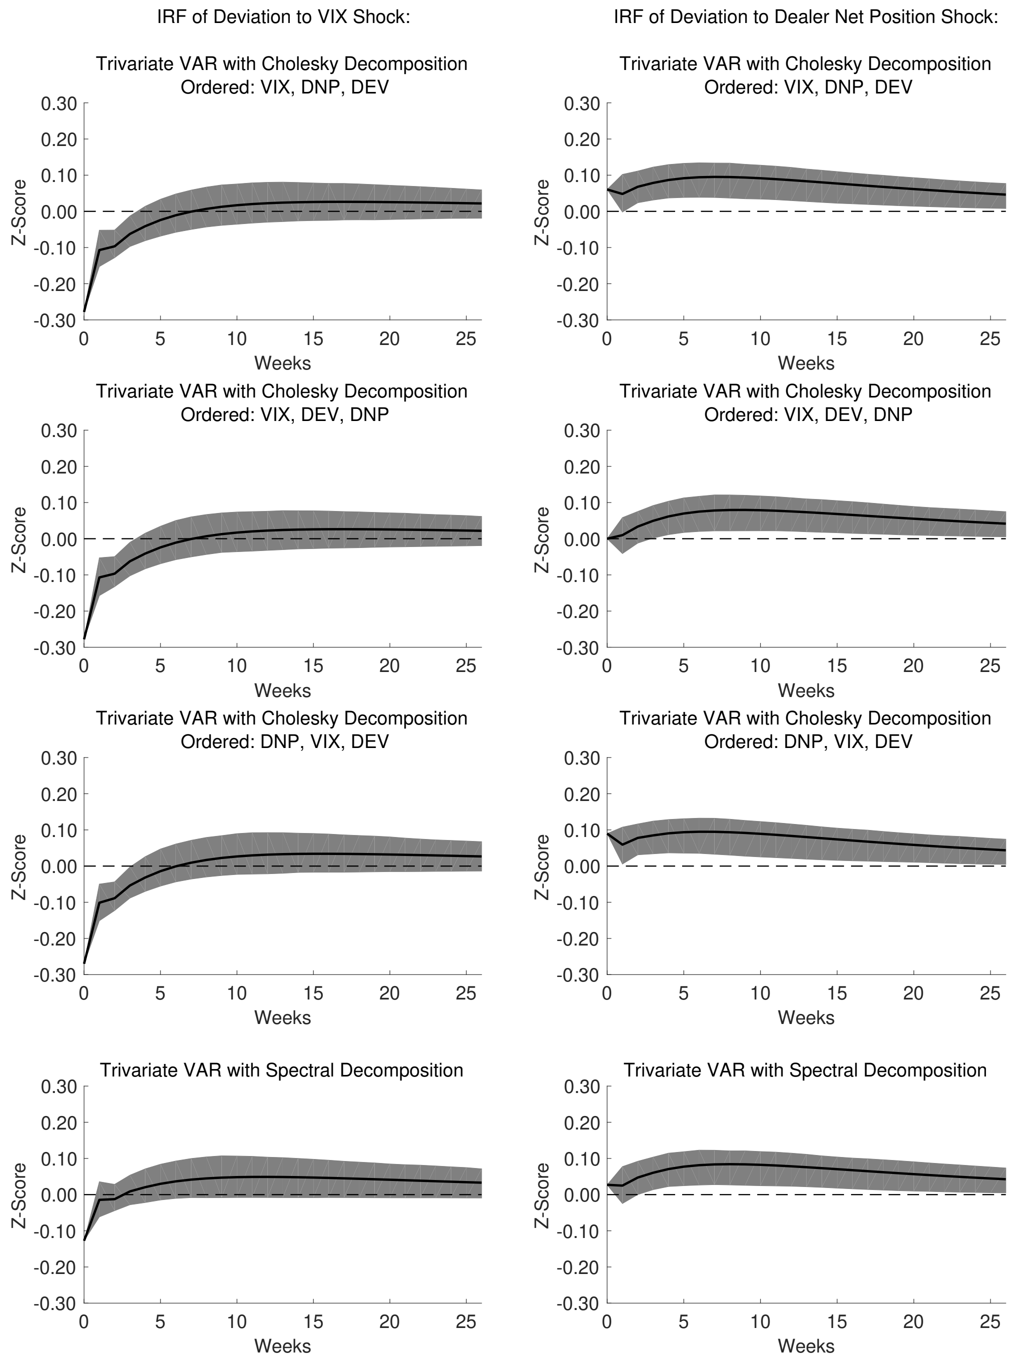

This figure plots the impulse response function of of the no-arbitrage deviation and dealer position to VIX and dealer position shocks across different trivariate VAR specifications. The first three rows report different orderings of the variables for a Cholesky decomposition. The fourth row reports the IRFs for a spectral decomposition. The variables in the vector autoregression are $y_t = [Deviation_t\; VIX_t\; Dealer\ Position_t]$. The VAR is estimated using weekly data from 2010 to 2018 ($T = 469$). Across specifications, the IRFs take on similar shapes. The deviation measure declines in response to a risk shock and increases in response to a demand shock. The impact of the risk shock dies out after a few weeks whereas the demand shocks are more persistent.

Figure A.7: Deviation Measure versus Risk and Demand Variables

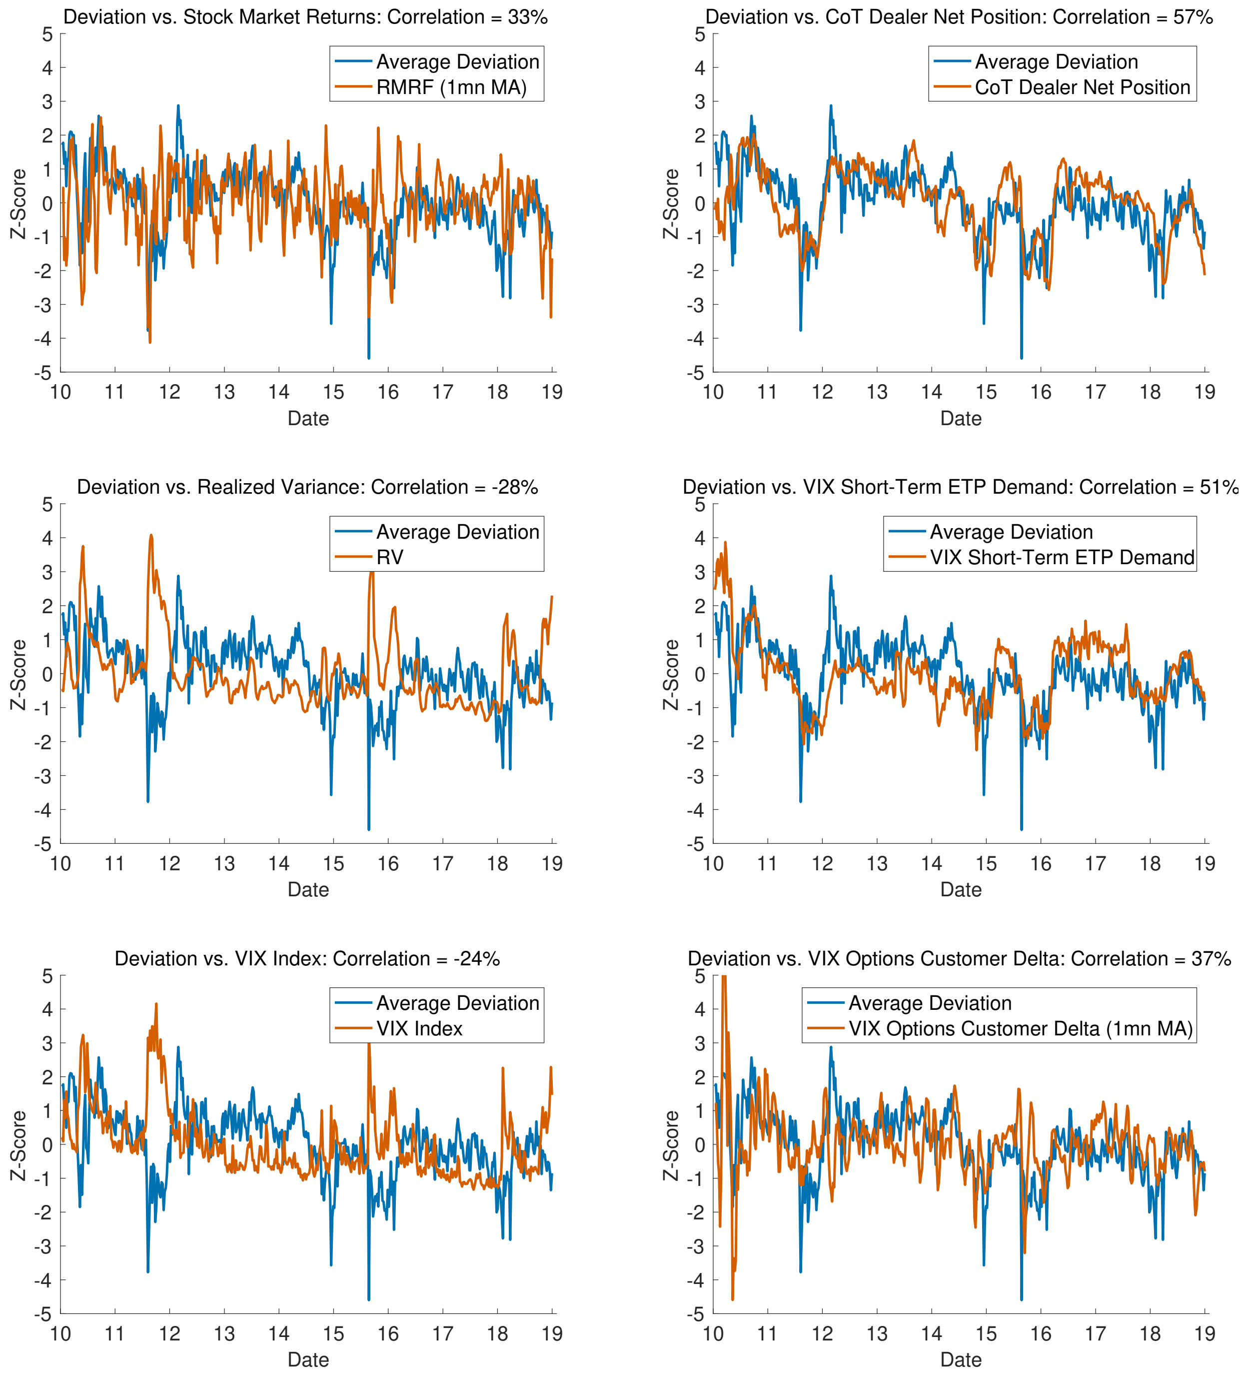

This figure plots the deviation measure against different risk and demand variables. The risk variables include stock market returns, realized variance, and the VIX Index. Increases in risk as measured by negative stock market returns or increases in volatility are negatively correlated with the deviation measure. The demand variables are Dealer Position from the CoT Report, VIX ETP demand, and VIX options customer delta. The demand variables are positively correlated with the deviation measure. The sample period is 2010 to 2018 using weekly data ($T = 469$).

Figure A.8: Impulse Response Functions of Deviation Measure to Risk and Demand Shocks in Bivariate VARs

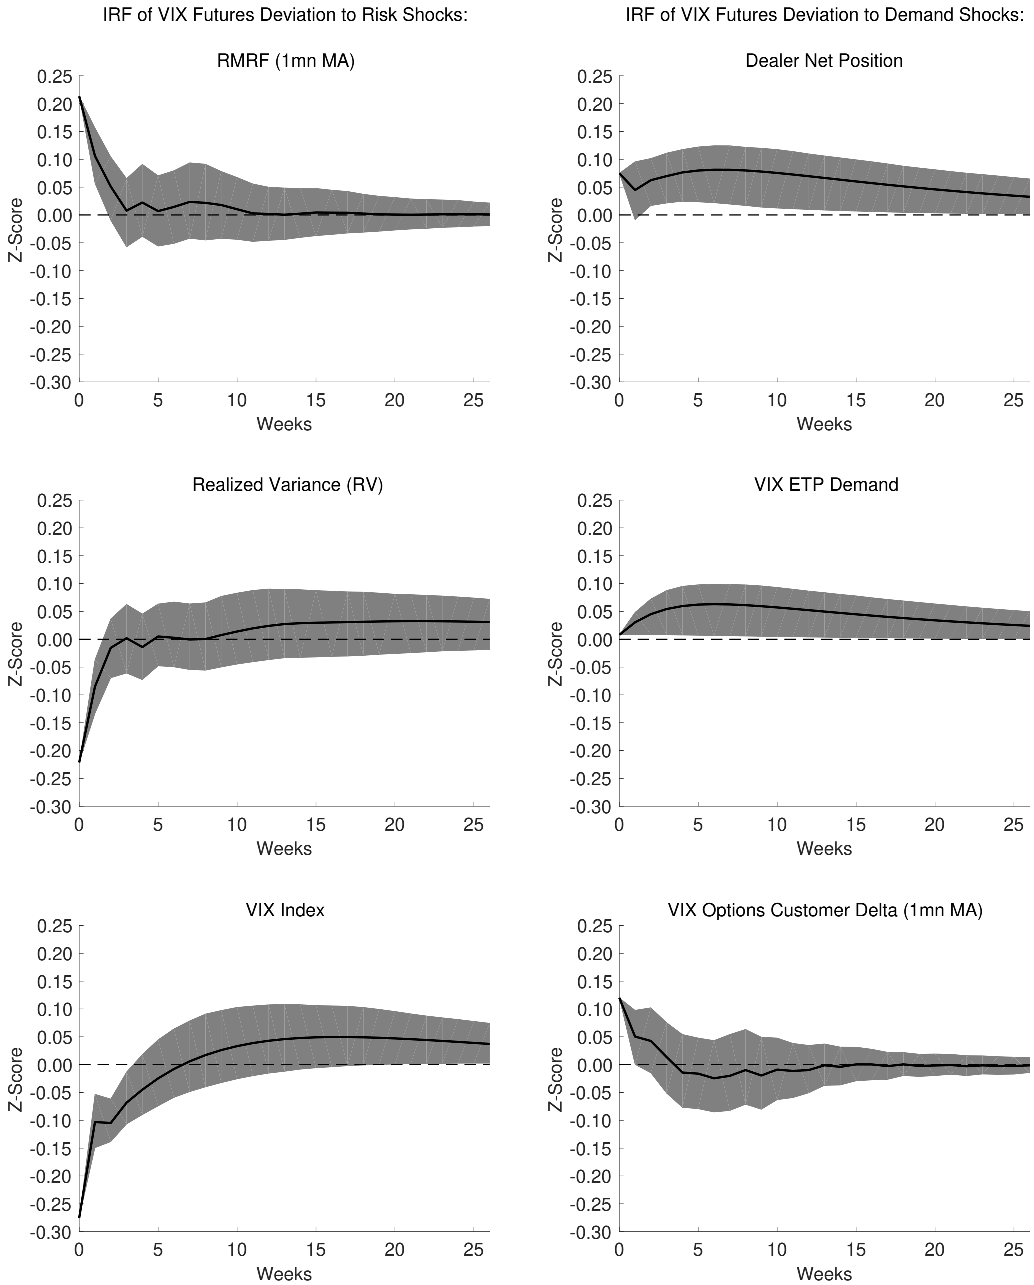

This figure plots impulse response functions from bivariate vector autoregressions to illustrate how the deviation measure reacts to risk and demand shocks. The left column reports the IRFs from bivariate VARs with risk variables. The right column reports the IRFs from bivariate VARs with demand variables. The IRFs are from a Cholesky decomposition with the deviation measure ordered second. The 95% confidence intervals in gray are block bootstrapped. The lag length is selected using the SBIC criterion. Similar to the trivariate VAR discussed in the paper and the time-series plots, the deviation measure decreases when risk increases and increases when demand increases. The magnitude of the response is larger for the risk shocks, but more persistent for the dealer position and VIX ETP demand shocks.

Figure A.9: Impulse Response Function of Deviation Measure Across Contracts

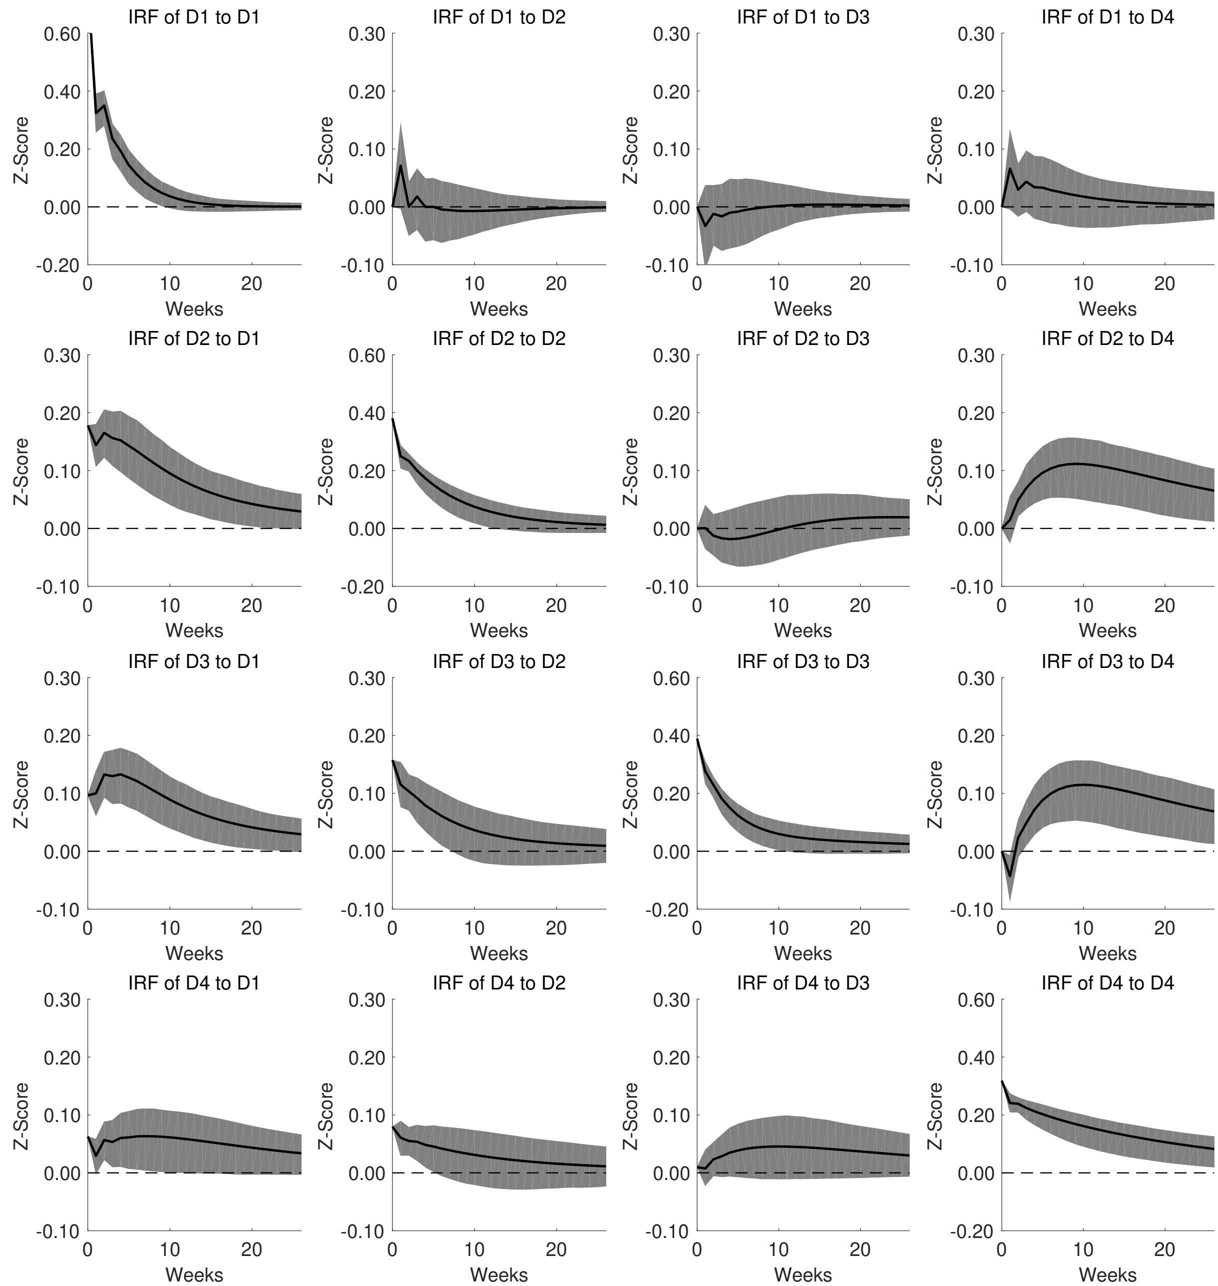

This figure reports the IRFs from a weekly VAR with $y_t = [Dev_{t,1}\ Dev_{t,2}\ Dev_{t,3}\ Dev_{t,4}]$ from 2010 to 2018. The goal of these IRFs is to highlight how shocks to different contracts propagate to other contracts. The results indicate that shocks to the front contract correspond to increases in the second and third contract that die out after a few months. Shocks to the fourth contract result in delayed increases in the second and third contract. The IRFs are computed using a Cholesky decomposition with the variables ordered as the deviation for the front, second, third, and fourth contract respectively. The results are reported using weekly data for comparison to the other IRFs reported in the paper and Appendix and focus on the front four contracts for parsimony. Similar qualitative results hold for daily data and for the front six contracts.

Figure A.10: Event Study of VIX Futures Deviation Around Large Risk and Demand Shocks

This figure plots the event time reaction of the deviation measure and VIX index to the largest increases and decreases in stock market returns, the VIX index, dealer position, and VIX options customer delta throughout the sample. The figure reports the change in the VIX index and average deviation measure across the front six contracts in the ten days before and after the event.

Figure A.11: VIX Futures Trading Strategy: Position and Post-Crisis Sharpe Ratios

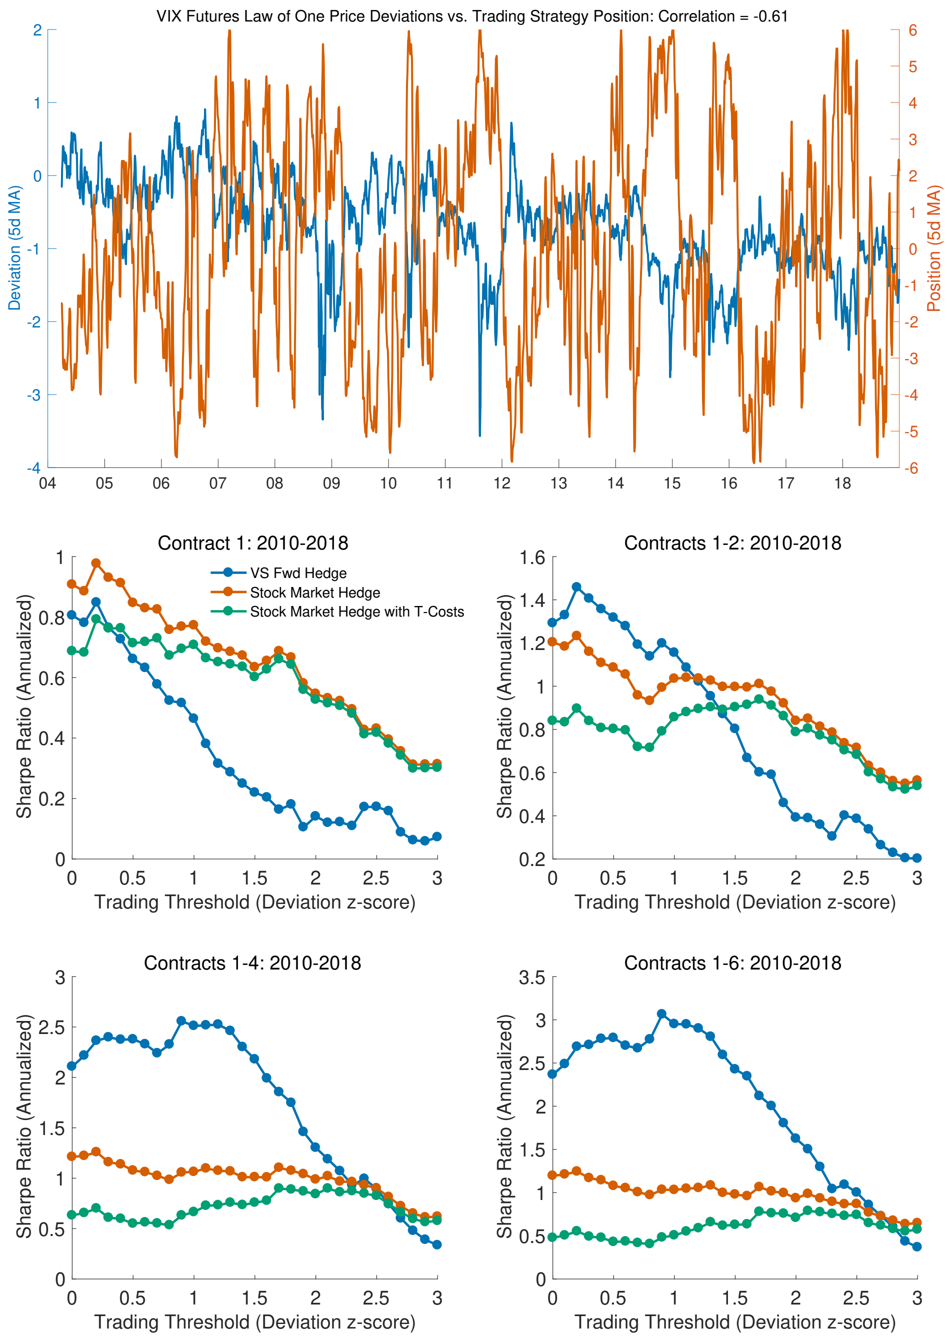

The top figure plots the number of net positions in the relative value trading strategy against the deviation measure over time. A long position corresponds to buying VIX futures that are hedged with variance swap forwards. Since the strategy trades the front six contracts, the number of net positions is bounded between -6 and 6. The plot illustrates the negative correlation between the deviation measure and the number of net positions. When the deviation measure is low (high), the strategy tends to buy (sell) VIX futures that are hedged with variance swap forwards. The bottom plot reports the SRs for the different strategies varying the number of contracts traded and threshold. This is analogous to Figure 7 but for the 2010 to 2018 post-crisis period.