Federal Reserve Bank of New York
Staff Reports

The Law of One Price in Equity Volatility Markets

Peter Van Tassel

Staff Report No. 953
December 2020

The Law of One Price in Equity Volatility Markets
Peter Van Tassel
Federal Reserve Bank of New York Staff Reports, no. 953
December 2020
JEL classification: G12, G13, C58

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 bounds. The deviations widen during periods of market stress and predict the returns of VIX futures. A relative value trading strategy based on the deviation measure earns a large Sharpe ratio and economically significant alpha-to-margin. There is evidence that systematic risk and demand pressure contribute to the variation in the no-arbitrage deviations over time.

Key words: limits-to-arbitrage, VIX futures, variance swaps, volatility, return predictability

---

Van Tassel: Federal Reserve Bank of New York (email: peter.vantassel@ny.frb.org). The author thanks Ing-Haw Cheng, Richard Crump, Bjorn Eraker, Travis Johnson, David Lando, David Lucca, Lasse Pedersen, Ivan Shaliastovich, and seminar participants at the American Finance Association, Wisconsin-Madison Junior Finance Conference, the Federal Reserve Bank of New York, and the Copenhagen Business School for comments. 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. The law states that assets with identical payoffs must have the same price. Otherwise, competitive traders will exploit the deviations to make arbitrage profits, thereby eliminating law of one price violations so that financial markets behave as if no arbitrage opportunities exist.

Financial markets are a prime testing ground for the law of one price. With low transaction costs and the ability to buy and sell similar securities, traditional theories predict that the law of one price should hold broadly across financial markets. A growing literature, however, challenges this view. There are now well-documented examples of law of one price violations in equity, fixed income, credit, and currency markets. When assets with closely related payoffs trade at significantly different prices for prolonged periods of time, there must be market frictions or inefficiencies that impede arbitrage activity, even for sophisticated traders and financial institutions. By documenting these anomalies, we can better interpret market prices and understand how financial markets function (Lamont and Thaler (2003), Gromb and Vayanos (2010)).

This paper contributes to the limits-to-arbitrage literature by providing 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 an opportunity to test 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 are accurate, leaving open the possibility of observing arbitrage violations.

This paper studies no-arbitrage relationships in equity volatility markets by comparing the prices of VIX futures to S&P 500 index options. The starting point is the observation that VIX futures prices are bounded above by variance swap forward rates (Carr and Wu 2006).[1] The paper uses S&P 500 index options to non-parametrically estimate variance swap forward rates, thus obtaining an option-implied upper bound for VIX futures. The no-arbitrage deviation measure, or arbitrage spread, is then defined for each futures contract

[1]The upper bound follows from the definition of the VIX index and Jensen's inequality. For each futures contract, the upper bound is the one-month variance swap forward rate starting on the futures expiration date, expressed in volatility units.

as the difference between the futures price and the option-implied upper bound.

Figure 1 plots the no-arbitrage deviation measure averaged across the front six futures contracts from March 26, 2004 to December 31, 2018.[2] Positive values indicate static arbitrage opportunities in which futures prices exceed their option-implied upper bounds. In these cases, traders can sell futures and pay fixed in forwards to lock in a riskless profit. Disaggregating the results by contract, the paper shows that there are frequent upper bound violations throughout the sample. In addition, the paper finds that the upper bound is relatively tight, with large negative deviations indicating that VIX futures are cheap relative to index option prices. Combining the upper bound with estimates from a term-structure model, the paper estimates the lower bounds for VIX futures prices and finds evidence of lower bound violations as well. Taken together, the evidence suggests that there are large and significant deviations between the prices of VIX futures and their option-implied bounds.

The paper then explores the return predictability of the deviation measure to provide further evidence that it is identifying an arbitrage spread. The paper finds that the deviation measure significantly predicts the returns of VIX futures relative to variance swap forwards. The predictability holds across contracts, sample periods, and horizons, and is robust to competing predictors like the variance risk premium. A trading strategy based on the deviation measure earns a large Sharpe ratio and economically significant alpha-to-margin across a range of assumptions, consistent with the law of one price interpretation. In addition, the paper finds that the deviation measure predicts the returns of VIX futures and variance swap forwards relative to the stock market, providing evidence that both the futures and options markets are contributing to the predictability of the deviation measure.

What drives the law of one price deviations and return predictability over time? The limits-to-arbitrage literature suggests channels like demand shocks and financial constraints (Shleifer et al. (1990), Liu and Longstaff (2003), Adrian and Shin (2013)). The paper explores these frictions by studying whether a range of risk and demand variables are related to the deviation measure. VAR and panel regression analysis indicate that the deviation measure is decreasing in systematic risk. When the stock market declines or volatility increases, the prices of VIX futures increase less than the prices of index options, leading to a decline in the deviation measure. In addition, the paper finds that the deviation measure is increasing in various proxies for the demand to buy VIX futures, consistent with demand pressure playing a role.

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 risk. Cheng

[2] There are only four contracts at the start of the sample. Six monthly contracts are available starting in late 2006 on a consistent basis. The paper analyzes how the results differ by contract and over time.

(2018) proposes a similar mechanism for the VIX premium. In addition to hedging demand, the VAR analysis highlights that dealer net positions increase following a positive shock to the deviation measure. To the extent that dealer positions are a veil for retail demand (Dong 2016), the results suggest that retail traders may chase momentum in the VIX futures market. The paper also provides event study analysis to investigate how the deviation measure behaves around margin changes, end-of-quarter dates, and FOMC and non-farm payroll announcements. There is some evidence that margin increases attenuate the deviation measure's relationship with risk and that the deviation measure declines following FOMC announcements and leading into quarter-end. Comparing the deviation measure to the VIX premium from Cheng (2018), the paper finds the two measures are related but have different interpretations. While the VIX premium reflects the implied volatility risk premium, the deviation measure represents a no-arbitrage violation or relative mispricing.

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

2 Equity Volatility Markets

2.1 Relation to the Literature

This paper contributes to the literature on anomalies and the limits of arbitrage. Examples of related studies 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), mortgage-backed securities (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 (Du et al. (2018)) in foreign exchange. Across asset classes, law of one price deviations receive significant attention because they pose a model-free challenge to the most basic assumption in many asset pricing models — that investors prefer more wealth to less.

In the equity volatility literature, studies of anomalies often focus on the index options market or the VIX futures market in isolation. Early papers in the option pricing literature emphasized the variance risk premium and implied volatility smile for out-of-the-money put options as irregularities not explained by the Black-Scholes-Merton model (Bates 2000). One strand of the literature developed more sophisticated option pricing models with stochastic volatility and jumps to account for these patterns (Heston (1993), Duffie et al. (2000), Madan et al. (1998)). Even with these improvements, the no-arbitrage models still struggle to explain some of the empirical properties of option prices (Bates 2003). Another strand of the literature studies 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, with demand for specific options impacting the prices of other options with related, unhedgeable features.

Building on these studies, the literature on VIX futures emphasizes the importance of demand shocks for understanding prices and risk premia. Cheng (2018) finds that the VIX premium, or implied volatility risk premium, declines when risk increases alongside decreases in dealer hedging demand. Dong (2016) argues that VIX ETP demand impacts the prices of VIX futures through dealer hedging activity. Mixon and Onur (2019) provide a related study of demand pressure using regulatory data that includes dealer net positions by contract at a daily frequency. 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. The estimated demand effects in Mixon and Onur (2019) are quantitatively small, however, in comparison to the no-arbitrage bounds for VIX futures.

This paper contributes to the literature by studying law of one price deviations across equity volatility markets. While the relationship between the pricing of VIX futures and index options is well understood from a theoretical perspective, this paper provides a comprehensive empirical study that documents large no-arbitrage deviations. The deviation measure allows for a direct measurement of arbitrage opportunities and finds a high frequency of violations.[3] In addition, this paper adds a new predictor to the VIX futures literature to complement the VIX premium of Cheng (2018) and the slope factor of Johnson (2017). The predictability of the deviation measure is strongest for VIX futures in excess of variance swap forwards, a relative value finding that is new to the literature. Finally, the paper provides new evidence on the importance of the risk and demand channels in driving the deviation measure. While the response to a risk shock is larger in magnitude over shorter horizons,

[3] 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 estimates synthetic variance swap forward rates from the S&P 500 index options market to obtain an upper bound, thus avoiding VIX options which are less liquid and only available starting in 2006.

demand shocks are more permanent according to VAR analysis. The results indicate that demand by itself is not sufficient to explain the negative response of the no-arbitrage deviation measure to increases in risk. Overall, the results raise new puzzles regarding the frequency of no-arbitrage violations and the significant return predictability of the deviation measure.

2.2 Equity Volatility Products and Market Participants

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)). S&P 500 index options started trading on the Chicago Board Options Exchange (CBOE) in 1983.

Variance swaps are over-the-counter derivatives that allow investors to hedge and speculate on realized volatility over different horizons. The only cashflow occurs at maturity and is equal to the difference between the fixed variance swap rate and the floating amount of realized variance that the underlying asset exhibits over the life of the swap. The fixed rate is priced to make the swap costless to enter at the time of trade.

Variance swaps gained traction in the late 1990s alongside theoretical developments that showed how to replicate variance swap payoffs with a static portfolio of options and a dynamic trading strategy in the underlying (Demeterfi et al. (1999), Bakshi and Madan (2000), and Britten-Jones and Neuberger (2000)). In 2003, the CBOE revised its definition of the VIX index to follow the no-arbitrage formula for pricing variance swaps (Carr and Wu (2006), CBOE (2019)). In recent work, Martin (2017) has extended the generality of variance swap pricing.

VIX futures were introduced by the CBOE in 2004. The payoff to a VIX futures contract is the difference between the futures price and the VIX index at maturity. The VIX index is the square root of a one-month synthetic variance swap rate. The swap rate is "synthetic" because it is computed from a portfolio of S&P 500 index options. This relationship binds together the pricing of VIX futures and index options.

Today, S&P 500 index options and VIX futures trade in large and liquid exchange-based markets. Figure 2 provides a brief summary of these markets including their growth over time and a breakdown of investor positioning in VIX futures. The top left plot illustrates the size of the markets which have exhibited significant growth in recent years. While the VIX futures market experienced a nearly 10-fold increase in open interest over the decade ending in 2018, the index options market was still 7-times larger than the VIX futures market in

2018 as measured by average open interest in units of Black-Scholes vega.[4]

Within the VIX futures market, the top right plot shows that the growth in open interest has coincided 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 VIX futures or index options directly. One explanation for the simultaneous growth of futures and ETPs is that dealers issue ETPs and hedge their exposures in the underlying futures market. The bottom left plot supports this hypothesis. Dealer positions from the CFTC's Commitment of Traders (CoT) report closely match the magnitude and time-series variation of VIX ETP net vega (Dong 2016).

To the extent that dealer positions are a veil for retail demand, who absorbs the retail demand shocks? The bottom right plot addresses this question by plotting net positions for different trader types including dealers, leveraged funds, asset managers, and other reportable traders from the CoT report. The largest net positions by magnitude belong to dealers and leveraged funds which exhibit a correlation of -91% from 2010 to 2018. One interpretation of this result is that hedge funds absorb retail demand shocks that are passed through by dealers from the VIX ETP market. An implication is that demand shocks may be an important factor for understanding no-arbitrage deviations across the VIX futures and index options markets. The paper investigates this hypothesis in Section 5.

3 No-Arbitrage Deviation Measure

3.1 VIX Futures Bounds

The price of a VIX futures contract is the risk-neutral expected value of the VIX index at maturity T,

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

Variance swap forwards provide an upper bound for this payoff. Consider a variance swap forward whose floating leg pays the realized variance of the S&P 500 index between the futures expiration date and the index options expiration date thirty calendar days later. Applying results from Carr and Wu (2009), a combination of index options with different strikes and maturities can be used to closely-approximate the one-month variance swap forward rate as,

$$Fwd_{t,T,T+1} = E_t^{\mathbb{Q}}[RV_{T,T+1}] \approx E_t^{\mathbb{Q}}[VIX_T^2]. \tag{2}$$

[4]S&P 500 index options had an average open interest of $3.4 billion in Black-Scholes-Merton 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.

It follows from Jensen's inequality that variance swap forwards provide an upper bound for the prices of VIX futures contracts,

$$\begin{aligned} Fwd_{t,T,T+1} &= E_t^{\mathbb{Q}}[VIX_T^2] \\ &\geq E_t^{\mathbb{Q}}[VIX_T]^2 \\ &= Fut_{t,T}^2. \end{aligned} \tag{3}$$

Violations of this bound are a static arbitrage opportunity. If the futures price exceeds the bound, an arbitrageur can sell ω of the VIX futures contract and buy the variance swap forward (pay fixed) to obtain the payoff Π_T equal to,

$$\Pi_T \equiv \text{Payoff at } T{=}VIX_T^2 - Fwd_{t,T,T+1} - \omega \cdot (VIX_T - Fut_{t,T}). \tag{4}$$

Since there are no cash flows at t, this trade is an arbitrage if $\Pi_T \geq 0$ for any $VIX_T > 0$. The Appendix shows that setting the hedge ratio to $\omega = 2 \cdot Fut_{t,T}$ results in an arbitrage when the upper bound is violated.

In addition, the Appendix shows that VIX futures are bounded below by volatility swap forward rates. While the square-root in the definition of the VIX is convenient for expressing the index in Black-Scholes-Merton implied volatility units that are familiar to option traders, the square-root introduces a wedge between the pricing of VIX futures and index options.

3.2 No-Arbitrage Deviation Measure

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

$$Deviation_{t,n} \equiv Fut_{t,n} - \sqrt{Fwd_{t,n}}. \tag{5}$$

This measure is the difference between the n-month futures price $Fut_{t,n}$ and its upper bound, the square-root of the one-month variance swap forward rate starting on the futures expiration date denoted as $Fwd_{t,n}$. The measure is expressed in annualized volatility units. It is based on the futures upper bound rather than the lower bound because variance swap forward rates can be estimated non-parametrically from index option prices with minimal assumptions.[5]

The deviation measure has a straightforward interpretation. When the measure is posi-

[5]The lower bound for the futures price is the one-month volatility swap forward rate starting on the futures expiration date. Compared to variance swaps, volatility swaps are traded less frequently and are more difficult to replicate from index option prices.

tive, the futures price is above its upper bound. This indicates a static arbitrage opportunity or law of one price violation. When the measure is negative and large in magnitude, the futures price is cheap relative to its forward rate. The paper finds that the upper bound is relatively tight. This means that large negative deviations are often law of one price violations for the lower bound. The paper estimates the lower bound as the non-parametric upper bound minus the difference between upper and lower bounds from a dynamic term-structure model following Van Tassel (2020).

Table 1 shows how to compute the deviation measure on two example dates and Figure 3 plots the results. The paper estimates the variance swap forward curve by assuming flat forward rates between index option maturities as in Carr and Wu (2009) and the construction of the VIX index.[6] The forward rate for each futures contract is the average forward rate between the futures expiration and the index options expiration thirty calendar days later. The deviation measure is the difference between the futures price and its corresponding one-month forward rate.

For example, consider the upper bound violations on February 27, 2012 in which the futures price is more than 1% above the forward rate for several contracts. For the front-month and second-month contracts, the variance swap forward rates are computed as $\sqrt{Fwd_{t,1}} = \sqrt{.0412} = 20.30\%$ and $\sqrt{Fwd_{t,2}} = \sqrt{(2/30) \cdot (.0412) + (28/30) \cdot .0597} = 24.17\%$. The resulting deviations are $Deviation_{t,1} = 21.40\% - 20.30\% = 1.10\%$ and $Deviation_{t,2} = 24.38\% - 24.17\% = 0.20\%$. The second example on February 8, 2018 features large negative deviations in which futures prices are below estimates of their lower bounds. These lower bound violations follow the large spike in the VIX and the turmoil in the VIX-ETP market on February 5, 2018.

3.3 Data

The paper estimates the no-arbitrage deviation measure from March 26, 2004 to December 31, 2018 using VIX futures prices and synthetic variance swap rates. The synthetic variance swap rates are estimated from OptionMetrics data following Van Tassel (2020) for traditional expirations on third Fridays of the month with at least two-weeks to maturity. The VIX futures prices are synchronized with the option quotes using either daily settlement values from the CBOE or intraday quotes from Thomson Reuters Tick History (TRTH).[7] To rule

[6]For example, let VS_{t,T_1} and VS_{t,T_2} be variance swap rates for the first and second maturities. The instantaneous forward rate between these maturities is assumed to be constant and equal to $Fwd_{t,T_1,T_2} = (VS_{t,T_2} \cdot T_2 - VS_{t,T_1} \cdot T_1)/(T_2 - T_1)$. The results are robust to the interpolation method.

[7]For the baseline measure, the VIX futures prices are CBOE settlement values until March 3, 2008 and then 4pm mid-quotes from TRTH to be synchronous with the index option quotes from OptionMetrics. The Appendix shows that there is little difference between the deviation measure defined using settlement prices

out concerns that the results are driven by the interpolation method, asynchronicity of futures and option prices, bias stemming from the upper bound, or data sources, the Appendix shows that the baseline deviation measure is highly correlated with alternative measures that are computed from different estimation methods and data sources. For example, measures that adjust for the bias from the upper bound using a term-structure model or regression are 97% to 100% correlated with the baseline measure. In addition, the baseline measure is highly correlated with alternative measures that are computed from interpolated variance swap rates, VIX settlement prices that are not synchronized with option quotes, Bloomberg synthetic variance swap rates, CBOE volatility indices, and over-the-counter variance swap quotes from Markit. Compared to the alternative datasets, the baseline measure is advantageous because it is available over the full sample period at a daily frequency.

3.4 Deviation Time-Series

Figure 1 plots the deviation measure averaged across the front six futures contracts from March 26, 2004 to December 31, 2018. The daily time-series highlights the variability of the measure. The one-month moving average shows its persistence. During the early years in the sample the average deviation is positive on a significant fraction of dates, indicating the presence of arbitrage opportunities. Since 2012 the average deviation is positive less often, but law of one price violations still occur for individual contracts, particularly for the front contracts. The plot also highlights how the deviation measure responds to changes in risk. Around events such as the financial crisis, equity flash crash, and S&P downgrade of U.S. debt, the deviation measure declines alongside increasing prices for VIX futures and index options. The decrease in the deviation measure indicates that the prices of VIX futures do not increase by as much as variance swap forward rates during these periods of heightened systematic risk.[8]

3.5 Deviation Summary Statistics

Table 2 presents summary statistics of the deviation measure for a balanced panel of the front six monthly contracts from January 3, 2007 to December 31, 2018. The last column averages the statistics across contracts. Panel A shows that the deviation measure is negative

versus synchronized prices (99% correlation for the average deviation across contracts). The baseline measure uses the synchronized observations to account for the days when the VIX makes a large move between 4pm and 4:15pm such as on February 5, 2018.

[8]The Appendix shows that similar results hold when comparing the prices of VIX futures to their no-arbitrage lower bound. The decline in the deviation measure when risk increases thus reflects the cheapening of VIX futures relative to index option prices and not just a widening of the no-arbitrage bounds.

on average with a bias that increases to around -1.5% for the longer-dated contracts. The negative bias is consistent with the definition of the deviation measure which is the futures price minus its option-implied upper bound. In terms of variability, the standard deviation of around 1% across contracts is the same order of magnitude as the bias. The result is that all contracts exhibit law of one price violations, with more frequent upper bound violations for the shorter-dated contracts that exhibit a smaller bias.

The summary statistics also highlight the persistence of the deviation measure. In the time-series, Panel A shows that the deviations are positively autocorrelated within contract across daily, weekly, and monthly horizons. In the cross-section, Panel B shows that the deviations are positively correlated across contracts. Panel C reports the average deviation by contract over time and reveals that the decline in the average deviation from the time-series plot in Figure 1 is driven by the longer-dated contracts rather than the front contract.

3.6 Frequency of Law of One Price Violations

Table 3 investigates the law of one price violations in more detail. Panel A.I reports the fraction of days when the futures price exceeds its no-arbitrage upper bound by contract. This occurs for 12% of contract-date observations from January 3, 2007 to December 31, 2018. The upper bound violations are most frequent for the front-month and second-month contracts which exhibit positive values of $Deviation_{t,n}$ on 31% and 16% of days respectively. Panel B.I reports the analogous results for lower bound violations. Lower bound violations occur on 25% of contract-date observations with more frequent violations 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 when VIX futures were less liquid? Panels A.II and B.II show that the answer is no. In a post-crisis sample from 2010 to 2018, the number of upper bound violations declines slightly from 12% to 9%, but the number of lower bound violations increases from 25% to 30%. Figure 4 illustrates this point graphically for the front-month contract. The deviation measure in the top plot reveals frequent law of one price violations throughout the sample period. In the bottom plot the VIX futures price and variance swap forward rate appear to track each other closely throughout the sample, but large deviations are observed after differencing. The results indicate that the arbitrage violations do not go away after VIX futures liquidity improved or after traders learned of the no-arbitrage relationships. Instead, they are pervasive throughout the sample.

In addition, the table demonstrates that many of the no-arbitrage violations are large in magnitude. Panel A.I shows that the futures price is above the upper bound by more than .25% (.50%) for 6% (3%) of observations. Panel B.I shows that the futures price is below the lower bound by more than .25% (.50%) for 16% (10%) of observations. Similar to before, Panels A.II and B.II show that large upper bound violations decrease slightly in the post-crisis period, while large lower bound violations increase slightly. Combined with the summary statistics, the results indicate that there are large and persistent law of one price deviations across the VIX futures and index option markets.

4 VIX Futures Return Predictability

This section shows that the no-arbitrage deviation measure is highly significant at predicting VIX futures returns in excess of synthetic variance swap forward and index option straddle and stock market returns. A trading strategy based on the deviation measure earns economically significant alpha-to-margin across a range of assumptions. The results are consistent with the interpretation that the deviation measure is identifying an arbitrage spread across the VIX futures and index options markets.

4.1 VIX Futures and Variance Swap Forward Returns

To investigate the return predictability of the deviation measure, define the excess return from selling the n-th VIX futures contract over horizon h as,

$$R_{t+h,n}^{Fut} = Fut_{t,n} - Fut_{t+h,n}, \tag{6}$$

where $Fut_{t,n}$ is the settlement value for the n-th contract on date t. Similarly, define the excess return from receiving fixed in a synthetic variance swap forward for the n-th futures contract as,

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

The variance swap forward associated with the n-th futures contract has a floating leg that pays realized variance from the futures expiration date to the index options expiration thirty calendar days later as in the definition of the deviation measure.

Table 4 reports summary statistics for the prices and returns of VIX futures and synthetic variance swap forwards from January 3, 2007 to December 31, 2018. Panel A shows that the unconditional term-structure of implied volatility is upward sloping on average for futures and forwards. The declining standard deviation across contracts reflects the mean reversion of implied volatility. The prices are positively skewed and exhibit excess kurtosis. Panel B shows that the average return for being exposed to implied volatility shocks through VIX futures and variance swap forwards is positive but not statistically significant. These results are similar to the findings in Dew-Becker et al. (2017) and Andries et al. (2015) who find a larger risk premium for being exposed to realized volatility in comparison to implied volatility. The returns also exhibit negative skewness and excess kurtosis which is common for volatility selling strategies.

4.2 Predicting Returns with the Deviation Measure

Table 5 reports return predictability regressions for VIX futures hedged with variance swap forwards over a weekly horizon ($h = 5$),

$$R^{Fut}_{t+h,n} - \hat{\beta}_n R^{Fwd}_{t+h,n} = \alpha_n + \gamma_n Deviation_{t,n} + \epsilon_{t+h,n} \quad (8)$$

The regression is a two-step procedure for each contract. First, the VIX futures return is regressed onto the variance swap forward return to obtain a hedge ratio $\hat{\beta}_n$. Second, the hedged return is regressed onto the deviation measure. The hedged or excess return is defined as the VIX futures return minus the variance swap forward return multiplied by the hedge ratio. The hedged return and deviation measure are z-scored or standardized in the second step for ease of interpretation. The hypothesis being tested in these regressions is that the deviation measure is identifying valuation differences across VIX futures and variance swap forwards. When the deviation measure is high (low), VIX futures are expensive (cheap) relative to variance swap forwards, so the returns from selling (buying) VIX futures and paying (receiving) fixed in variance swap forward rates should be high.

The results show that the deviation measure predicts hedged returns. Panel A considers the full sample period from 2004 to 2018. In the first stage regressions in Panel A.I variance swap forward returns are highly significant for VIX futures returns, exhibiting an average explanatory power of 67% across contracts. In the second stage regressions in Panel A.II the deviation measure significantly predicts hedged returns. A one-standard deviation increase in the deviation measure for the front (sixth) contract predicts a .23 (.25) standard deviation higher hedged return with an R^2 of 5.2% (6.0%) over a weekly horizon. The average t-statistic and R^2 across contracts are 5.7 and 6.7%, with the deviation measure significantly predicting returns for each contract. To interpret the magnitude of the predictability, the standard deviation of the hedged returns is .71 across contracts on average. Thus, a .25 standard deviation higher hedged return corresponds to an increase of $.71 \times .25 = .18$ futures points or $.18/.05 = 3.6$ bid-ask spreads. Panel A.III shows that the results are robust to including control variables such as the VIX index and realized variance over the past month to proxy for the variance risk premium.[9] If anything, including control variables increases the strength of the deviation measure as a predictor.

Are the predictability 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. The average t-statistic and R^2 of the deviation in Panel B.II are 7.5 and 13.3%, even stronger than the full sample results. The Appendix includes additional results showing that the predictability holds over daily and monthly horizons, for percentage and logarithmic returns, 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.

Expanding on the results, Table 6 shows that the predictability is robust to the hedging portfolio. In this case, VIX futures are hedged with CRSP value-weighted stock market returns and at-the-money forward, delta-hedged straddle returns for the nearest index option maturities before and after the futures contract expiration. The loadings in the first stage regression have the expected signs. The returns from selling VIX futures load positively on stock market returns, consistent with the leverage effect (Black 1976), and load positively (negatively) on straddle returns with an expiration date after (before) the futures contract expiration, which proxy for the variance swap forward return. As before, the deviation measure significantly predicts hedged returns in the second stage regressions. The results show that the predictability is not driven by the definition of variance swap forward returns or by the estimation of variance swap forward rates.

4.3 VIX Futures Drive the Return Predictability

What drives the predictability results: the VIX futures or index options market? On one hand, the large and established index options market may provide a fair-value measure for VIX futures. On the other hand, given that VIX futures are traded directly, rather than being synthesized from option portfolios across maturities, the VIX futures market may provide a fair value measure for synthetic variance swap forwards. One way to test these

[9] The control variables include the VIX index, realized variance over the past month, CRSP value-weighted stock market returns over the past week, and VIX futures trading volume for the n-th contract normalized by open interest. The coefficients on the control variables are omitted to save space and to focus on how including these variables impacts the predictability of the deviation measure.

hypotheses is by investigating the predictability of the deviation measure for VIX futures and variance swap forward returns in excess of returns from a separate, but related market. For example, if VIX futures are mispriced and synthetic variance swap forwards are accurately valued, the deviation measure should remain significant at predicting VIX futures excess returns but not variance swap forward excess returns.

Tables A.6 and A.7 in the Appendix investigate these questions for VIX futures and variance swap forward returns relative to stock market returns. The approach is analogous to Tables 5 and 6. As before, VIX futures and variance swap forward returns are first regressed onto stock market returns. The hedged returns are then regressed onto the deviation measure. In the first stage, stock market returns are highly significant with an average R^2 of 54% for VIX futures. This explanatory power is lower than in Tables 5 and 6 when variance swap forwards and index option straddles are included, but illustrates that the stock market still provides significant explanatory power by itself. In the second stage, the deviation measure predicts the returns of VIX futures and variance swap forwards with the expected sign relative to the stock market. In the post-crisis sample period, the deviation measure is statistically significant for almost all of the contracts in both markets. While the strength of the predictability is somewhat lower than before, the results suggest that the VIX futures and index options markets are both contributing to the predictability of the deviation measure.

4.4 Trading Strategy and Alpha-to-Margin Estimates

The previous section shows that the deviation measure predicts VIX futures excess returns across various regression specifications. How can this predictability be interpreted? Is it robust out-of-sample?

One way to investigate these questions is by estimating the alpha-to-margin for a trading strategy based on the deviation measure (Garleanu and Pedersen 2011). To that end, consider a relative value strategy that sells (buys) hedged VIX futures when the deviation measure exceeds a high (low) threshold. To normalize the deviation measure within contract and over time, the strategy converts the deviation measure into a rolling z-score using one-year of lagged data,

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

Similarly, the hedge ratio for each contract $\beta_{t,n}$ is computed from rolling regressions with one-year of lagged data. This approach is similar to the first stage in the return predictability regressions, but accounts for time-variation in the hedge ratios and ensures the hedge ratios are in the investor's information set.

The strategy trades when the deviation z-score exceeds a threshold, $|Z_{t,n}| \geq \mathcal{T}$. If there are multiple contracts to trade on the same day, the strategy forms an equal-weighted portfolio. The return from selling a hedged VIX future over horizon h for the n-th contract is defined as,

$$R_{t+h,n} \equiv \frac{Multiplier \cdot (Fut_{t,n} - Fut_{t+h,n}) - \beta_{t,n}^{hedge} \cdot R_{t+h,n}^{hedge}}{Margin_{t,n}} - Rf_t \cdot \frac{h}{365} - tcost. \quad (10)$$

This definition uses the maximum leverage to margin. The payoff is the change in the VIX futures price times the contract multiplier minus the hedging return net of financing and transaction costs. The financing cost is the risk-free rate times the actual number of days h over 365 calendar 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.

Figure 5 plots the performance of the trading strategy with a threshold of $\mathcal{T} = .50$ z-scores and a one-week horizon $h = 5$ that trades the front six VIX futures contracts that are hedged with variance swap forwards. The margin is the exchange-based initial margin for a VIX futures contract from the CBOE and the transaction cost is assumed to be zero for now. The performance is compared to the stock market with both return series being normalized to 10% annualized volatility. From March 2004 to December 2018, the VIX futures trading strategy earns a Sharpe Ratio (SR) of 3.0 versus .5 for CRSP value-weighted stock market returns. Across market environments and even during the financial crisis, the relative value strategy based on the deviation measure performs well and exhibits minimal drawdowns. The results indicate that the predictability of the deviation measure is economically significant and robust out-of-sample - consistent with the interpretation that the deviation measure is identifying an arbitrage spread.

How sensitive is the performance of the relative value strategy to the specification? Table 7 answers this question by varying the hedging portfolio and transaction cost assumptions. The first column (1) matches Figure 5. The weekly returns have annualized volatility of $.0139 \times \sqrt{52} = .10$ by construction and an annualized SR of $.41 \times \sqrt{52} = 2.96$ or 3.0 rounded to one significant digit. The maximum drawdown of 8.7% contrasts a maximum drawdown of around 50% for the stock market over the same period of time. The second column (2) uses S&P 500 E-mini futures as a hedge instead of variance swap forward returns.[10] Similar to the return predictability regressions, the performance declines somewhat when using the

[10] The motivation for using S&P 500 E-mini futures returns rather than CRSP value-weighted returns in the trading strategy is that the initial margin and transaction costs for trading E-mini futures are directly observable and incorporated into some specifications.

less precise hedge, but the strategy still obtains a large, annualized SR of 1.38. The returns remain positively skewed with a small maximum drawdown of 6.5%. The third column (3) adds transaction costs of \$50 and \$12.50 per trade for VIX futures and E-mini futures which further decreases the annualized SR to .62. The transaction costs are the minimum tick size times the contract multiplier, which equals the median bid-ask spread for these liquid futures contracts from recent years. The subsequent columns (4-6) repeat the analysis for a post-crisis sample and show that similar results hold. For all specifications except column (1), the strategy delivers positively skewed, fat-tailed returns.

Figure 6 illustrates these results graphically. The top plot shows the performance of the trading strategy from columns (1-3) in Table 7 versus the stock market. The bottom plot shows the robustness of the strategy to varying the number of contracts traded and the trading threshold. Even when trading fewer contracts and for different thresholds near .50 z-scores, the strategy continues to earn a large SR. Once the threshold becomes sufficiently large, the SR begins to decline as the strategy sits in cash more often and exploits fewer opportunities.

The results so far highlight how the deviation measure can be used to convert the return predictability regressions into a trading strategy that performs well as measured by its SR. But what about the magnitude of the returns? Table 8 provides a perspective on the size of the trading strategy returns by reporting alpha-to-margin estimates. The rows vary the trading strategy assumptions. The columns vary the factor model and sample period.

Since the returns are hedged, the strategy is largely uncorrelated with traditional risk factors and thus earns large alpha-to-margin. For example, column (1) in Panel A of Table 8 shows that the baseline strategy from Figure 5 has a CAPM alpha-to-margin of 5.23% per week. The returns in Table 7 are de-levered to obtain 10% annualized volatility. The unlevered mean return is 5.31%. This is close to the alpha-to-margin estimate, indicating that the maximally leveraged strategy delivers large returns that are largely unexplained by the market factor. The second (2) and third (3) columns show that this result is robust to including the Fama-French three factors and Carhart momentum factor (FFC4) and a six-factor model (FFCV6) that adds realized and implied volatility factor returns for one-month variance swaps and front-month VIX futures. The Appendix reports the factor loadings. The following columns (4-6) repeat the analysis for a post-crisis sample from 2010 to 2018. Across specifications, the alpha-to-margin remain high, ranging from 4.5% to 5.2% per week.

The subsequent panels (B-D) report the alpha-to-margin when the assumptions about the trading strategy are varied. In particular, Panel B hedges with E-mini futures. Panel C adds transaction costs, and Panel D adds the margin from the E-mini contract.[11] The CAPM

[11]The margin assumption in Panel D is conservative. If an exchange or bilateral counterparty offered

alpha-to-margin estimates in column (1) decline from 5.23% in Panel A to 3.32% in Panel B, 1.59% in Panel C, and .78% in Panel D. While the hedging portfolio, transaction costs, and margin assumptions all impact the estimates, the CAPM alpha-to-margin estimate of .78% per week in Panel D is is still economically and statistically significant. Similar results hold for the other specifications.

In summary, the results indicate that trading against the no-arbitrage deviation measure earns significant returns with minimal exposure to traditional risk factors. The large alpha-to-margin estimates and high SRs are robust across a range of specifications and sample periods. The strongest results correspond to the relative value strategy that hedges with synthetic variance swap forwards.

5 Discussion

5.1 Deviation Measure versus Risk and Demand Factors

What drives the no-arbitrage 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. In addition, when risk increases, financing constraints may bind more tightly and force arbitrageurs to exit positions. The resulting effects can be anomalies such as return predictability and no-arbitrage violations.

To investigate how the deviation measure relates to risk and demand factors, the paper estimates a vector autoregression (VAR) for $y_t = [DEV_t \, VIX_t \, DNP_t]$. The variables in the VAR are the average deviation across the front six contracts (DEV), the CBOE Volatility Index (VIX), and the dealer net position (DNP) in VIX futures from the CFTC's CoT report. The DEV variable is averaged across contracts to keep the VAR parsimonious. The VIX and DNP variables are proxies for risk and demand factors.

Figure 2 motivates using DNP as a demand variable by illustrating its high correlation with VIX ETP demand from Dong (2016), a proxy for retail demand. The advantage of DNP relative to VIX ETP demand is that it only depends on quantities, not on prices. The DNP variable is normalized as a fraction of open interest. This bounds DNP between 0 and 1 which removes the time trend in net position size that reflects the growth in the VIX

portfolio margin, the capital requirement for the hedged trade would be lower than the margin for a position in VIX futures because of the reduced risk due to the hedge. In contrast, Panel D assumes that initial margin is required for both the VIX futures and E-mini futures contracts. The average initial margin for VIX futures and E-mini futures are around $4,400 and $4,100 over the sample period so including the E-mini margin roughly halves the alpha-to-margin estimates versus Panel C.

futures market in recent years.

The sample period for the VAR analysis is 2010 to 2018 using weekly observations on CoT release dates when the DNP variable is reported. The sample period and start date are motivated by 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.[12] 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 7 presents the time-series relationship between the variables. 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 to find a relationship with 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 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 in the next section.[13]

5.2 VAR Impulse Response Functions

The paper estimates a trivariate VAR for the standardized variables $y_t = [DEV_t \, VIX_t \, DNP_t]$ at a weekly frequency from 2010 to 2018. Figure 8 reports impulse response functions (IRFs)

[12] 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.

[13] The Appendix expands on this analysis by showing 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. As such, the results are not limited to the specific variables selected here.

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.

The top row in Figure 8 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 one to two months. 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 8 reveals how the DNP demand 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 their option-implied upper bounds. 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 in Figure 8 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, when 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 synthetically 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 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. The insignificant responses are consistent with the null hypothesis of no effect on the VIX which seems most plausible given the large size of the index options market relative to the VIX futures market.

5.3 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 isolate within contract variation and broader control variables to account for the economic environment.

Table 9 reports panel regressions and finds similar relationships between the deviation measure and the risk and demand variables despite the different identification approach from the VAR. The regression specification is,

$$Deviation_{t+h,n} = \beta \cdot \Delta x_{t+h} + \rho \cdot Deviation_{t,n} + \delta \cdot Controls_t + \text{FEs} + \epsilon_{t+h,n}. \qquad (11)$$

The horizon is one-week $h = 5$ as in 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. Since the standard deviation of the deviation variable is around 1% on average, this makes the results roughly comparable to the IRFs from the VAR where the variables are standardized. The first three columns (1-3) in Table 9 show how the deviation measure responds to changes in the explanatory variable controlling for the persistence of the deviation measure. The next three columns (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 contract number (i.e. first, second, etc. 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 -.12 in Panel A or -.16 in Panel B. In the VAR, the IRF of the deviation measure to a VIX shock has a point estimate of around -.05 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 the VIX. 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, the deviation measure is increasing in demand pressure. A one-standard deviation increase in dealer net position (DNP) 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 (ETP) and in the delta of VIX options traded by retail customers (VixOpt). Similar to the DNP variable, the VIX ETP and VIX 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.

5.4 Other Channels

Beyond risk and demand, the limits-to-arbitrage literature highlights a number of events 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 dollar funding pressure that is associated with heightened arbitrage deviations in foreign exchange markets (Du et al. 2018). Salient events like FOMC announcements feature surprising return predictability patterns (Lucca and Moench 2015). How does the deviation measure relate to these events?

Figure 9 investigates this question by reporting event study plots to illustrate how the average deviation across the front six contracts and VIX index change around various events. The figure reports the average change with a 95% pointwise confidence interval for the deviation measure for the ten trading days before and after the events.

The top row plots the results for margin changes. The initial margin for the front month VIX futures contract has 41 increases and 33 decreases from 2004 to 2018. Before a margin increase (decrease) the VIX tends to increase (decrease) by around 4 (1) points. Based on the panel regressions, the expected change in the deviation measure given the change in the VIX is a decrease (increase) of around .5 (.1). For the margin increases, the deviation measure only declines by around .25, which is about half the expected change. Before the margin decrease, the deviation measure increases by around .20 which is somewhat more than expected. This provides some evidence that margin increases (decreases) attenuate (amplify) the relationship between the deviation measure and the VIX.

The subsequent plots highlight the change in the deviation measure around FOMC and non-farm payroll announcements and around month-end and quarter-end dates. The results reveal a decrease in the deviation measure following the FOMC announcement and an increase leading into non-farm payrolls. This suggests that investors may reduce VIX futures hedges after the FOMC and increase hedges ahead of non-farm payrolls. There is also a decline in the deviation measure in the week prior to month-end and quarter-end, which could be consistent with hedgers deleveraging and reducing positions during this period.

5.5 Deviation Measure versus VIX premium

How does the VIX futures no-arbitrage deviation measure relate to the VIX premium from Cheng (2018)? The VIX premium measures the difference between the one-month futures price and a statistical forecast of the VIX index at maturity.[14] It can be interpreted as the implied volatility risk premium or the expected return from going short $1 notional of a VIX futures contract over a one-month horizon. In contrast, the deviation measure is a relative pricing or no-arbitrage discrepancy across the VIX futures and index options markets.

Figure 10 plots the VIX premium against the weighted average of the deviation measure across the front two contracts for a constant, one-month maturity. The VIX premium data is from Ing-Haw Cheng's website and is presented for the expanding window model. The deviation measure is computed as the weighted average deviation across the front two

[14] The VIX premium is constructed by rolling from the front contract to the second contract at the end of the month. The definition adjusts for the number of days to maturity to scale the premium to a one-month horizon.

contracts for a one-month maturity.[15] The plot highlights a strong time-series correlation between the VIX premium and the deviation measures, showing that, despite the different interpretations and estimation techniques, the measures are positively correlated.

Table 10 provides further analysis by regressing the VIX premium onto the deviation measure, the VIX index, and realized variance (RV). As before, the variables are standardized in each of the regressions. In column (1) the deviation measure explains around 11% of the VIX premium and is statistically significant. A one-standard deviation increase in the deviation measure is associated with a .34 standard deviation increase in the VIX premium. Column (2) shows that the VIX index and realized variance are also related to the VIX premium, with explanatory power of around 25%. Column (3) adds all of the variables together and finds that the deviation measure remains significant. Columns (4-6) repeat the analysis during a post-crisis sample. In this period, the deviation measure has higher explanatory power than the VIX and RV and remains significant. These results are consistent with the strong time-series correlation from Figure 10, showing the similarity in the measures despite their different interpretations.

6 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, revealing the presence of arbitrage opportunities. A deviation measure equal to the difference between the VIX futures price and its corresponding variance swap forward rate is found to significantly predict VIX futures excess returns. A relative value trading strategy that exploits the deviation measure earns a significant Sharpe ratio with minimal exposure to traditional risk factors.

The results are surprising because the no-arbitrage relationships investigated in the paper are well known to option traders. Even if the arbitrage trade is difficult to implement in practice, there is still significant return predictability when using the deviation measure to predict VIX futures returns hedged with S&P 500 futures, both of which are liquid futures contracts. The trading strategy obtains a sizeable Sharpe ratio and alpha-to-margin estimate under a range of assumptions. 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 sys-

[15] The one-month deviation measure is $Deviation_{t,1} \cdot w_t + Deviation_{t,2} \cdot (1 - w_t)$ so $\tau_{1,t} \cdot w_t + \tau_{2,t} \cdot (1 - w_t) = 1/12$. Early in the sample there are some dates when the front contract has a maturity greater than one-month in which case the weight is set to $w_t = 1$.

tematic risk and demand pressure as well as other channels like margin changes, economic announcements, and quarter-end effects. 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 risk and risk premia.

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.

Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the 1976 Meetings of the American Statistical Association*, 171–181.

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 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.

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.

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.

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.

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., 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. (2020). Equity volatility term premia. *Working Paper*.

Table 1: Estimating the Law of One Price Deviation Measure

This table provides examples for how the paper estimates the law of one price deviation measure on February 27, 2012 and February 8, 2018. First, the paper estimates the variance swap forward curve by assuming flat forward rates between observed index option maturities. Second, the paper computes the deviation measure, $Deviation_{t,n}$, for the n-month futures contract as the difference between the futures price, $Fut_{t,n}$, and the corresponding variance swap forward rate, $\sqrt{Fwd_{t,n}}$, in volatility units. The variance swap forward rate is the average forward rate between the futures expiration date and the index options expiration date thirty calendar days later.

Estimating the Deviation Measure on February 27, 2012

Variance swap forward Curve, annualized variance in percentage units

$Fwd_{t,T_{n-1},T_n} = (VS_{t,T_n} \cdot T_n - VS_{t,T_{n-1}} \cdot T_{n-1})/(T_n - T_{n-1})$

Option Expiration	3/16/12	4/20/12	5/18/12	6/15/12	9/21/12	12/21/12
Maturity (Years)	0.05	0.15	0.22	0.30	0.57	0.82
VS_{t,T_n}	2.99	3.74	4.51	5.15	6.21	6.72
Fwd_{t,T_{n-1},T_n}	2.99	4.12	5.97	7.01	7.38	7.88

Deviation Measure, annualized volatility in percentage units

$Deviation_{t,n} = Fut_{t,n} - \sqrt{Fwd_{t,n}}$

Futures Expiration	3/21/12	4/18/12	5/16/12	6/20/12	7/18/12	8/22/12
Futures Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
$Fut_{t,n}$	21.40	24.38	25.87	27.02	28.17	28.78
$\sqrt{Fwd_{t,n}}$	20.30	24.17	26.34	27.17	27.17	27.17
$Deviation_{t,n}$	1.10	0.20	-0.46	-0.15	1.00	1.60

Estimating the Deviation Measure on February 8, 2018

Variance swap forward Curve, annualized variance in percentage units

$Fwd_{t,T_{n-1},T_n} = (VS_{t,T_n} \cdot T_n - VS_{t,T_{n-1}} \cdot T_{n-1})/(T_n - T_{n-1})$

Option Expiration	3/16/18	4/20/18	5/18/18	6/15/18	9/21/18	12/21/18
Maturity (Years)	0.10	0.19	0.27	0.35	0.62	0.87
VS_{t,T_n}	11.94	9.26	8.31	7.84	6.40	6.18
Fwd_{t,T_{n-1},T_n}	11.94	6.51	5.90	6.16	4.53	5.63

Deviation Measure, annualized volatility in percentage units

$Deviation_{t,n} = Fut_{t,n} - \sqrt{Fwd_{t,n}}$

Futures Expiration	2/14/18	3/21/18	4/18/18	5/16/18	6/20/18	7/18/18
Futures Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
$Fut_{t,n}$	28.88	22.83	21.08	20.12	19.27	19.15
$\sqrt{Fwd_{t,n}}$	34.56	25.51	24.37	24.79	21.30	21.30
$Deviation_{t,n}$	-5.68	-2.68	-3.30	-4.67	-2.02	-2.15

Table 2: Summary Statistics for the Law of One Price Deviation Measure

Panel A reports summary statistics for the law of one price deviation measure for the front six VIX futures contracts. Panel B reports the correlation of the deviation measure across contracts. Panel C reports the average deviation for each contract by year. The last column reports the average statistics across contracts. The sample period is January 3, 2007 to December 31, 2018 for a balanced panel.

Summary Statistics of $Deviation_{t,n} = Fut_{t,n}$ - $\sqrt{Fwd_{t,n}}$

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)	Avg.
Panel A: Summary Statistics							
Mean	-0.38	-0.79	-0.98	-1.50	-1.44	-1.41	-1.08
Standard Deviation	0.98	0.90	1.10	1.00	0.95	1.03	0.99
t-statistic	-5.62	-10.38	-10.68	-17.32	-15.70	-13.94	-12.28
Skewness	-4.69	-1.21	0.38	-1.24	-0.53	-0.40	-1.28
Kurtosis	44.01	9.86	9.50	9.24	5.00	4.41	13.67
Minimum	-12.55	-6.99	-7.30	-9.45	-7.78	-6.13	-8.37
25th-Percentile	-0.62	-1.14	-1.57	-2.02	-1.97	-1.99	-1.55
Median	-0.25	-0.69	-0.97	-1.46	-1.41	-1.39	-1.03
75th-Percentile	0.09	-0.27	-0.41	-0.86	-0.83	-0.74	-0.50
Maximum	3.83	5.52	6.60	2.76	1.74	3.42	3.98
Autocorrelation 1-day	0.65	0.76	0.76	0.82	0.85	0.86	0.78
Autocorrelation 5-day	0.40	0.58	0.62	0.67	0.72	0.77	0.63
Autocorrelation 21-day	0.26	0.36	0.31	0.33	0.47	0.47	0.37
Panel B: Correlation Matrix							
Deviation$_{t,1}$	1.00	0.26	0.06	0.20	0.20	0.18	0.18
Deviation$_{t,2}$	0.26	1.00	0.23	0.55	0.49	0.49	0.40
Deviation$_{t,3}$	0.06	0.23	1.00	0.29	0.28	0.39	0.25
Deviation$_{t,4}$	0.20	0.55	0.29	1.00	0.58	0.44	0.41
Deviation$_{t,5}$	0.20	0.49	0.28	0.58	1.00	0.67	0.44
Deviation$_{t,6}$	0.18	0.49	0.39	0.44	0.67	1.00	0.43
Panel C: Average Annually							
2007	-0.46	-0.27	-0.53	-0.46	-0.33	-0.29	-0.39
2008	-1.09	-0.59	-0.20	-1.30	-1.10	-1.16	-0.91
2009	-0.31	-0.82	-0.42	-1.73	-1.81	-1.76	-1.14
2010	-0.32	-0.46	-0.45	-0.98	-1.16	-1.27	-0.77
2011	-0.42	-1.05	-1.49	-2.12	-2.06	-2.12	-1.54
2012	-0.12	-0.57	-0.88	-1.16	-0.89	-0.64	-0.71
2013	-0.07	-0.43	-0.57	-0.95	-0.89	-0.85	-0.63
2014	-0.22	-0.71	-1.02	-1.68	-1.49	-1.49	-1.10
2015	-0.54	-1.43	-1.71	-1.96	-1.97	-2.12	-1.62
2016	-0.28	-1.08	-1.46	-1.85	-1.79	-1.73	-1.36
2017	-0.16	-0.90	-1.42	-1.90	-1.83	-1.67	-1.31
2018	-0.55	-1.13	-1.69	-1.96	-1.88	-1.84	-1.51

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 for varying thresholds for sample periods from 2007-2018 in A.I and from 2010-2018 in A.II. Panel B reports the analogous frequency of lower bound violations. The last column reports the average frequency across contracts.

Frequency of Law of One Price Violations

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)	Avg.
Panel A.I: Futures $>$ Upper Bound + Threshold from 2007-2018							
Threshold = 0	0.31	0.16	0.12	0.03	0.04	0.06	0.12
Threshold = .25	0.14	0.08	0.07	0.02	0.02	0.04	0.06
Threshold = .50	0.06	0.03	0.04	0.01	0.01	0.02	0.03
Panel A.II: Futures $>$ Upper Bound + Threshold from 2010-2018							
Threshold = 0	0.32	0.11	0.06	0.01	0.02	0.03	0.09
Threshold = .25	0.14	0.05	0.02	0.01	0.01	0.02	0.04
Threshold = .50	0.06	0.02	0.01	0.00	0.01	0.02	0.02
Panel B.I: Futures $<$ Lower Bound - Threshold from 2007-2018							
Threshold = 0	0.12	0.17	0.29	0.37	0.29	0.24	0.25
Threshold = -.25	0.07	0.11	0.19	0.25	0.20	0.15	0.16
Threshold = -.50	0.04	0.07	0.11	0.17	0.13	0.10	0.10
Panel B.II: Futures $<$ Lower Bound - Threshold from 2010-2018							
Threshold = 0	0.11	0.20	0.36	0.47	0.37	0.30	0.30
Threshold = -.25	0.06	0.12	0.24	0.32	0.25	0.19	0.20
Threshold = -.50	0.03	0.07	0.14	0.21	0.16	0.12	0.12

Table 4: Summary Statistics for the Prices and Returns of VIX Futures and Variance Swap Forwards

Panel A reports summary statistics for VIX futures prices and for their corresponding variance swap forward rates. Panel B reports summary statistics for one-week excess returns for VIX futures and variance swap forwards ($h = 5$). The sample period is January 3, 2007 to December 31, 2018.

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	21.31	21.65	21.95	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	17.55	18.58	19.27	19.67	20.08	20.27	19.24
Panel A.II: VS Forward Rates $\sqrt{Fwd_{t,n}}$ (annualized volatility units)							
Mean	20.45	21.60	22.29	23.15	23.38	23.62	22.41
Standard Deviation	8.99	7.90	7.03	6.76	6.50	6.32	7.25
Skewness	2.28	1.92	1.71	1.47	1.32	1.25	1.66
Kurtosis	9.94	7.66	6.84	5.45	4.77	4.56	6.54
Median	17.97	19.56	20.67	21.42	21.70	21.90	20.54
Panel B.I: VIX Futures Returns (percent) $R^{Fut}_{t+h,n} = Fut_{t,n} - Fut_{t+h,n}$							
Mean	0.17	0.15	0.09	0.06	0.06	0.05	0.10
Standard Deviation	2.43	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	1.22	1.47
Skewness	-1.56	-1.20	-1.03	-0.86	-0.80	-0.74	-1.03
Kurtosis	12.76	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) $R^{Fwd}_{t+h,n} = Fwd_{t,n} - Fwd_{t+h,n}$							
Mean	0.48	0.47	0.39	0.54	0.05	0.13	0.34
Standard Deviation	16.31	12.90	10.27	8.47	7.29	6.64	10.31
Sharpe Ratio	0.03	0.04	0.04	0.06	0.01	0.02	0.03
t-statistic	0.78	0.89	1.06	1.69	0.20	0.56	0.86
Skewness	-3.98	-2.92	-1.73	-0.60	-1.27	-1.05	-1.92
Kurtosis	65.65	45.73	27.82	16.22	18.70	16.38	31.75
Median	0.93	1.15	0.68	0.61	0.26	0.30	0.66

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

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

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

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample Period from 2004 to 2018						
Panel A.I: First Stage - Hedge Ratios						
β_n	12.72***	13.26***	12.15***	12.90***	13.55***	13.56***
	(1.43)	(1.26)	(1.02)	(0.72)	(0.71)	(0.87)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.680	0.727	0.618	0.654	0.682	0.672
Panel A.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.23***	0.25***	0.30***	0.28***	0.25***	0.25***
	(0.05)	(0.04)	(0.05)	(0.04)	(0.04)	(0.04)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.052	0.063	0.092	0.076	0.061	0.060
Panel A.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.32***	0.34***	0.34***	0.32***	0.31***	0.28***
	(0.04)	(0.04)	(0.06)	(0.04)	(0.04)	(0.04)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.073	0.106	0.120	0.091	0.089	0.075
Panel B: Post-Crisis Sample from 2010 to 2018						
Panel B.I: First Stage - Hedge Ratios						
β_n	19.46***	20.34***	18.49***	15.49***	16.15***	15.05***
	(1.34)	(0.79)	(0.70)	(0.75)	(0.66)	(0.57)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.775	0.823	0.757	0.711	0.722	0.725
Panel B.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.47***	0.35***	0.38***	0.32***	0.35***	0.30***
	(0.08)	(0.04)	(0.04)	(0.05)	(0.04)	(0.05)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.225	0.123	0.143	0.101	0.119	0.087
Panel B.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.56***	0.42***	0.42***	0.33***	0.35***	0.29***
	(0.09)	(0.05)	(0.05)	(0.05)	(0.04)	(0.05)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.240	0.158	0.159	0.106	0.126	0.090

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

Table 6: Deviation Measure Predicts VIX Futures Returns Hedged with Stock Market and Index Option Straddle Returns

This table reports return predictability regressions for VIX futures hedged with stock market and straddle returns over a weekly horizon, $h = 5$. The hedging portfolio includes CRSP value-weighted returns and at-the-money forward delta-hedged straddle returns for the nearest index option maturities before and after the futures contract expiration.

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample Period from 2004 to 2018						
Panel A.I: First Stage - Hedge Ratios						
$RMRF_{t+h}$	0.58***	0.44***	0.32***	0.26***	0.22***	0.20***
	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
$R^{Strad}_{t+h,T_{n+1}}$	8.68***	13.68***	21.99***	17.44***	25.64***	29.79***
	(0.82)	(1.17)	(1.68)	(5.20)	(2.88)	(2.42)
R^{Strad}_{t+h,T_n}	-1.61***	-2.58***	-6.90***	-3.67	-8.87***	-12.19***
	(0.62)	(0.47)	(0.81)	(3.24)	(2.04)	(1.60)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.694	0.751	0.765	0.741	0.770	0.759
Panel A.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.17**	0.14***	0.09**	0.12*	0.14***	0.15***
	(0.08)	(0.05)	(0.04)	(0.07)	(0.04)	(0.04)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.030	0.020	0.008	0.015	0.020	0.021
Panel A.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.24***	0.12***	0.08**	0.09*	0.10***	0.10***
	(0.08)	(0.04)	(0.04)	(0.05)	(0.04)	(0.04)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.053	0.052	0.046	0.038	0.042	0.045
Panel B: Post-Crisis Sample from 2010 to 2018						
Panel B.I: First Stage - Hedge Ratios						
$RMRF_{t+h}$	0.65***	0.53***	0.41***	0.32***	0.26***	0.24***
	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)
$R^{Strad}_{t+h,T_{n+1}}$	7.67***	11.47***	19.80***	22.97***	27.12***	26.47***
	(0.91)	(0.67)	(1.30)	(1.78)	(2.16)	(1.91)
R^{Strad}_{t+h,T_n}	-0.69	-1.70***	-6.17***	-8.10***	-10.88***	-10.34***
	(0.48)	(0.33)	(0.69)	(1.04)	(1.44)	(1.25)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.698	0.821	0.849	0.842	0.832	0.826
Panel B.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.05	0.14**	0.16***	0.13***	0.13***	0.09**
	(0.06)	(0.06)	(0.05)	(0.05)	(0.05)	(0.05)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.002	0.019	0.024	0.016	0.017	0.008
Panel B.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.13**	0.14**	0.11**	0.09**	0.10**	0.05
	(0.06)	(0.06)	(0.05)	(0.04)	(0.04)	(0.04)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.036	0.036	0.035	0.035	0.033	0.029

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

Table 7: Summary Statistics for Deviation-Based Trading Strategy

This table reports summary statistics for weekly returns from the deviation-based trading strategy that are levered for 10% annualized volatility. The columns vary the specifications which include hedging with variance swap forwards (1), E-mini S&P 500 futures (2), and E-mini futures including transaction costs (3). Columns (4-6) report analogous results for a 2010-2018 post-crisis sample.

Weekly Return Summary Statistics

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Mean	0.57	0.27	0.12	0.58	0.23	0.09
Standard Deviation	1.39	1.39	1.39	1.39	1.39	1.39
Skewness	-1.12	1.46	1.43	2.32	1.60	1.56
Kurtosis	54.46	14.02	14.01	28.11	16.42	16.37
Sharpe Ratio	0.41	0.19	0.09	0.42	0.17	0.06
t-statistic	14.85	5.89	2.65	12.39	4.34	1.63
Maximum Drawdown	8.74	6.54	15.30	4.00	6.36	12.83
Hedge	Fwd	E-mini	E-mini	Fwd	E-mini	E-mini
Transaction Costs	No	No	Yes	No	No	Yes
Post-Crisis	No	No	No	Yes	Yes	Yes
Observations	3697	3697	3697	2245	2245	2245

Table 8: Alpha-to-Margin Estimates for Deviation-Based Trading Strategy

This table reports weekly alpha-to-margin estimates for the deviation-based trading strategy using the maximum leverage to required initial margin. Panels A-D vary the hedge and whether transaction costs and the E-mini margin are included. Columns 1-3 vary the factor model. Columns 4-6 report results for a 2010-2018 post-crisis sample. The alpha-to-margin estimates are large and significant across the various specifications. The Appendix reports the factor loadings.

Weekly Alpha-to-Margin Estimates

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Variance swap forward hedge						
Alpha	5.23***	5.23***	5.00***	4.68***	4.72***	4.55***
	(0.37)	(0.37)	(0.42)	(0.39)	(0.39)	(0.43)
Adjusted R^2	0.009	0.009	0.022	0.005	0.007	0.014
Panel B: E-mini hedge						
Alpha	3.32***	3.35***	3.89***	2.68***	2.76***	3.19***
	(0.52)	(0.52)	(0.49)	(0.60)	(0.61)	(0.61)
Adjusted R^2	0.027	0.029	0.110	0.010	0.026	0.102
Panel C: E-mini hedge and t-costs						
Alpha	1.59***	1.62***	2.15***	1.11*	1.19**	1.60***
	(0.51)	(0.51)	(0.49)	(0.60)	(0.60)	(0.61)
Adjusted R^2	0.026	0.029	0.106	0.010	0.026	0.098
Panel D: E-mini hedge, t-costs, and margin						
Alpha	0.78***	0.78***	1.08***	0.52*	0.56*	0.77***
	(0.26)	(0.26)	(0.24)	(0.29)	(0.29)	(0.30)
Adjusted R^2	0.026	0.028	0.110	0.008	0.025	0.086
Factor Model	CAPM	FFC4	FFCV6	CAPM	FFC4	FFCV6
Post-Crisis	No	No	No	Yes	Yes	Yes
Observations	3697	3697	3697	2245	2245	2245

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

Table 9: 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 January 3, 2007 to December 31, 2018

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Explanatory Variable	RMRF	RV	VIX	RMRF	RV	VIX
Δx_{t+h}	0.10***	-0.08***	-0.12***	0.11***	-0.22***	-0.15***
	(0.03)	(0.02)	(0.02)	(0.02)	(0.05)	(0.02)
$\text{Deviation}_{t,n}$	0.73***	0.72***	0.74***	0.49***	0.50***	0.49***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Observations	17683	17683	17683	17683	17683	17683
Adjusted R^2	0.544	0.541	0.549	0.604	0.602	0.612
Controls and FEs	No	No	No	Yes	Yes	Yes

Panel B: Risk Factors from January 4, 2010 to December 31, 2018

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Explanatory Variable	RMRF	RV	VIX	RMRF	RV	VIX
Δx_{t+h}	0.17***	-0.11***	-0.16***	0.18***	-0.36***	-0.19***
	(0.02)	(0.03)	(0.02)	(0.02)	(0.05)	(0.02)
$\text{Deviation}_{t,n}$	0.79***	0.77***	0.80***	0.51***	0.50***	0.50***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Observations	13254	13254	13254	13254	13254	13254
Adjusted R^2	0.635	0.618	0.635	0.695	0.690	0.699
Controls and FEs	No	No	No	Yes	Yes	Yes

Panel C: Demand Factors from January 4, 2010 to December 31, 2018

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Explanatory Variable	DNP	ETP	VixOpt	DNP	ETP	VixOpt
Δx_{t+h}	0.05***	0.02**	0.07***	0.05***	0.03**	0.08***
	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.02)
$\text{Deviation}_{t,n}$	0.79***	0.79***	0.78***	0.51***	0.51***	0.51***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Observations	13254	13254	13254	13254	13254	13254
Adjusted R^2	0.614	0.613	0.617	0.675	0.674	0.677
Controls and FEs	No	No	No	Yes	Yes	Yes

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

Table 10: Deviation Measure versus VIX premium

This table compares the deviation measure to the VIX premium from Cheng (2018). The deviation measure is the weighted average of the deviation across the front two contracts for a constant, one-month maturity to match the maturity of the VIX premium. The table regresses the VIX premium onto the one-month deviation measure, VIX, and realized variance.

VIX premium Regressed onto Deviation, VIX, and RV

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Deviation 1mn	0.34***		0.33***	0.39***		0.46***
	(0.09)		(0.06)	(0.05)		(0.06)
VIX		0.68***	0.80***		0.43***	0.54***
		(0.12)	(0.12)		(0.15)	(0.11)
RV		-0.93***	-0.87***		-0.84***	-0.54***
		(0.17)	(0.18)		(0.19)	(0.15)
Observations	3697	3697	3697	2245	2245	2245
Adjusted R^2	0.115	0.251	0.331	0.168	0.057	0.244
Sample Period	Full	Full	Full	Post	Post	Post

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

Figure 1: Law of One Price Deviation between VIX Futures and Index Option Prices

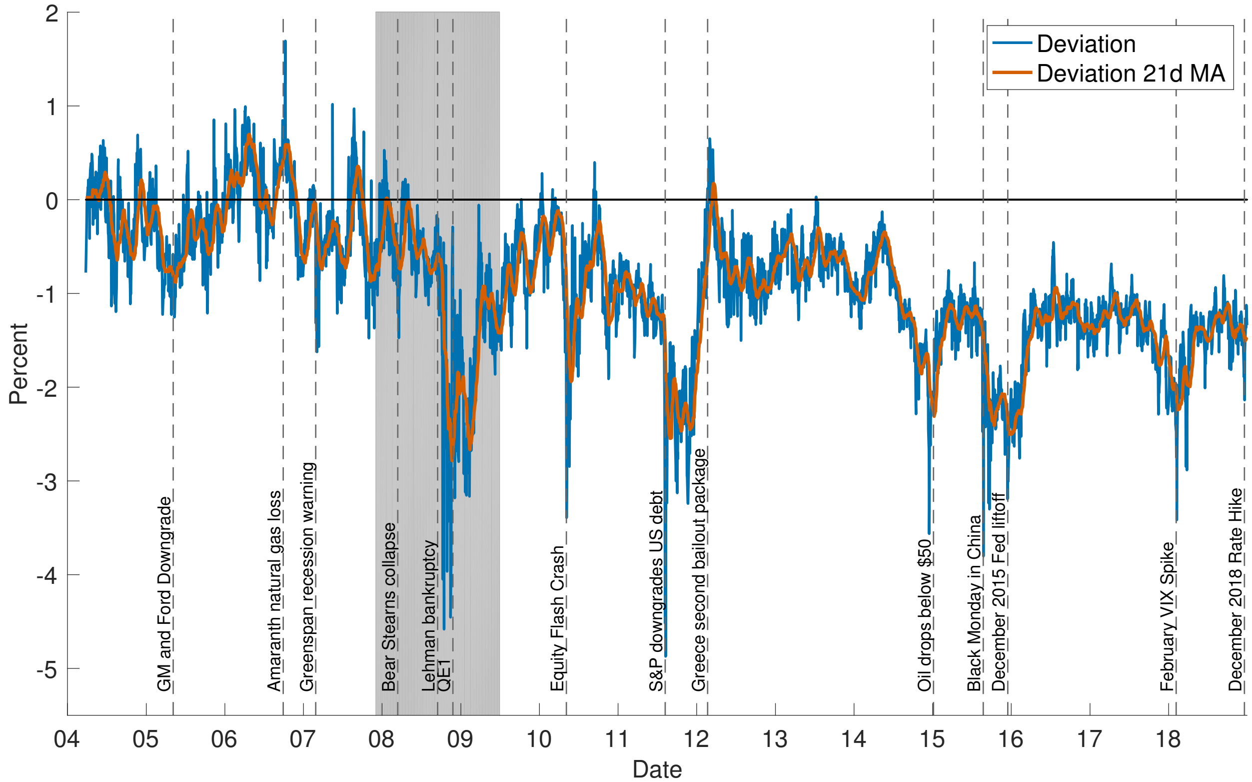

This figure plots the deviation measure averaged across the front six VIX futures contracts as a daily time-series and a one-month moving average from March 26, 2004 to December 31, 2018. The no-arbitrage deviation measure for the n-month VIX futures contract is the difference between the futures price and its corresponding variance swap forward rate, $Deviation_{t,n} = Fut_{t,n} - \sqrt{Fwd_{t,n}}$. Positive values are law of one price violations. Gray shading indicates an NBER recession.

Figure 2: Equity Volatility Market Size and Dealer Positions

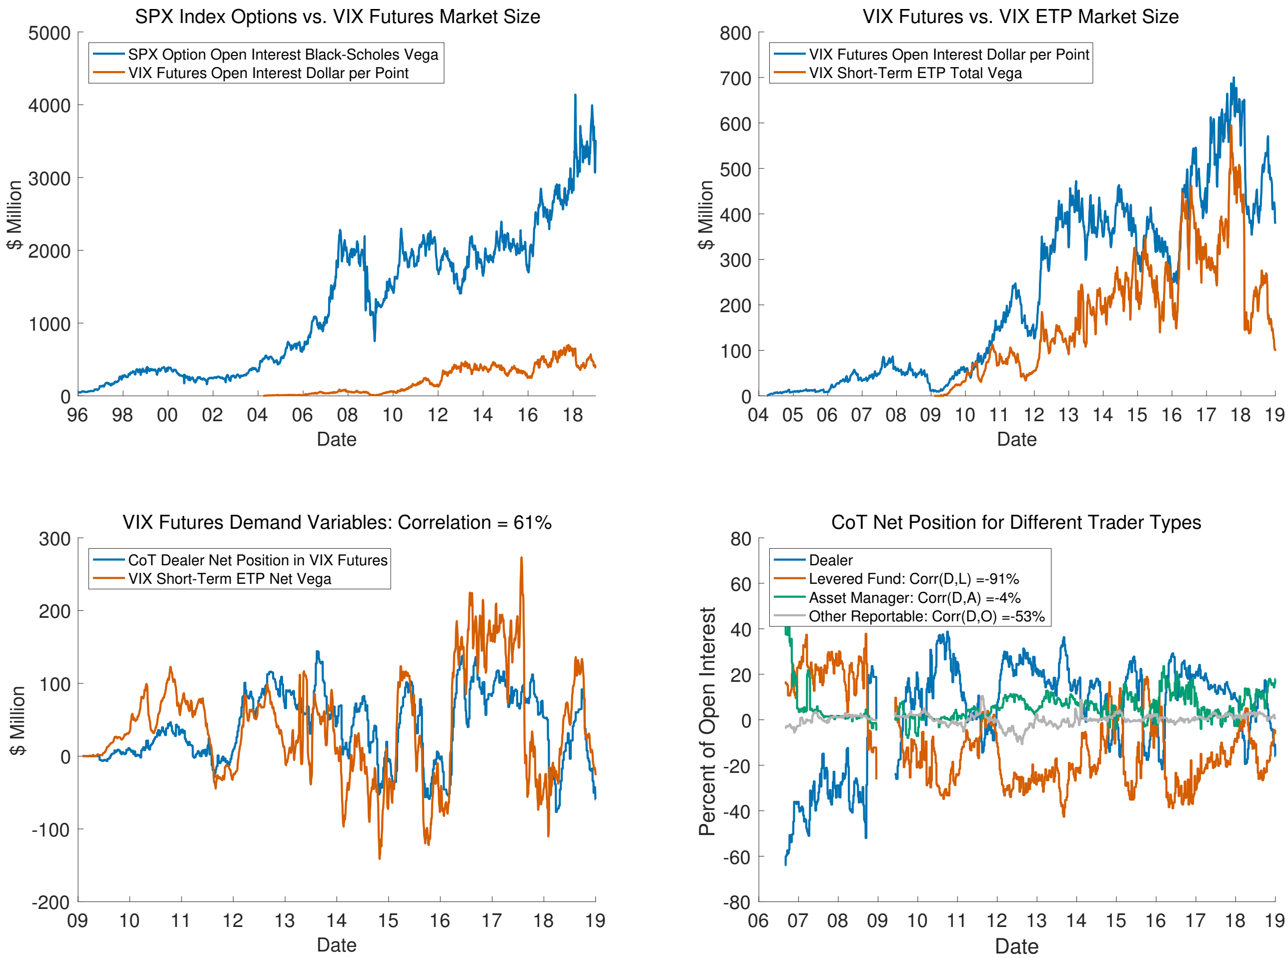

This figure provides an overview of equity volatility markets including the market size and investor positioning in the S&P 500 index options and VIX futures markets. The index options market is much larger than the VIX futures market as measured by open interest in Black-Scholes vega (top left). Since 2009, there has been substantial growth in the VIX futures and ETP markets which has coincided (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.

Figure 3: Estimating the Law of One Price Deviation Measure

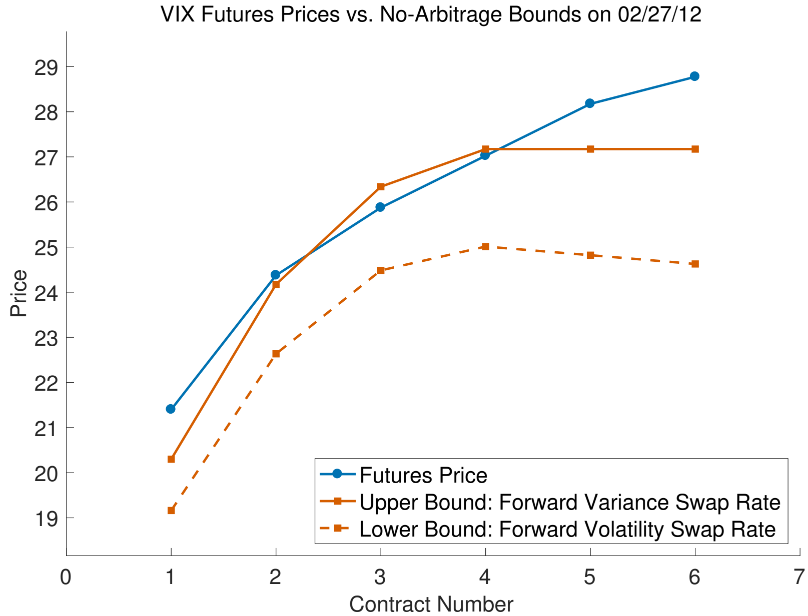

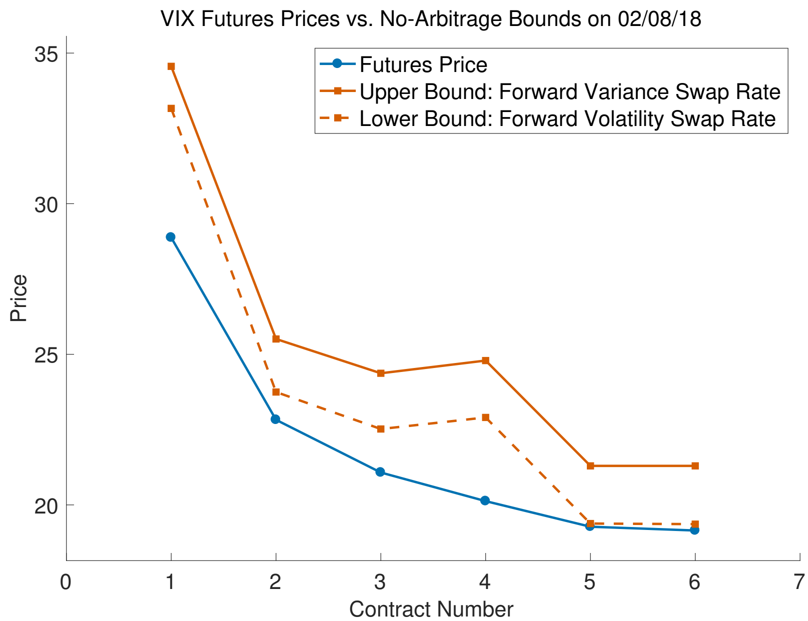

This figure plots the prices of VIX futures and variance swap forwards from Table 1 that are used to construct the law of one price deviation measure on February 27, 2012 and February 8, 2018. The top plot illustrates static arbitrage opportunities in which the prices of VIX futures are above their upper bounds for the first, fifth, and sixth contracts. In these cases, an arbitrageur could sell VIX futures and buy variance swap forwards (pay fixed) to lock in a riskless profit. The bottom plot illustrates arbitrage opportunities in which the prices of VIX futures are below their lower bounds. The lower bound is equal to the upper bound minus the difference in bounds from a dynamic term-structure model.

Figure 4: Law of One Price Deviation for the Front Month VIX Futures Contract

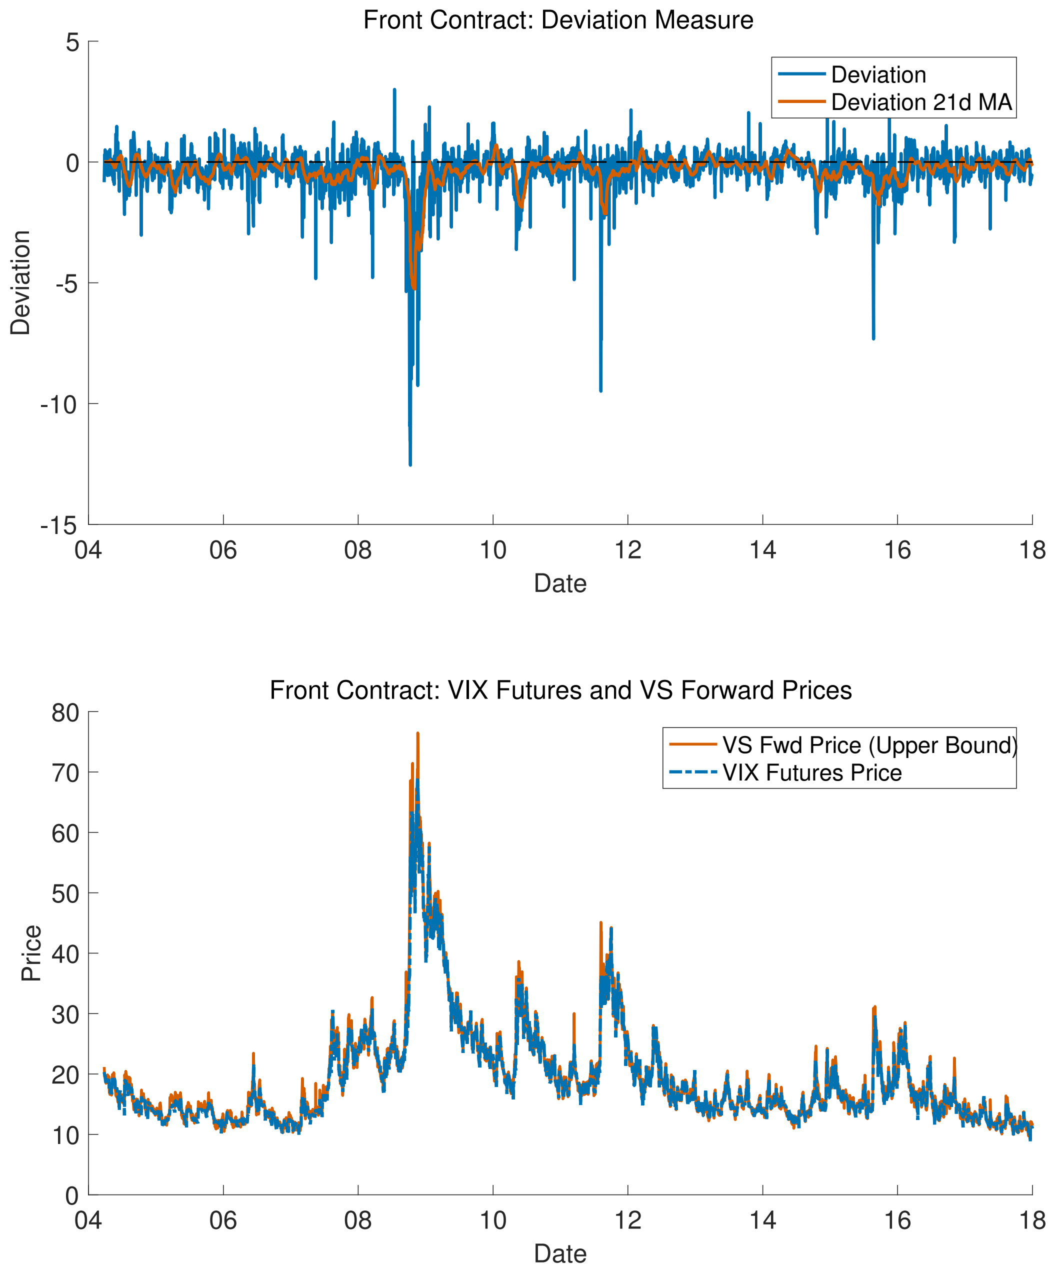

The top figure plots the deviation measure for the front month VIX futures contract. The bottom figure plots the VIX futures price and variance swap forward rate that are used to compute the deviation measure. While the futures price and forward price tend to track each other closely in the bottom plot, there are periods with prolonged and significant law of one price deviations as evidenced by the top plot.

Figure 5: Performance of Deviation-Based Trading Strategy

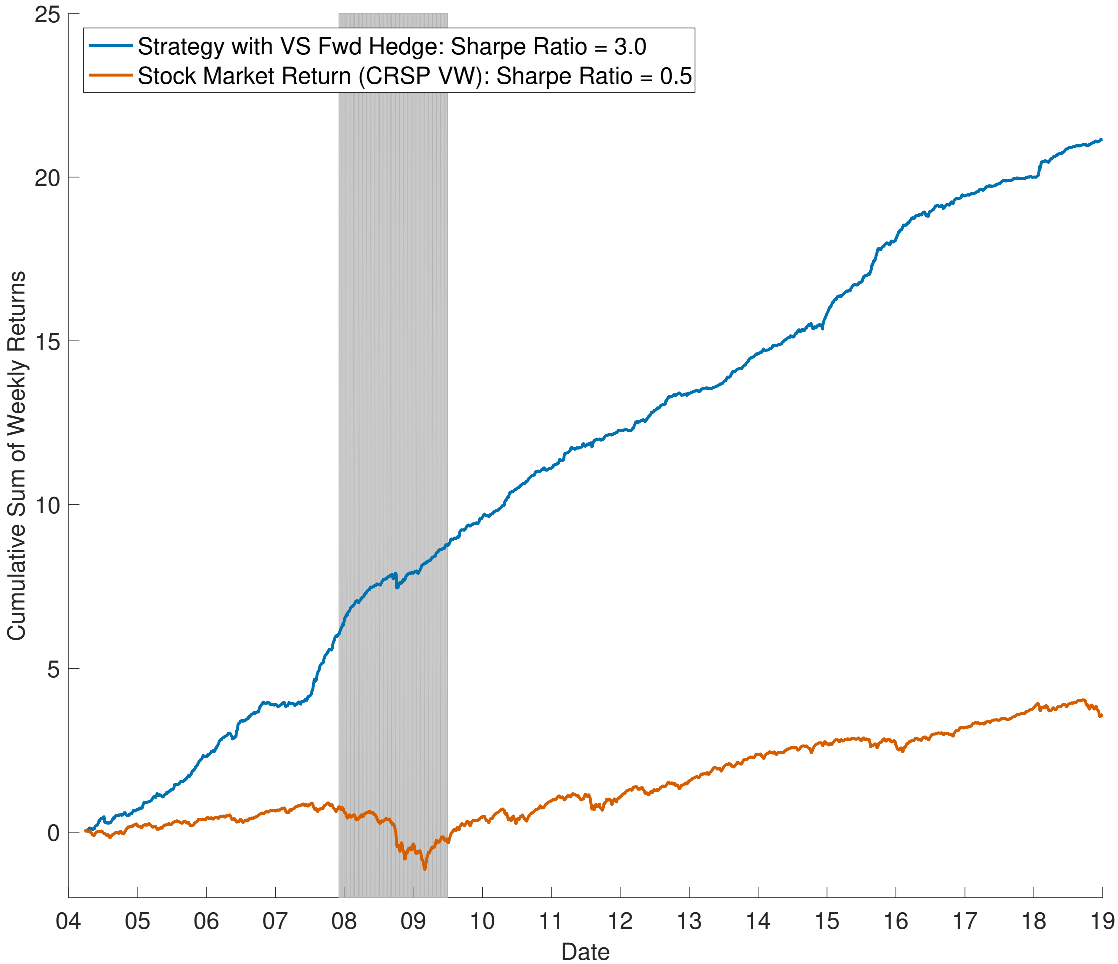

This figure plots the performance of the relative value trading strategy in VIX futures and variance swap forwards against stock market returns. The relative value strategy goes long (short) VIX futures when the deviation measure exceeds a low (high) threshold of $\mathcal{T} = .50$ z-scores for the 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 6: Robustness of Deviation-Based Trading Strategy

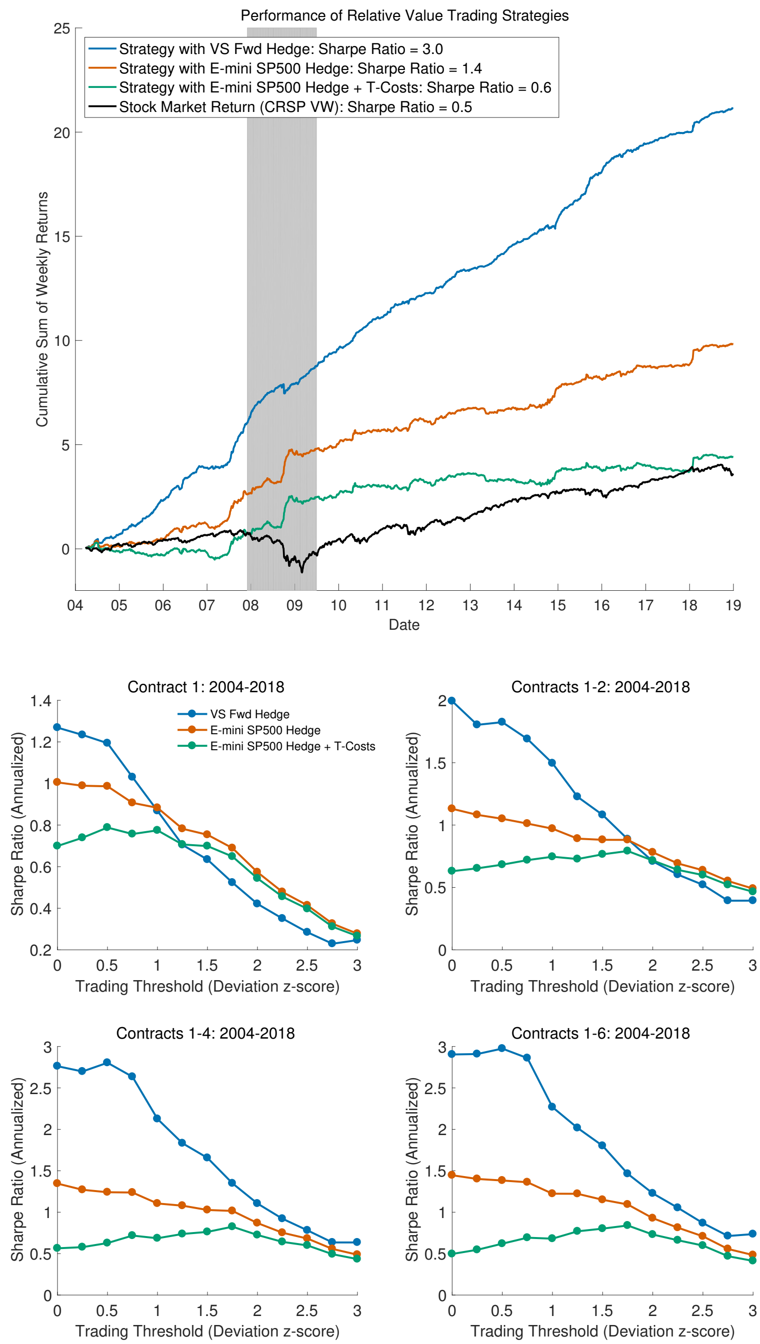

This figure illustrates the robustness of the relative value trading strategy to hedging with stock market returns, including transaction costs, trading different numbers of contracts, and varying the trading threshold. The top plot reports the time-series of cumulative returns for the different strategies 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 7: Time-Series of Deviation, VIX, and Dealer Net Position for VAR

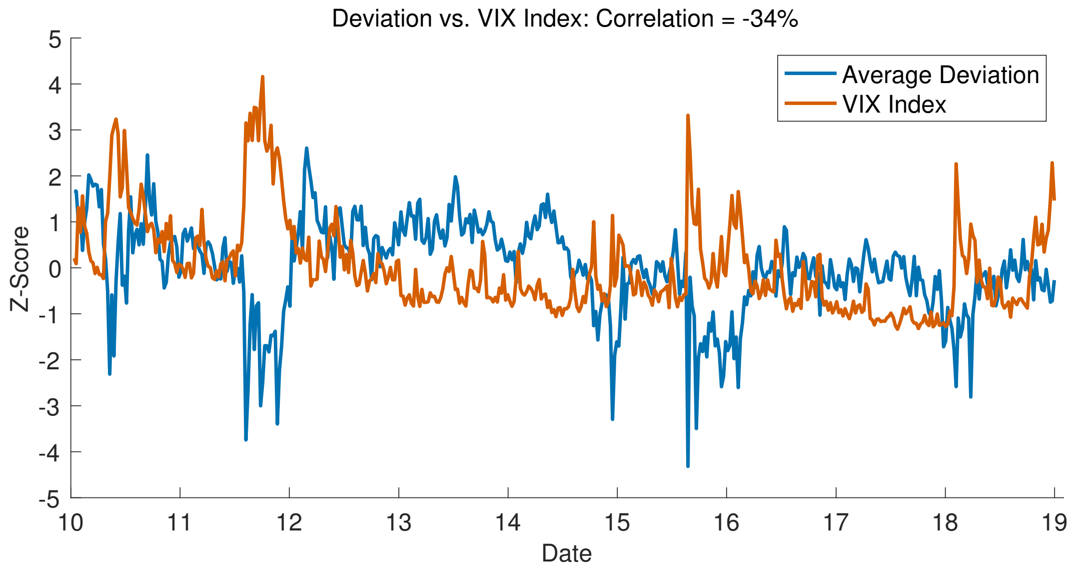

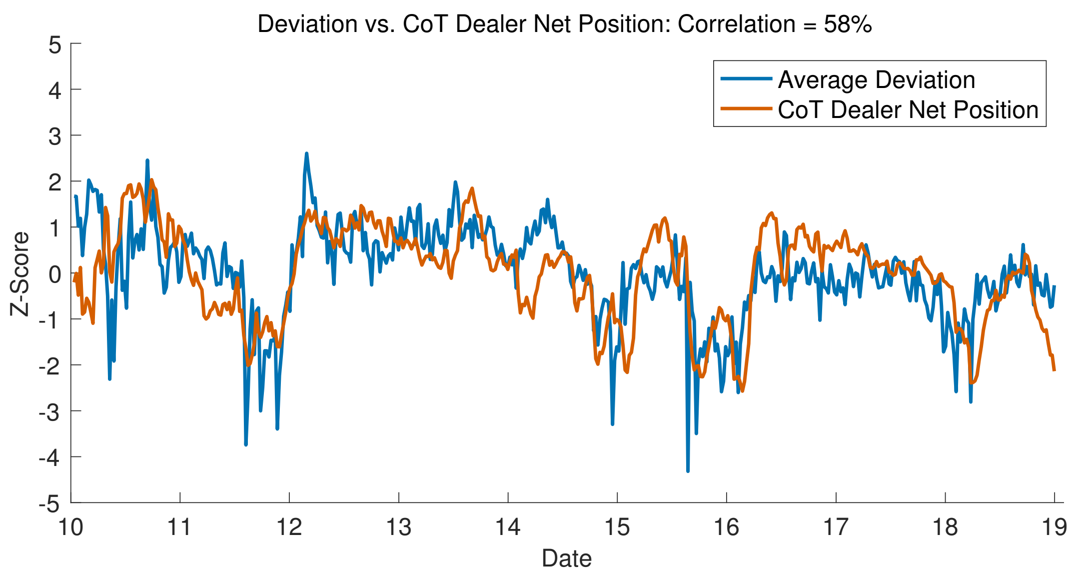

This figure plots the deviation measure averaged across the front six VIX futures contracts against the VIX index and dealer net position at a weekly frequency from 2010 to 2018. The VAR is estimated with these variables which are standardized for comparison. The deviation measure is negatively correlated 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 8: Impulse Response Functions of the No-Arbitrage Deviation Measure and Dealer Net Position to Shocks in the VIX index and Dealer Net Position

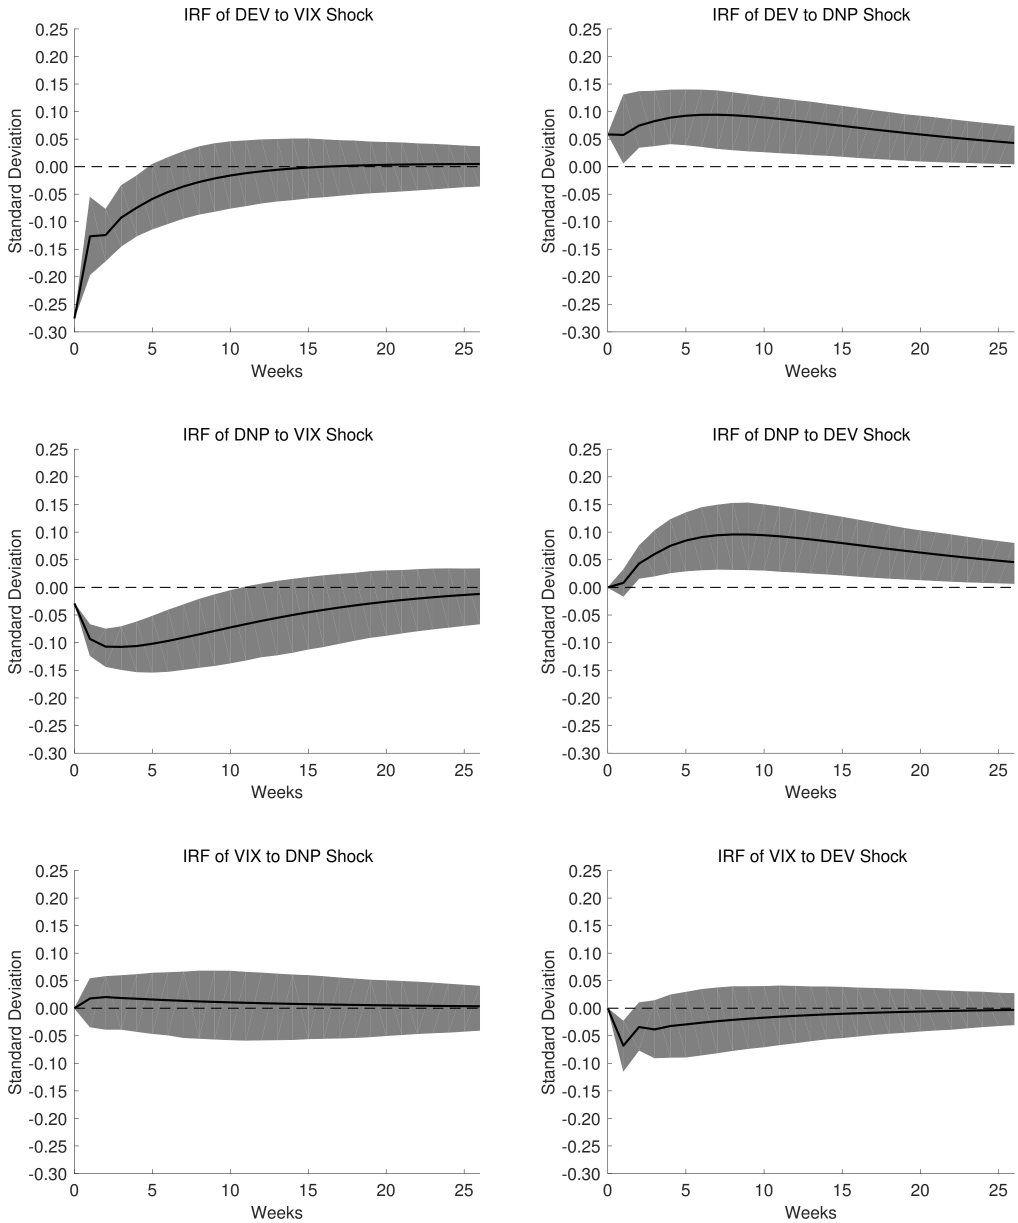

This figure reports the impulse response functions (IRFs) from the VAR. The IRFs are identified from a Cholesky decomposition with the ordering: VIX, DNP, DEV. The VAR is estimated using weekly data from 2010 to 2018 ($T = 469$). The IRFs in the first row show that DEV is decreasing (increasing) in VIX (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 9: Event Study of Deviation Measure Around VIX Futures Margin Changes, FOMC days, Nonfarm Payrolls, Month-End, and Quarter-End Dates

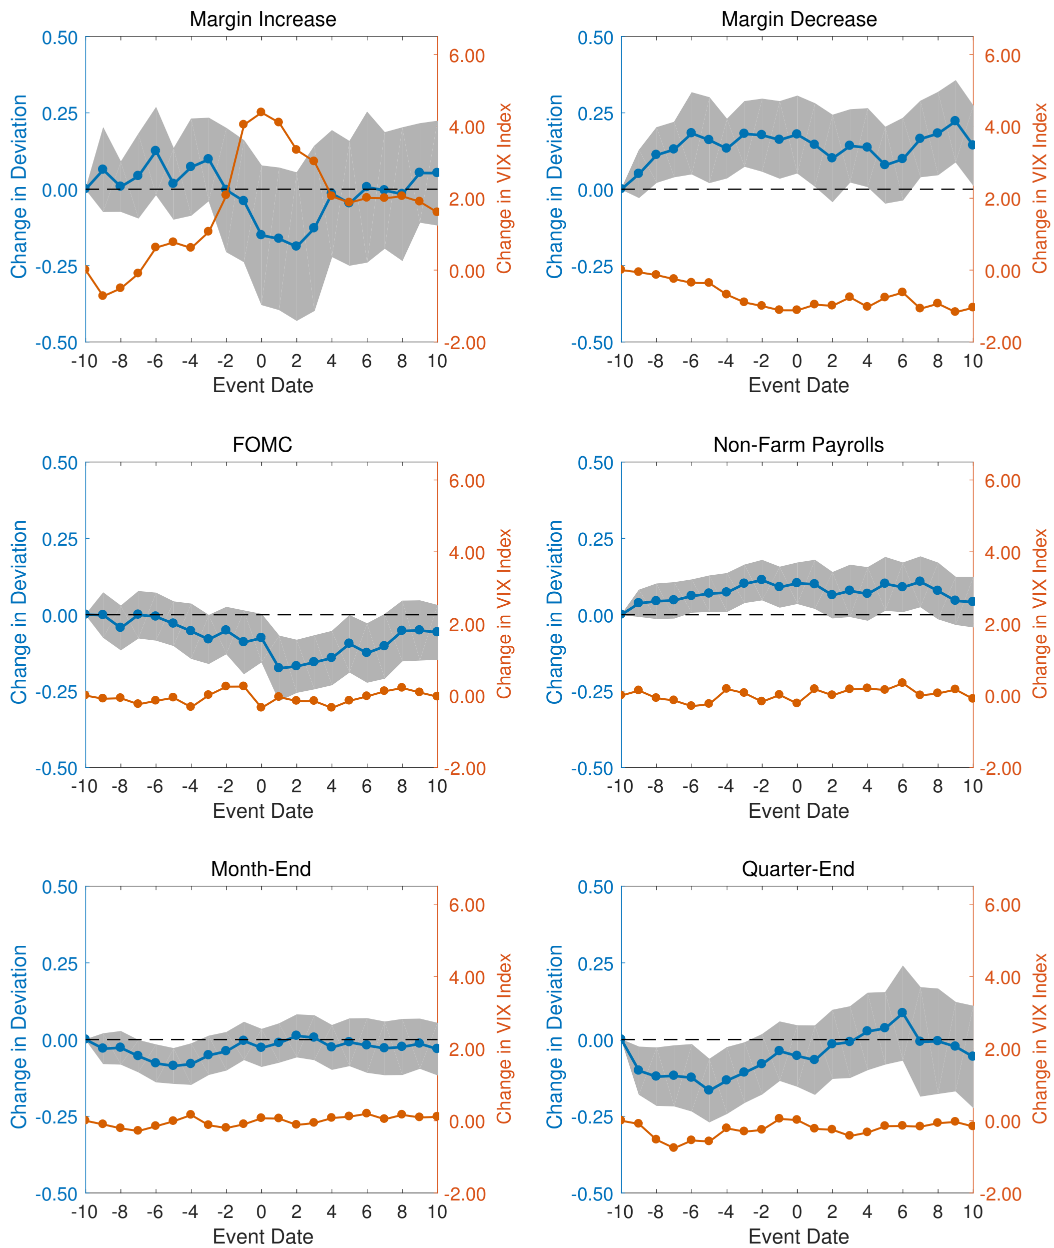

This figure plots the event time reaction of the deviation measure and VIX index to changes in the initial margin for the front month VIX futures contract, FOMC announcements, non-farm payrolls announcements, end-of-month dates, and end-of-quarter dates. The figure reports the change in the VIX index 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.

Figure 10: Deviation Measure versus VIX premium

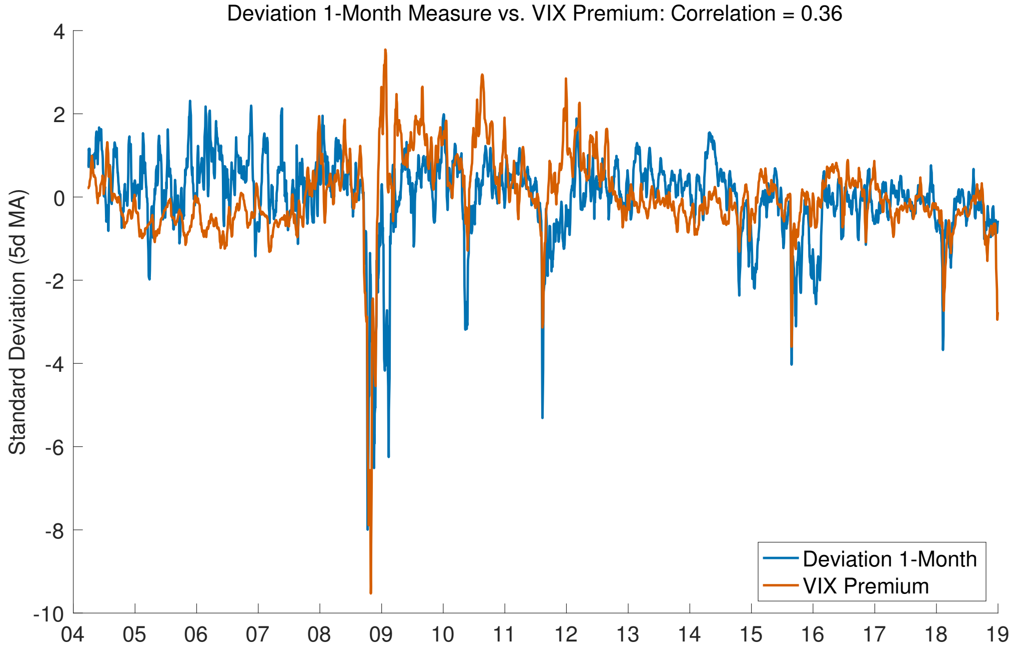

This figure plots the VIX premium from Cheng (2018) against the weighted average of the deviation measure across the front two contracts for a constant, one-month maturity. The measures exhibit a positive and significant time-series correlation.

Appendix:

"The Law of One Price in Equity Volatility Markets"

Peter Van Tassel

December 2020

A Appendix

A.1 Arbitrage Hedge Ratio and VIX Futures Lower Bound

Consider an arbitrage opportunity in which the price of a VIX futures contract exceeds the corresponding variance swap forward rate $Fut_{t,T} > Fwd_{t,T,T+1}$. The payoff function $\Pi_T(VIX_T)$ for the arbitrage trade can be decomposed as,

$$\begin{aligned} \Pi_T(VIX_T) &= VIX_T^2 - Fwd_{t,T,T+1} - \omega \cdot (VIX_T - Fut_{t,T}) \\ &= VIX_T^2 - Fut_{t,T}^2 + (Fut_{t,T}^2 - Fwd_{t,T,T+1}) - \omega \cdot (VIX_T - Fut_{t,T}) \\ &= \underbrace{(Fut_{t,T}^2 - Fwd_{t,T,T+1})}_{\text{Arbitrage Violation}} + \underbrace{g(VIX_T) - g(Fut_{t,T}) - g'(Fut_{t,T})(VIX_T - Fut_{t,T})}_{\text{Taylor series approximation when hedge ratio } \omega = g'(Fut_{t,T})}. \end{aligned} \tag{12}$$

The first term reflects the arbitrage violation and is thus non-negative by assumption. Setting the hedge ratio to $\omega = g'(Fut_{t,T}) = 2 \cdot Fut_{t,T}$, the second term is non-negative because it is the residual from a first-order Taylor series approximation of a convex function $g(x) = x^2$ around the futures price. This shows that the hedge ratio $\omega = 2 \cdot Fut_{t,T}$ ensures $\Pi_T \geq 0 \; \forall VIX_T$.

The lower bound for the price of a VIX futures contract is a volatility swap forward rate. To see this, let variance swaps be modeled as the risk-neutral expected value of realized variance,

$$VS_{t,T} = E_t^{\mathbb{Q}}[RV_{t,T}]. \tag{13}$$

The VIX is defined to closely approximate the square-root of a one-month variance swap $VIX_t \equiv \sqrt{VS_{t,t+1}}$. It follows by Jensen's inequality that,

$$\begin{aligned} Fut_{t,T} &= E_t^{\mathbb{Q}}[VIX_T] \\ &= E_t^{\mathbb{Q}}\left[\sqrt{VS_{T,T+1}}\right] \\ &= E_t^{\mathbb{Q}}\left[\sqrt{E_T^{\mathbb{Q}}[RV_{T,T+1}]}\right] \\ &\geq E_t^{\mathbb{Q}}\left[E_T^{\mathbb{Q}}\left[\sqrt{RV_{T,T+1}}\right]\right] \\ &= E_t^{\mathbb{Q}}\left[\sqrt{RV_{T,T+1}}\right] \\ &\equiv Fvol_{t,T}, \end{aligned} \tag{14}$$

where the last line defines a volatility swap forward rate. Thus, VIX futures prices are bounded above by variance swap forward rates and below by volatility swap forward rates.

A.2 No-Arbitrage Deviation Measure Estimation

The deviation measure has two components: the VIX futures price and the variance swap forward rate. The baseline measure in the paper takes the futures price to be the settlement price until March 3, 2008 and then the mid-quote at 4pm ET from TRTH data to be synchronous with the option quotes from OptionMetrics. The forward rate is computed from synthetic variance swap rates by assuming flat forward rates between index option maturities. The synthetic variance swap rates are estimated following the approach in Van Tassel (2020) for traditional expirations on the third Friday of the calendar month with at least two weeks to maturity.

Is the deviation measure robust to this specification? Table A.1 answers this question by reporting the correlation of the baseline deviation measure with alternative measures in levels and in monthly and weekly changes. The alternative measures vary the estimation method and data sources for computing the deviation measure. Across specifications, the baseline deviation measure from the paper is found to be highly correlated with the alternative measures. The advantage of the baseline measure is its availability throughout the full sample period.

The alternative measures in Panel A consider different ways of removing the bias from the upper bound. Panel A.I uses the residual from a regression of the deviation measure onto a constant and time-to-maturity. While the regression-adjusted measure has lower bias, it is almost perfectly correlated with the baseline measure in the time-series dimension. Panel A.II uses a term-structure model from Van Tassel (2020) to remove the bias. In this case the deviation measure is defined as,

$$Devation_{t,n}^{ts} \equiv Deviation_{t,n} + (UB_{t,n}^{model} - Fut_{t,n}^{model}), \tag{15}$$

where $UB_{t,n}^{model}$ and $Fut_{t,n}^{model}$ are the upper bound and futures price from the estimated term-structure model. As with the regression-adjusted measure, the term-structure adjusted measure has a correlation of 95% to 99% with the baseline measure in levels and changes. These high correlations indicate that the time-series properties of the bias-adjusted measures are similar to those of the baseline measure. The implication is that similar results hold when using these measures for return predictability or for constructing trading strategies, as will be shown shortly.

Panel B considers alternative estimation methods and data sources. Panel B.I linearly interpolates synthetic variance swap rates at index option maturities onto a fixed monthly grid of maturities. The interpolation is linear in total variance following Carr and Wu (2009). The forward rate is then computed from the interpolated rates. Panel B.II uses VIX settlement prices for the futures price rather than using mid-quotes after March 3, 2008. Panel B.III uses Bloomberg synthetic variance swap rates which are available on a fixed monthly grid. Panel B.IV uses the CBOE 1, 3, and 6-month indices and Bloomberg data for longer maturities. Panel B.V uses over-the-counter variance swap quotes from Markit which are available at a monthly frequency. The table indicates the available sample period for each data source.

Across the different measures, the correlation for the average deviation across contracts is often 90% or greater. Figures A.1 and A.2 illustrate this result. In the top plot in Figure A.1 from February 27, 2012, the futures price is above the forward rate for the fifth and sixth contracts for all of the measures and for the front contract for all measures except the CBOE VIX indices. Similarly, in the bottom plot from February 8, 2018, the futures prices are below the forward rates and an estimate of the lower bound across specifications.

The top plot in Figure A.2 shows that the alternative specifications are more than 95% correlated with the baseline measure as a five-day moving average. The bottom plot reports three time-series that correspond to the baseline deviation measure (futures price minus upper bound), the deviation measure that is bias-adjusted with the term-structure model, and the futures price minus a lower bound. These measures are also highly correlated and show that the variation in the deviation measure is not driven by changes in the size of the no-arbitrage bounds.

A.3 VIX Futures Return Predictability Robustness

The next set of tables illustrates the robustness of the return predictability regressions to the deviation measure specification, forecast horizon, and return definitions. Table A.2 repeats the return predictability regression from Panel A.II in Table 5 for different deviation measures. Whether the deviation measure is lagged, computed as a five-day moving average, bias-adjusted, computed from linearly interpolated variance swap rates, or estimated from different datasets over different sample periods, the no-arbitrage deviation remains significant at predicting VIX futures returns hedged with variance swap forward rates. Throughout these regressions, only the predictor variable on the right-hand side changes. The hedged return on the left-hand side remains the same.

Table A.3 shows that the predictability holds over daily, weekly, and monthly horizons for the full sample period and for a post-crisis sample from 2010 to 2018. The explanatory power increases with the forecast horizon for the longer-dated contracts as measured by the R2. These regressions require at least two days to maturity and use the hold-to-maturity return if the contract expires before the forecast horizon for the first and second contract.

Tables A.4 and A.5 are analogous to Table 5 in the paper. They show that the return predictability results continue to hold when using percentage and logarithmic returns. In the paper, the motivation for defining returns as the change in prices is that this more closely corresponds to the payoff of a futures contract or a forward swap on a fixed notional, and thus is more directly comparable to the relative value trading strategy. Obtaining percentage or logarithmic returns would require a trading strategy that dynamically adjusts the position size to account for the level of the futures and forward prices. For the sake of simplicity, the return predictability regressions and trading strategy avoid this complication in the paper.

A.4 VIX Futures vs. Variance Swap Forward Return Predictability

Tables A.6 and A.7 report return predictability regressions for VIX futures and variance swap forwards in excess of CRSP value-weighted stock market returns. The approach is analogous to Table 5 in the paper and speaks to the question raised in the paper: which market drives the return predictability of the deviation measure, the VIX futures market or the index options market? The results indicate that the deviation measure is significant at forecasting VIX futures and variance swap forward excess returns relative to the stock market, particularly in the post-crisis sample period. The sign on the deviation measure is positive for VIX futures and negative for variance swap forwards as expected. For example, a high deviation measure indicates that variance swap forwards are cheap relative to VIX futures, predicting low returns from receiving fixed in variance swap forwards, hence the negative sign in Table A.7. The results provide evidence that both VIX futures and synthetic variance swap forwards implied by the index options market are contributing to the predictability of the deviation measure.

A.5 Alpha-to-Margin Loadings and Trading Strategy Position

Table A.8 reports the factor loadings from the alpha-to-margin regressions in Table 8. While there are some significant factors, the CAPM and FFC4 models only explain the trading strategy returns with an R2 of 1%-3%. The realized and implied volatility factors increase the explanatory power somewhat in Panels B and C to 10%-11% with the negative loadings on the volatility factors increasing the alpha-to-margin by around 50 basis points. Except for this change, the primarily insignificant factor loadings and low R2 highlight how the hedge in the trading strategy removes most of the systematic risk stemming from traditional risk factors.

Figure A.3 presents additional results for the deviation-based trading strategy. The top plot shows the net position in VIX futures as a five-day moving average. A long position corresponds to buying VIX futures that are hedged with variance swap forwards. The net position is the number of long contracts minus the number of short contracts. The results are presented for the strategy that hedges with variance swap forwards, uses a threshold of $\mathcal{T} = .50$ z-scores, and trades the front six contracts. The strategy goes long (short) hedged VIX futures when VIX futures look cheap (expensive) relative to variance swap forwards. This generates a negative correlation with the deviation measure as reported in the plot. The bottom plot repeats the analysis from Figure 6 for a 2010 to 2018 sample. As before, the strategy obtains a high Sharpe Ratio for thresholds around .5 z-scores and when trading fewer than six contracts, showing that the results continue to hold post-crisis.

A.6 Deviation Measure by Contract

Figure A.4 reports the time-series of the deviation measure alongside the futures price and forward rate by contract. This plot breaks out the average deviation from Figure 1 into the contributions from each contract and shows when the fifth and sixth contracts enter the sample. As the plot makes clear, violations of the upper bound occur for all contracts but are most pervasive for the front contract, consistent with the summary statistics in Table 3. The plot also highlights how the futures price and forward rate seem to track each other well in the left plot before computing the deviation measures in the right plot. The lower average deviation for the longer-dated contracts is consistent with Table 2.

Figures A.5 and A.6 build on Table 2 by reporting histograms of the futures price minus the no-arbitrage upper and lower bounds. Figure A.5 reports histograms of the deviation measure, or the futures price minus the upper bound. While the distributions are negatively skewed, the positive mass for each contract indicates the presence of arbitrage opportunities that match the frequency estimates in Panel A.I from Table 2 for a threshold of zero. Similarly, Figure A.6 reports histograms of the futures price minus the lower bound in which the negative mass indicates the presence of arbitrage opportunities, matching Panel B.I from Table 2 for a threshold of zero.

A.7 Risk and Demand Factors and VAR Robustness

Figure A.7 plots the average deviation measure against the risk and demand factors. The data is weekly from 2010 to 2018 to match the release of the CoT report ($T = 469$). The risk factors include stock market returns, realized variance, and the VIX index. The stock market return is the CRSP value-weighted return over the past month. The realized variance estimate is the realized variance of the S&P 500 Index over the past twenty-one trading days using one-minute TRTH high frequency data and a two-scale approach as in Van Tassel

(2020). The demand factors include the dealer net position in VIX futures from the CoT report, VIX futures ETP demand from Dong (2016) as a proxy for retail demand, and the net delta from customer VIX options trades over the past month.[16] The demand variables are normalized by VIX futures open interest to account for the growth in the VIX futures market over the sample period. The variables in the plot are standardized for comparison and the time-series correlation with the deviation measure is included in the subplot titles.

The plots in Figure A.7 reveal a similar pattern to Figure 7 in the paper. An increase in risk as measured by negative stock market returns or higher volatility is associated with a decline in the deviation measure. An increase in demand for VIX futures is associated with an increase in the deviation measure.

Figure A.8 expands on these results by reporting IRFs from bivariate VARs that are estimated with the deviation measure and each of the different risk and demand variables. Similar to the time-series pots, the IRFs show that an increase in risk is associated with a decline in the deviation measure that reverts after a few weeks. An increase in demand is associated with a more persistent increase in the deviation measure for shocks to the dealer net position and VIX ETP demand variables. The impact of a shock to the VIX options customer delta variable is positive and reverts more quickly.

Figure A.9 illustrates the robustness of the IRFs from the paper that describe how the deviation measure responds to VIX and dealer net position shocks. The top row repeats the ordering from the paper. The second and third row vary the ordering and the final row uses a spectral decomposition. The qualitative results are similar across the specifications. Overall, these results show that the findings in the paper are robust to using different risk and demand variables and various VAR specifications.

[16] The VIX ETP demand variable is computed from Bloomberg data following Dong (2016). It equals the leverage-weighted market capitalization of short-term ETPs net of short interest, $D_{ETP} = \sum_{i \in ST}(Shrout_i - ShortInt_i) \cdot P_i \cdot M_i$, where $Shrout_i$ is shares outstanding, $ShortInt_i$ is short interest, P_i is price, and M_i is the direction and leverage multiplier. The variable D_{ETP} includes data for the VXX, VIIX, VIXY, UVXY, TVIX, XIV, SVXY, IVOP, XXV, and VXXB ETPs. The multipliers are equal to $M = [1\ 1\ 1\ 2\ 2\ \text{-}1\ \text{-}1\ \text{-}1\ \text{-}1\ 1]$. The multipliers for UVXY and SVXY change to 1.5 and -.5 after February 28, 2018. The paper divides D_{ETP} by the VIX index to express the variable in \$ million of vega. Similarly, the total vega in Figure 2 is $\sum_{i \in ST} Shrout_i \cdot P_i \cdot |M_i|/VIX$. The data for constructing the VIX options customer demand variable is from the CBOE open-close data merged with Black-Scholes-Merton Greeks from OptionMetrics. This data is used in Cheng (2018) as a demand variable for VIX futures contracts and is described in Garleanu et al. (2009).

Table A.1: Correlation of Deviation Measure with Alternative Specifications

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 Estimates

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	1.00	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	0.98	0.97	0.96	0.96	0.95
Correlation in monthly changes	0.99	0.99	0.99	0.99	0.99	0.99	0.98
Correlation in weekly changes	0.99	0.99	0.99	0.99	0.99	0.99	0.97
Panel B: Alternative Estimation Methods and Data Sources							
Panel B.I: Linear VS Interpolation (Mar04-Dec18)							
Correlation in levels	0.90	0.89	0.76	0.82	0.85	0.83	0.99
Correlation in monthly changes	0.89	0.85	0.66	0.64	0.65	0.62	0.97
Correlation in weekly changes	0.88	0.86	0.61	0.59	0.68	0.62	0.97
Panel B.II: VIX settlement prices, not synchronized (Mar04-Dec18)							
Correlation in levels	0.97	0.97	0.99	0.99	0.99	0.99	0.98
Correlation in monthly changes	0.96	0.96	0.98	0.99	0.98	0.99	0.95
Correlation in weekly changes	0.95	0.94	0.97	0.98	0.97	0.98	0.90
Panel B.III: Bloomberg data for VS forward (Nov08-Dec18)							
Correlation in levels	0.55	0.72	0.54	0.58	0.69	0.68	0.91
Correlation in monthly changes	0.59	0.58	0.37	0.40	0.49	0.43	0.81
Correlation in weekly changes	0.62	0.49	0.31	0.34	0.40	0.41	0.73
Panel B.IV: CBOE VIX, VIX3M, VIX6M for VS forward (Nov08-Dec18)							
Correlation in levels	0.58	0.52	0.56	0.58	0.74	0.69	0.94
Correlation in monthly changes	0.65	0.33	0.39	0.51	0.54	0.46	0.87
Correlation in weekly changes	0.66	0.08	0.31	0.39	0.46	0.42	0.79
Panel B.V: Markit OTC quotes for VS forward (Monthly Sep06-Dec15)							
Correlation in Levels	0.53	0.60	0.57	0.71	0.80	0.77	0.87
Correlation in monthly changes	0.41	0.52	0.26	0.62	0.66	0.64	0.81

Table A.2: Deviation Measure Return Predictability Across Specifications

This table reports return predictability regressions for VIX futures hedged with variance swap forwards over a weekly horizon using different specifications of the deviation measure. Panel A repeats Panel A.II from Table 5 in the paper. The other panels follow the same approach. Across specifications, the deviation measure significantly predicts VIX futures hedged returns.

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

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Deviation from paper ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation	0.23***	0.25***	0.30***	0.28***	0.25***	0.25***
	(0.05)	(0.04)	(0.05)	(0.04)	(0.04)	(0.04)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.052	0.063	0.092	0.076	0.061	0.060
Panel B: Deviation 1-day lag ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation - 1d Lag	0.18**	0.14***	0.15***	0.14***	0.10***	0.08***
	(0.08)	(0.04)	(0.03)	(0.03)	(0.03)	(0.02)
Observations	3354	3681	3694	3647	3205	3037
Adjusted R^2	0.018	0.023	0.028	0.036	0.024	0.018
Panel C: Deviation 5-day moving average ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation - 5d MA	0.21**	0.15***	0.17***	0.14***	0.10***	0.09***
	(0.09)	(0.04)	(0.03)	(0.03)	(0.02)	(0.02)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.023	0.025	0.036	0.038	0.022	0.020
Panel D: Regression convexity adjustment ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation - Regression Adj.	0.30***	0.24***	0.27***	0.20***	0.16***	0.15***
	(0.07)	(0.04)	(0.05)	(0.03)	(0.03)	(0.03)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.049	0.063	0.090	0.073	0.058	0.056
Panel E: Term-structure model convexity adjustment ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation - Term-Structure Adj.	0.29***	0.26***	0.26***	0.20***	0.17***	0.15***
	(0.07)	(0.04)	(0.04)	(0.03)	(0.03)	(0.03)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.046	0.076	0.083	0.073	0.065	0.062
Panel F: Linear VS Interpolation ($h = 5$, Mar04-Dec18, daily overlapping)						
Deviation - Linear VS Interp.	0.24***	0.19***	0.15***	0.11***	0.12***	0.12***
	(0.06)	(0.06)	(0.03)	(0.02)	(0.03)	(0.02)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.035	0.039	0.027	0.023	0.036	0.039
Panel G: Bloomberg data ($h = 5$, Nov08-Dec18, daily overlapping)						
Deviation - Bloomberg Data	0.43***	0.17**	0.21***	0.15***	0.11***	0.11***
	(0.11)	(0.07)	(0.05)	(0.03)	(0.04)	(0.03)
Observations	2266	2510	2510	2510	2510	2510
Adjusted R^2	0.096	0.027	0.045	0.036	0.027	0.035
Panel H: VIX, VIX3M, VIX6M Indices ($h = 5$, Nov08-Dec18, daily overlapping)						
Deviation - VIX Indices Data	0.18**	0.08	0.15***	0.12***	0.13***	0.11***
	(0.08)	(0.05)	(0.04)	(0.04)	(0.03)	(0.02)
Observations	2266	2510	2510	2510	2510	2509
Adjusted R^2	0.016	0.005	0.024	0.024	0.037	0.032

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

Table A.3: Deviation Measure Return Predictability Across Horizons

This table reports return predictability regressions for VIX futures hedged with variance swap forwards over one-day, one-week, and one-month horizons. Panel A (B) reports results for the full (post-crisis) sample period. The deviation measure continues to significantly forecast returns across horizons. The explanatory power increases with the forecast horizon for the longer-dated contracts.

Return Predictability Regression: $R^{Fut}_{t+h,n} - \hat{\beta}_n R^{Fwd}_{t+h,n} = \gamma_n' x_{t,n} + \epsilon_{t+h,n}$						
Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A.I: Sample period 2004-2018, daily returns $h = 1$						
Deviation	0.32***	0.19***	0.17***	0.18***	0.17***	0.15***
	(0.05)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.100	0.036	0.027	0.033	0.030	0.023
Panel A.II: Sample period 2004-2018, weekly returns $h = 5$						
Deviation	0.23***	0.25***	0.30***	0.28***	0.25***	0.25***
	(0.05)	(0.04)	(0.05)	(0.04)	(0.04)	(0.04)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.052	0.063	0.092	0.076	0.061	0.060
Panel A.III: Sample period 2004-2018, monthly returns $h = 21$						
Deviation	0.12	0.32***	0.38***	0.38***	0.26***	0.38***
	(0.08)	(0.06)	(0.07)	(0.07)	(0.07)	(0.06)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.013	0.105	0.147	0.140	0.070	0.144
Panel B.I: Sample period 2010-2018, daily returns $h = 1$						
Deviation	0.37***	0.23***	0.21***	0.21***	0.20***	0.15***
	(0.08)	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.135	0.051	0.045	0.042	0.038	0.023
Panel B.II: Sample period 2010-2018, weekly returns $h = 5$						
Deviation	0.47***	0.35***	0.38***	0.32***	0.35***	0.30***
	(0.08)	(0.04)	(0.04)	(0.05)	(0.04)	(0.05)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.225	0.123	0.143	0.101	0.119	0.087
Panel B.III: Sample period 2010-2018, monthly returns $h = 21$						
Deviation	0.47***	0.53***	0.54***	0.43***	0.48***	0.51***
	(0.05)	(0.06)	(0.06)	(0.06)	(0.05)	(0.06)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.220	0.277	0.292	0.187	0.230	0.265

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

Table A.4: Deviation Measure Return Predictability – Percentage Returns

This table reports return predictability regressions for VIX futures hedged with variance swap forwards for percentage returns. It is analogous to Table 5 in the paper. As before, the deviation measure significantly predicts returns across contracts and sample periods.

Predictability Regression: $R^{Fut}_{t+h,n} - \hat{\beta}_n R^{Fwd}_{t+h,n} = \gamma'_n x_{t,n} + \epsilon_{t+h,n}$

Return Definitions: $R^{Fut}_{t+h,n} \equiv \frac{Fut_{t,n} - Fut_{t+h,n}}{Fut_{t,n}}, R^{Fwd}_{t+h,n} \equiv \frac{Fwd_{t,n} - Fwd_{t+h,n}}{Fwd_{t,n}}$

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample Period from 2004 to 2018						
Panel A.I: First Stage - Hedge Ratio						
β_n	0.38***	0.38***	0.36***	0.36***	0.38***	0.36***
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	2837	3697	3697	3655	3207	3197
Adjusted R^2	0.853	0.782	0.704	0.697	0.708	0.689
Panel A.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.37***	0.35***	0.34***	0.30***	0.27***	0.25***
	(0.04)	(0.04)	(0.03)	(0.03)	(0.04)	(0.04)
Observations	2837	3697	3697	3655	3207	3197
Adjusted R^2	0.133	0.122	0.115	0.090	0.072	0.061
Panel A.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.46***	0.44***	0.38***	0.34***	0.32***	0.28***
	(0.04)	(0.05)	(0.03)	(0.03)	(0.04)	(0.04)
Observations	2837	3697	3697	3655	3207	3197
Adjusted R^2	0.157	0.173	0.134	0.104	0.095	0.076
Panel B: Post-Crisis Sample from 2010 to 2018						
Panel B.I: First Stage - Hedge Ratio						
β_n	0.38***	0.40***	0.41***	0.37***	0.40***	0.37***
	(0.02)	(0.03)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	1705	2245	2245	2245	2245	2245
Adjusted R^2	0.869	0.817	0.773	0.732	0.736	0.735
Panel B.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.33***	0.29***	0.31***	0.32***	0.33***	0.27***
	(0.05)	(0.04)	(0.03)	(0.04)	(0.04)	(0.05)
Observations	1705	2245	2245	2245	2245	2245
Adjusted R^2	0.106	0.084	0.097	0.104	0.110	0.071
Panel B.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.44***	0.38***	0.37***	0.34***	0.34***	0.27***
	(0.05)	(0.05)	(0.04)	(0.04)	(0.04)	(0.05)
Observations	1705	2245	2245	2245	2245	2245
Adjusted R^2	0.140	0.150	0.118	0.112	0.117	0.074

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

Table A.5: Deviation Measure Return Predictability – Log Returns

This table reports return predictability regressions for VIX futures hedged with variance swap forwards for log-returns. It is analogous to Table 5 in the paper. As before, the deviation measure significantly predicts returns across contracts and sample periods.

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

Return Definition: $R^{Fut}_{t+h,n} = \ln(Fut_{t,n}/Fut_{t+h,n}), R^{Fwd}_{t+h,n} \equiv \ln(Fwd_{t,n}/Fwd_{t+h,n})$

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample Period from 2004 to 2018						
Panel A.I: First Stage - Hedge Ratio						
β_n	0.43***	0.41***	0.38***	0.36***	0.39***	0.37***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	2837	3697	3697	3655	3207	3197
Adjusted R^2	0.875	0.811	0.716	0.692	0.709	0.692
Panel A.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.48***	0.40***	0.35***	0.32***	0.28***	0.26***
	(0.07)	(0.04)	(0.03)	(0.03)	(0.04)	(0.04)
Observations	2837	3697	3697	3655	3207	3197
Adjusted R^2	0.229	0.157	0.125	0.101	0.076	0.065
Panel A.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.55***	0.48***	0.39***	0.35***	0.33***	0.28***
	(0.06)	(0.05)	(0.03)	(0.03)	(0.04)	(0.04)
Observations	2837	3697	3697	3655	3207	3197
Adjusted R^2	0.253	0.206	0.146	0.111	0.101	0.079
Panel B: Post-Crisis Sample from 2010 to 2018						
Panel B.I: First Stage - Hedge Ratio						
β_n	0.45***	0.45***	0.43***	0.38***	0.40***	0.38***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	1705	2245	2245	2245	2245	2245
Adjusted R^2	0.904	0.868	0.781	0.731	0.735	0.737
Panel B.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.52***	0.38***	0.36***	0.35***	0.34***	0.28***
	(0.05)	(0.04)	(0.03)	(0.04)	(0.04)	(0.04)
Observations	1705	2245	2245	2245	2245	2245
Adjusted R^2	0.268	0.144	0.127	0.122	0.117	0.078
Panel B.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.60***	0.47***	0.41***	0.36***	0.35***	0.28***
	(0.06)	(0.06)	(0.04)	(0.05)	(0.04)	(0.05)
Observations	1705	2245	2245	2245	2245	2245
Adjusted R^2	0.287	0.206	0.146	0.126	0.124	0.079

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

Table A.6: Deviation Measure Predictability for VIX Futures Returns Hedged with the Stock Market

This table reports return predictability regressions for VIX futures hedged with the stock market returns over a weekly horizon, $h = 5$. The regressions are analogous to Table 5 except for hedging with CRSP value-weighted returns instead of variance swap forward returns in the first step. Even with this less precise hedge, as evidenced by the slightly lower explanatory power in the first step regressions, the deviation measure remains significant at predicting VIX futures hedged returns for almost all of the contracts in the full sample and post-crisis periods.

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

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample Period from 2004 to 2018						
Panel A.I: First Stage - Hedge Ratios						
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	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.574	0.566	0.551	0.534	0.519	0.496
Panel A.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.22***	0.11**	0.05	0.13**	0.08*	0.04
	(0.08)	(0.06)	(0.05)	(0.05)	(0.05)	(0.05)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.046	0.012	0.002	0.016	0.007	0.002
Panel A.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.26***	0.08*	0.05	0.09**	0.05	0.00
	(0.08)	(0.05)	(0.05)	(0.04)	(0.04)	(0.04)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.063	0.048	0.040	0.040	0.033	0.032
Panel B: Post-Crisis Sample from 2010 to 2018						
Panel B.I: First Stage - Hedge Ratios						
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	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.592	0.676	0.669	0.651	0.620	0.606
Panel B.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.09	0.13**	0.15***	0.14***	0.15***	0.05
	(0.06)	(0.06)	(0.05)	(0.05)	(0.05)	(0.05)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.008	0.018	0.023	0.019	0.022	0.002
Panel B.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	0.14**	0.14**	0.13***	0.11**	0.13***	0.01
	(0.06)	(0.06)	(0.05)	(0.05)	(0.05)	(0.05)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.033	0.023	0.024	0.029	0.030	0.013

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

Table A.7: Deviation Measure Predictability for VS Forward Returns Hedged with the Stock Market

This table reports return predictability regressions for variance swap forwards hedged with the stock market returns over a weekly horizon, $h = 5$. The regressions are analogous to Table A.6. In comparison to VIX futures, the deviation measure exhibits less predictability for variance swap forwards for most maturities, particularly in the post-crisis period. At the same time, the deviation measure is significant for some contracts with the expected, negative, sign.

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

Contract (n)	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample Period from 2004 to 2018						
Panel A.I: First Stage - Hedge Ratios						
β_n	0.04***	0.03***	0.02***	0.02***	0.02***	0.02***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.509	0.383	0.321	0.379	0.409	0.389
Panel A.II: Second Stage - Predicting Returns with Deviation						
Deviation	0.02	-0.07	-0.26***	-0.10	-0.13**	-0.18***
	(0.12)	(0.07)	(0.06)	(0.07)	(0.05)	(0.04)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.000	0.004	0.065	0.011	0.017	0.034
Panel A.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	-0.00	-0.16***	-0.29***	-0.18***	-0.21***	-0.25***
	(0.09)	(0.05)	(0.07)	(0.05)	(0.05)	(0.05)
Observations	3355	3697	3697	3655	3207	3197
Adjusted R^2	0.024	0.089	0.100	0.045	0.060	0.073
Panel B: Post-Crisis Sample from 2010 to 2018						
Panel B.I: First Stage - Hedge Ratios						
β_n	0.04***	0.03***	0.02***	0.02***	0.02***	0.02***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.550	0.613	0.572	0.530	0.537	0.524
Panel B.II: Second Stage - Predicting Returns with Deviation						
Deviation	-0.29***	-0.12**	-0.18***	-0.14***	-0.16***	-0.23***
	(0.07)	(0.06)	(0.06)	(0.05)	(0.06)	(0.06)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.083	0.014	0.031	0.020	0.024	0.050
Panel B.III: Second Stage - Predicting Returns with Deviation and Controls						
Deviation	-0.28***	-0.16***	-0.22***	-0.17***	-0.18***	-0.26***
	(0.10)	(0.05)	(0.06)	(0.05)	(0.05)	(0.06)
Observations	2029	2245	2245	2245	2245	2245
Adjusted R^2	0.109	0.025	0.046	0.033	0.042	0.062

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

Table A.8: Factor Loadings for Alpha-to-Margin Estimates

This table reports the factor loadings for the alpha-to-margin estimates in Table 8. The loadings for Panel D in Table 8 are omitted to save space but are qualitatively similar to the Panel C loadings.

Weekly Alpha-to-Margin Estimates

Specification	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Variance swap forward hedge						
Alpha	5.23***	5.23***	5.00***	4.68***	4.72***	4.55***
	(0.37)	(0.37)	(0.42)	(0.39)	(0.39)	(0.43)
RMRF	0.51	0.53	-0.01	0.40	0.37	0.15
	(0.39)	(0.45)	(0.29)	(0.27)	(0.28)	(0.40)
HML		0.14	0.21		0.06	0.02
		(0.29)	(0.29)		(0.37)	(0.36)
SMB		-0.32	-0.25		0.10	0.13
		(0.52)	(0.48)		(0.32)	(0.33)
MOM		-0.01	0.06		-0.39	-0.45*
		(0.17)	(0.17)		(0.25)	(0.25)
VS1			10.41			11.34
			(6.50)			(8.55)
VX1			-0.00			-0.00
			(0.00)			(0.00)
Adjusted R^2	0.009	0.009	0.022	0.005	0.007	0.014
Panel B: E-mini hedge						
Alpha	3.32***	3.35***	3.89***	2.68***	2.76***	3.19***
	(0.52)	(0.52)	(0.49)	(0.60)	(0.61)	(0.61)
RMRF	-1.13***	-1.27***	1.13***	-0.75*	-0.78	1.82***
	(0.35)	(0.40)	(0.42)	(0.44)	(0.50)	(0.58)
HML		-0.32	-0.43		-1.06**	-0.77**
		(0.41)	(0.34)		(0.48)	(0.39)
SMB		0.68	0.36		0.39	0.15
		(0.49)	(0.40)		(0.62)	(0.52)
MOM		-0.32	-0.38		-1.21***	-1.02***
		(0.27)	(0.24)		(0.42)	(0.36)
VS1			-17.13***			-18.11**
			(5.42)			(8.09)
VX1			-0.00***			-0.00**
			(0.00)			(0.00)
Adjusted R^2	0.027	0.029	0.110	0.010	0.026	0.102
Panel C: E-mini hedge and t-costs						
Alpha	1.59***	1.62***	2.15***	1.11*	1.19**	1.60***
	(0.51)	(0.51)	(0.49)	(0.60)	(0.60)	(0.61)
RMRF	-1.12***	-1.27***	1.09***	-0.74*	-0.77	1.76***
	(0.35)	(0.39)	(0.42)	(0.44)	(0.50)	(0.58)
HML		-0.34	-0.45		-1.05**	-0.76**
		(0.41)	(0.34)		(0.48)	(0.39)
SMB		0.70	0.39		0.41	0.18
		(0.48)	(0.40)		(0.61)	(0.52)
MOM		-0.35	-0.40*		-1.22***	-1.03***
		(0.27)	(0.24)		(0.41)	(0.36)
VS1			-16.59***			-17.47**
			(5.47)			(8.15)
VX1			-0.00***			-0.00**
			(0.00)			(0.00)
Adjusted R^2	0.026	0.029	0.106	0.010	0.026	0.098
Factor Model	CAPM	FFC4	FFCV6	CAPM	FFC4	FFCV6
Post-Crisis	No	No	No	Yes	Yes	Yes
Observations	3697	3697	3697	2245	2245	2245

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

Figure A.1: Estimating the Law of One Price Deviation Measure – Robustness to Estimation Method and Data

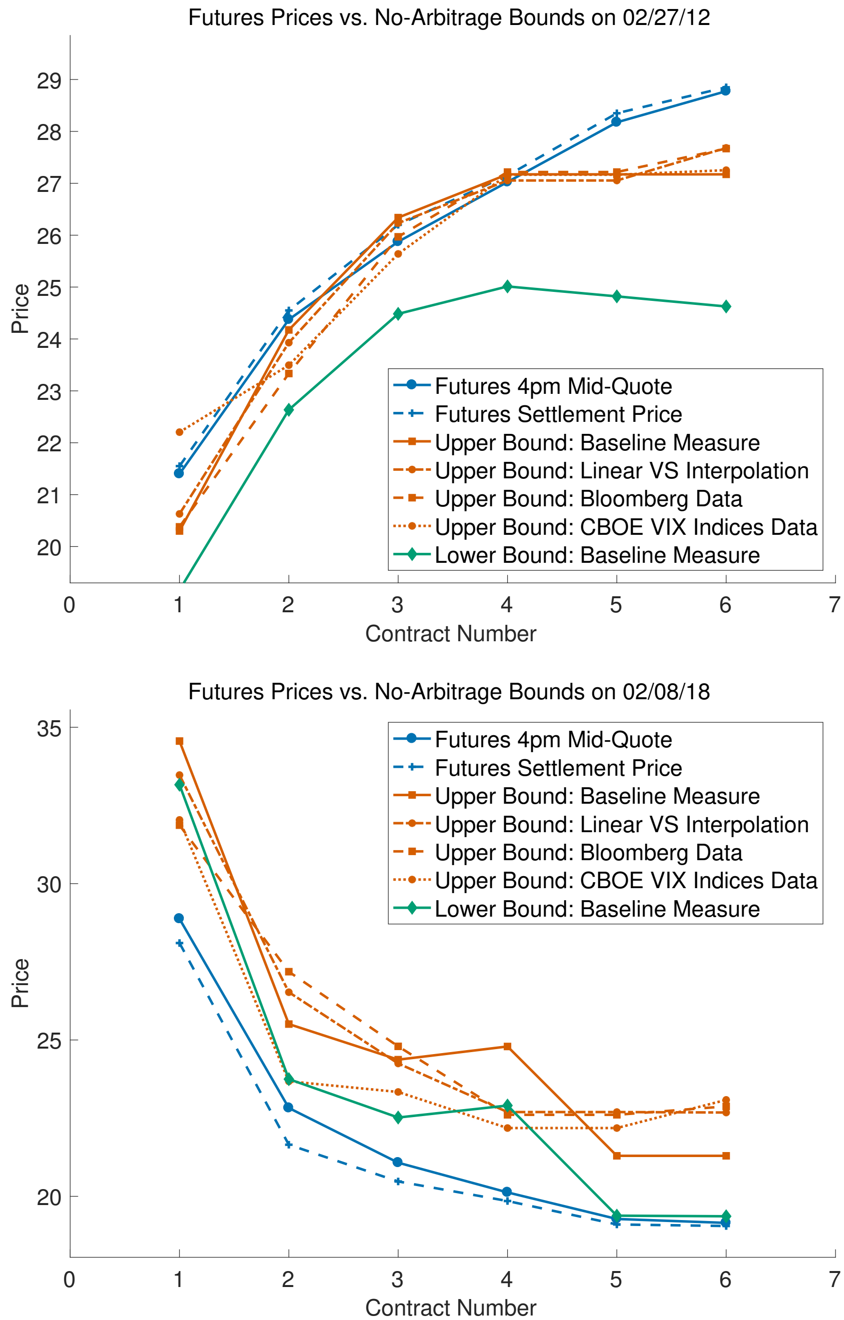

This figure compares VIX futures prices at 4pm and at 4:15pm (settlement) to variance swap forward rates estimated from a variety of alternative data sources on two different days. The results illustrate the robustness of Figure 3. Regardless of which of the four data sources are used to compute variance 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 below estimates of the upper and lower bounds.

Figure A.2: Deviation Measure Robustness Across Estimation Methods and Datasets

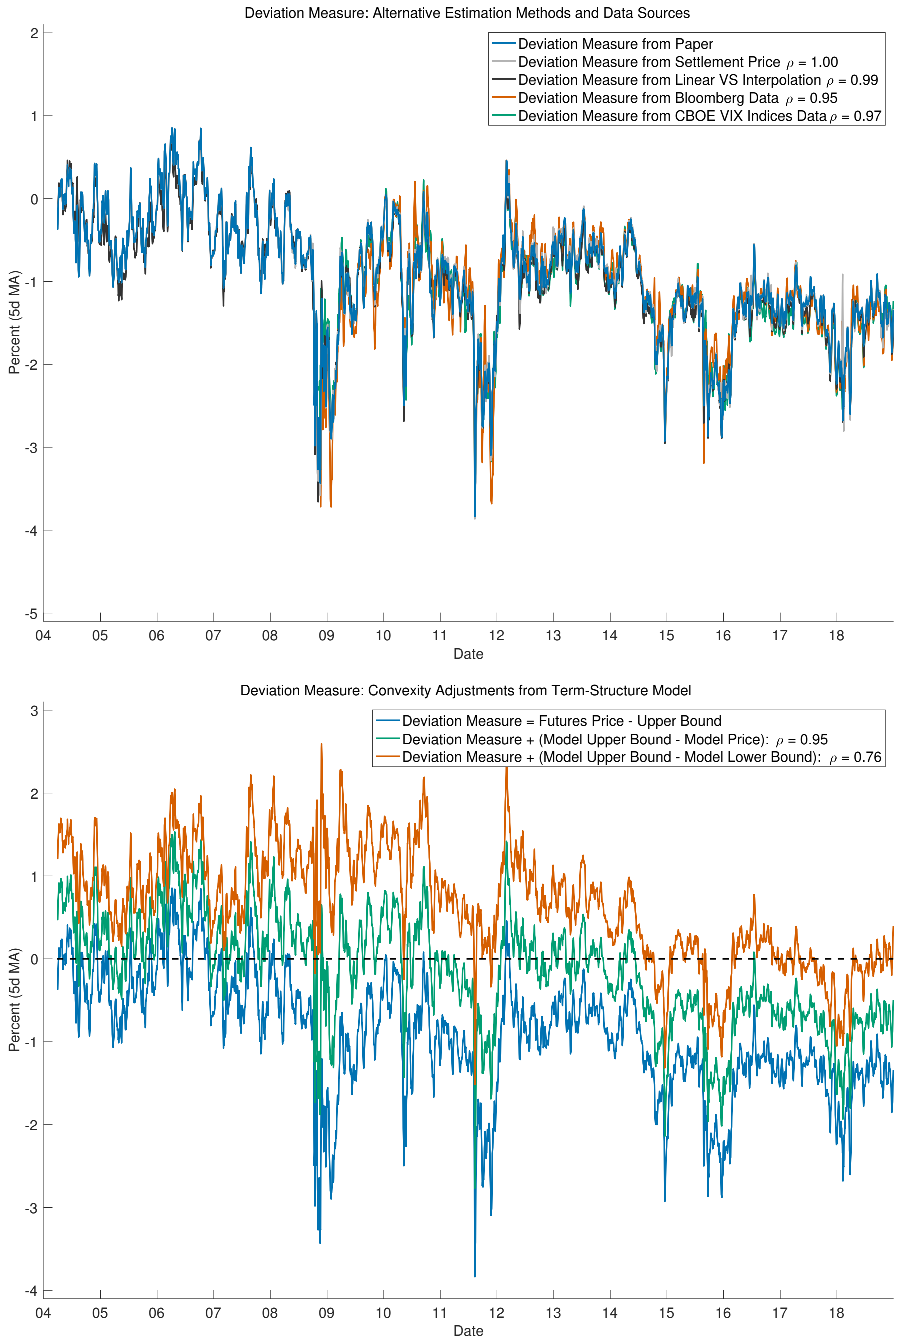

As in Figure 1, this figure reports the average deviation across the front six futures contracts as a five-day moving average. The top plot reports the baseline deviation measure from the paper against alternative measures from different estimation approaches and data sources. The baseline measure is highly correlated with the alternative measures as in Table A.1. The bottom plot reports the baseline deviation measure plus convexity adjustments from a term-structure model. As in the top plot, the measures remain highly correlated after the convexity adjustments. Adding the difference between the model upper bound and model futures price reduces the bias in the deviation measure (green line). Positive values relative to the upper bound (blue line) and negative values relative to the lower bound (red line) are law of one price violations.

Figure A.3: VIX Futures Trading Strategy: Position and Post-Crisis Sharpe Ratios

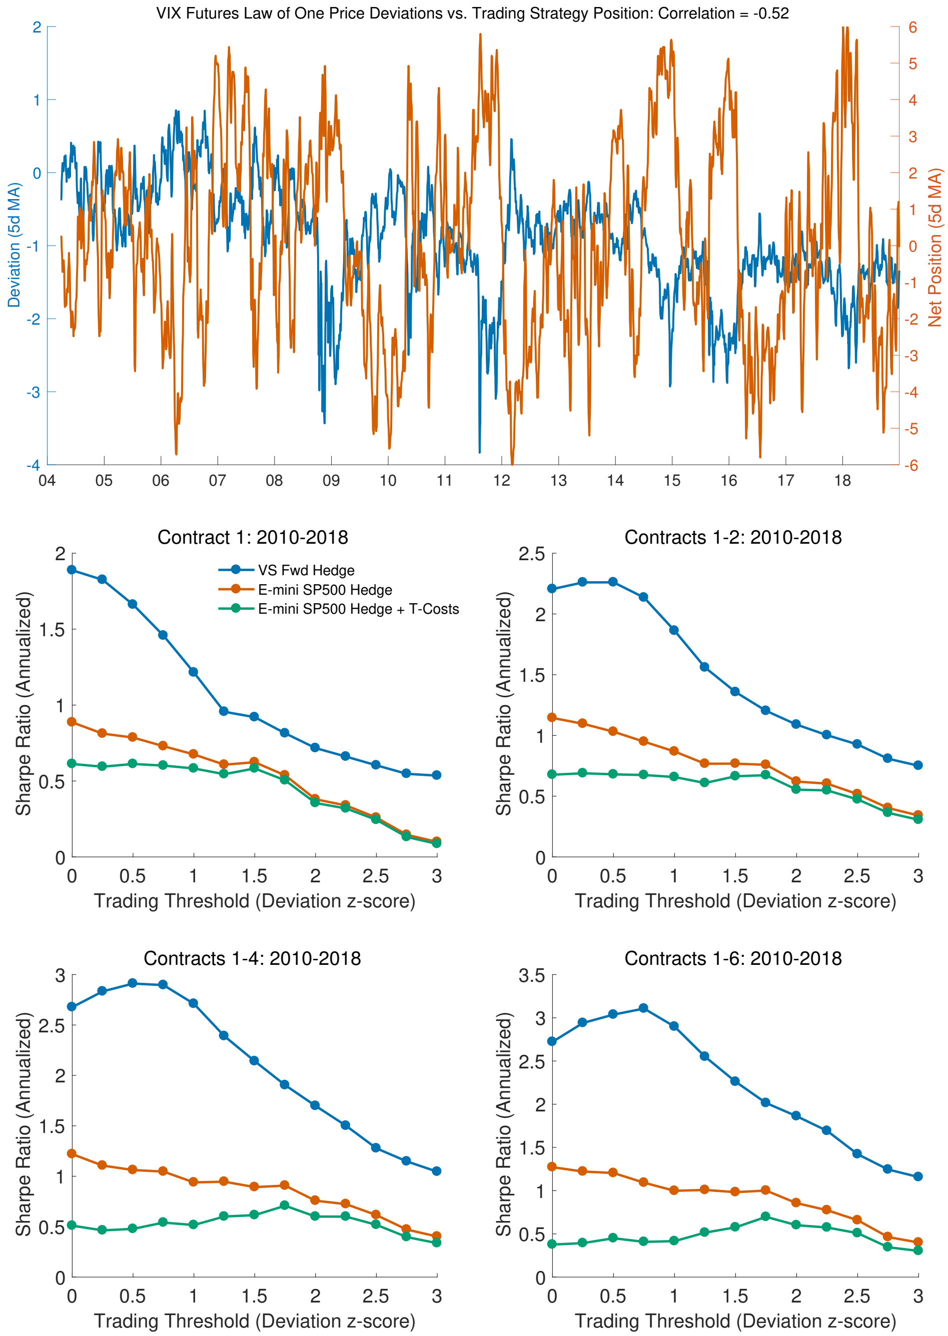

The top figure plots the net position in the relative value trading strategy against the deviation measure over time. A long position corresponds to buying VIX futures that are hedged with variance swap forwards. Since the strategy trades the front six contracts, the number of net positions is bounded 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 6 but for the 2010 to 2018 post-crisis period.

Figure A.4: Law of One Price Deviations for the Front Six VIX Futures Contracts

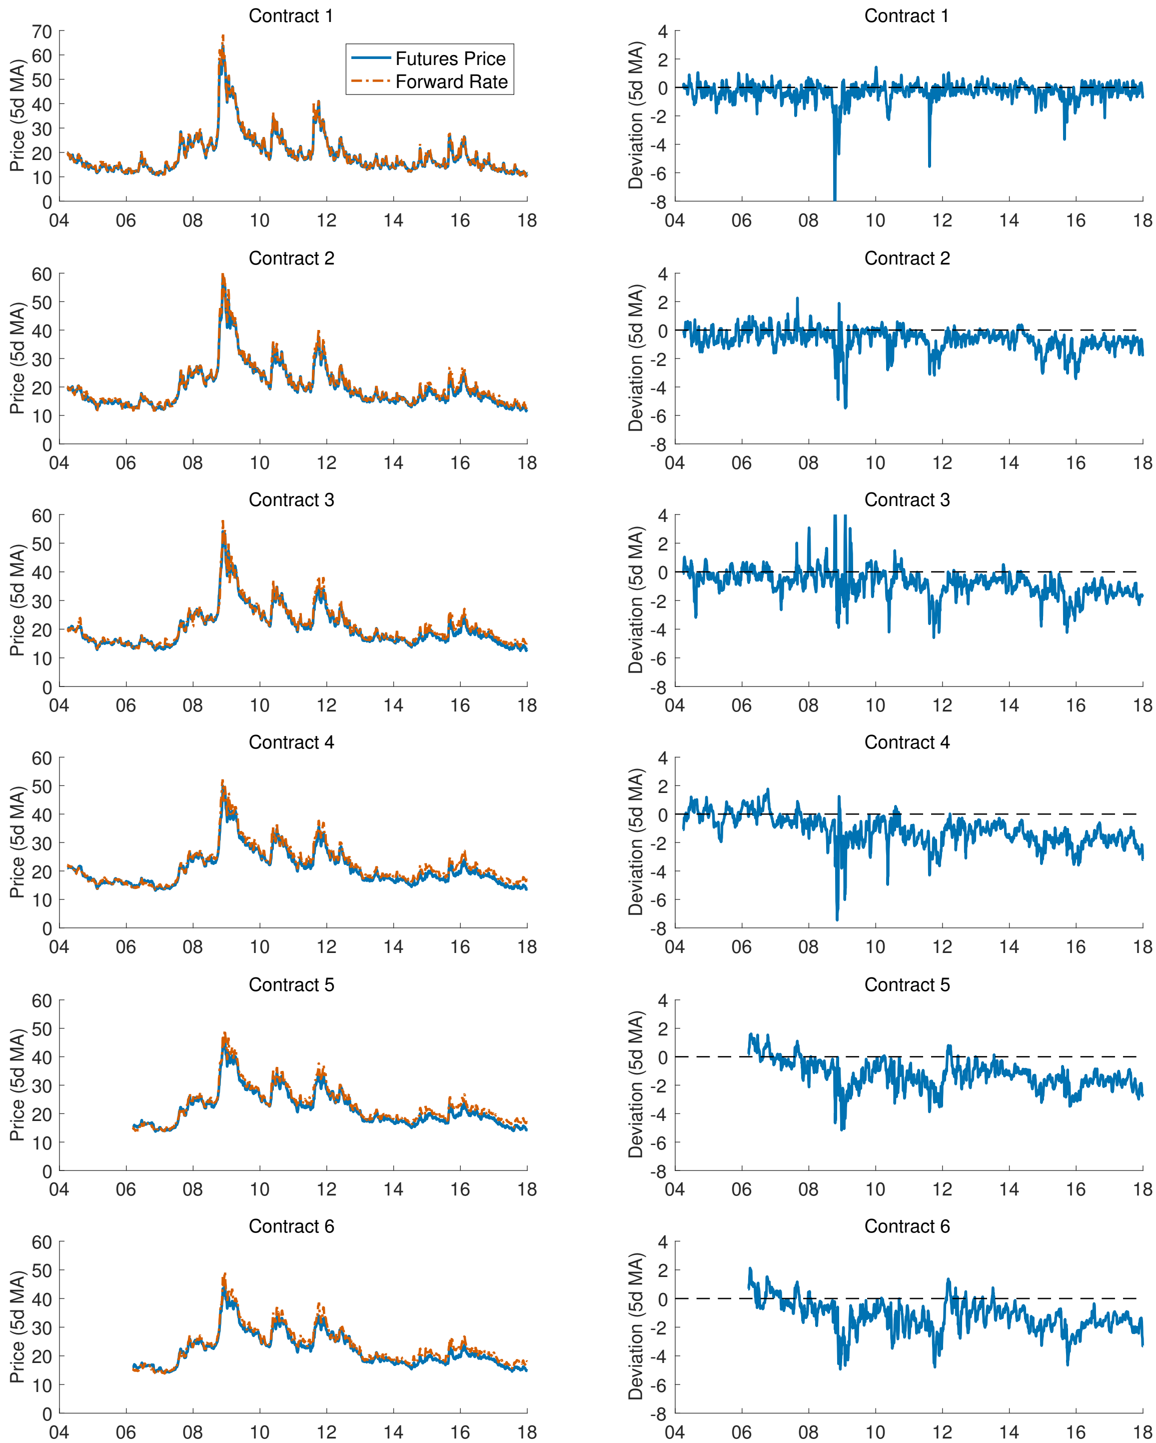

This figure plots the VIX futures price and variance swap forward rate against the deviation measure on the right for each of the front six futures contracts. The prices and deviation measure for the longer-dated contracts become available later in the sample, motivating the 2007-2018 and 2010-2018 sample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

Figure A.5: Distribution of VIX Futures Prices Relative to the No-Arbitrage Upper Bound from January 3, 2007 to December 31, 2018

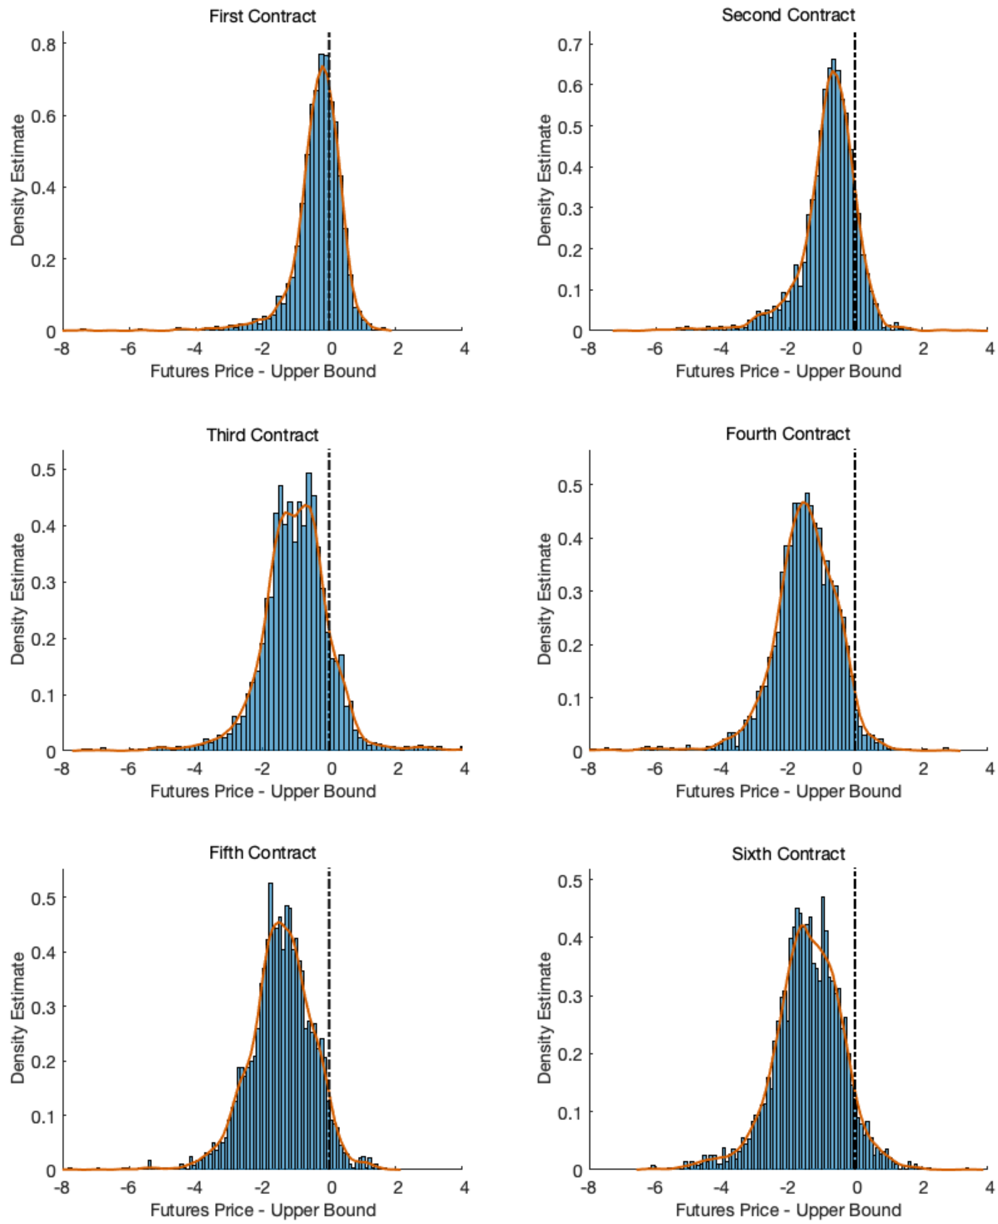

This figure plots histograms and kernel density estimates of the law of one price deviation measure by VIX futures contract from January 3, 2007 to December 31, 2018. The distribution is negatively skewed for the front contract and second contract. The histograms indicate the presence of law of one price 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.6: Distribution of VIX Futures Prices Relative to the No-Arbitrage Lower Bound from January 3, 2007 to December 31, 2018

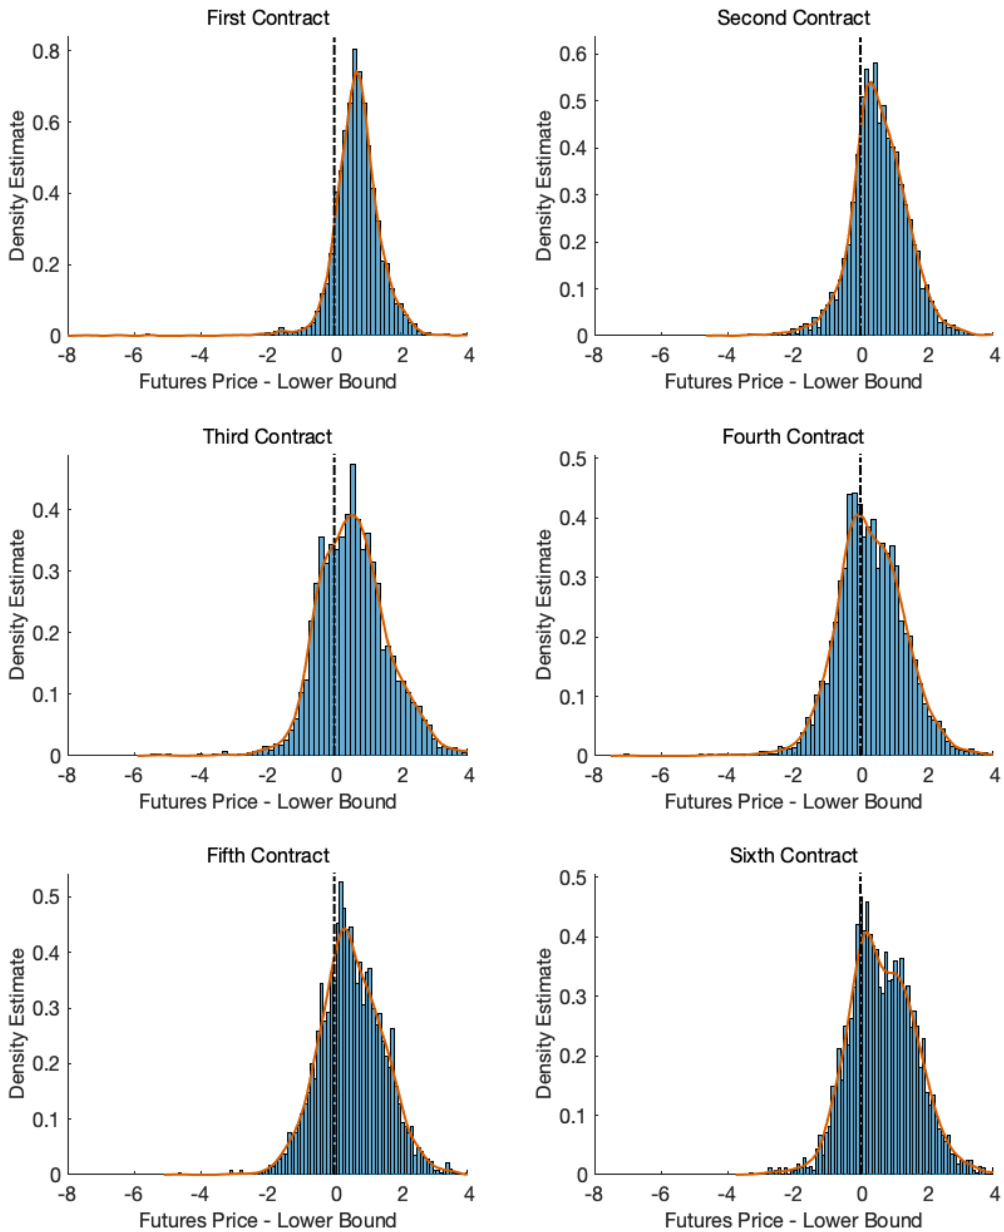

This figure plots histograms and kernel density estimates of VIX futures prices relative to their law of one price lower bound from January 3, 2007 to December 31, 2018. The lower bound is computed as $Fwd_{t,n} - (UB_{t,n} - LB_{t,n})$ where $UB_{t,n}$ is the variance swap forward rate and $LB_{t,n}$ 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 **?**. 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.7: Deviation Measure versus Risk and Demand Variables

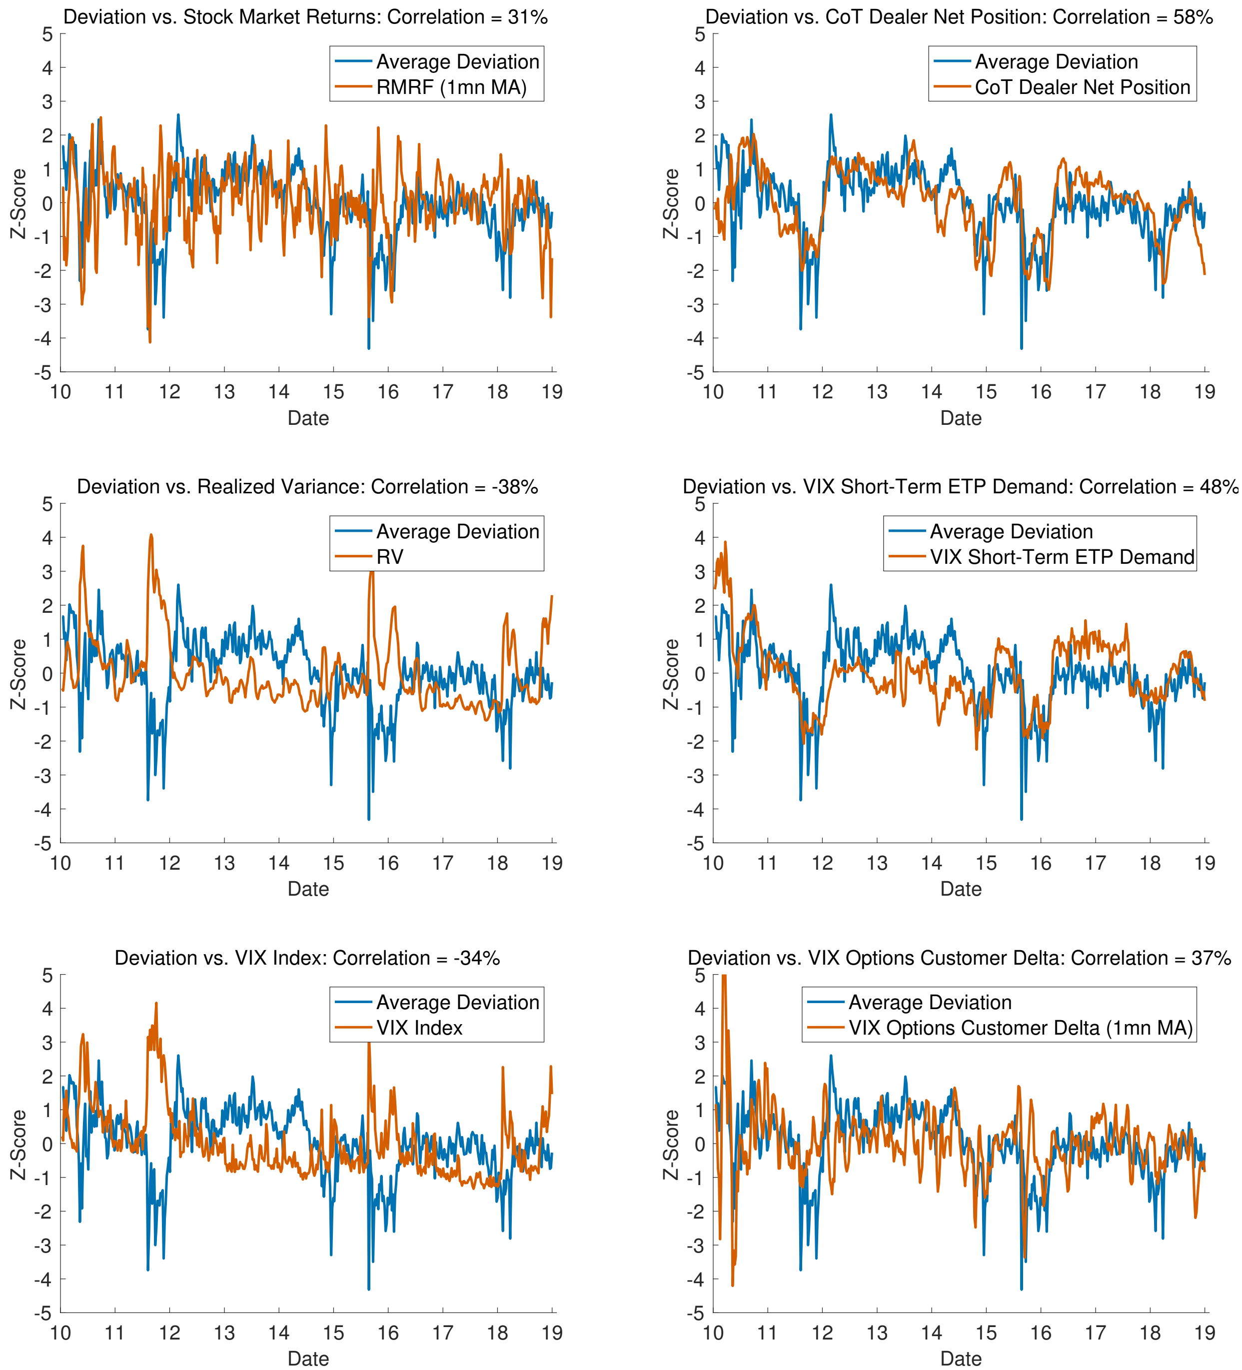

This figure plots the deviation measure against different risk and demand variables. The risk variables include stock market returns, realized variance, and the VIX index. Increases in risk as measured by negative stock market returns or increases in volatility are negatively correlated with the deviation 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.

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

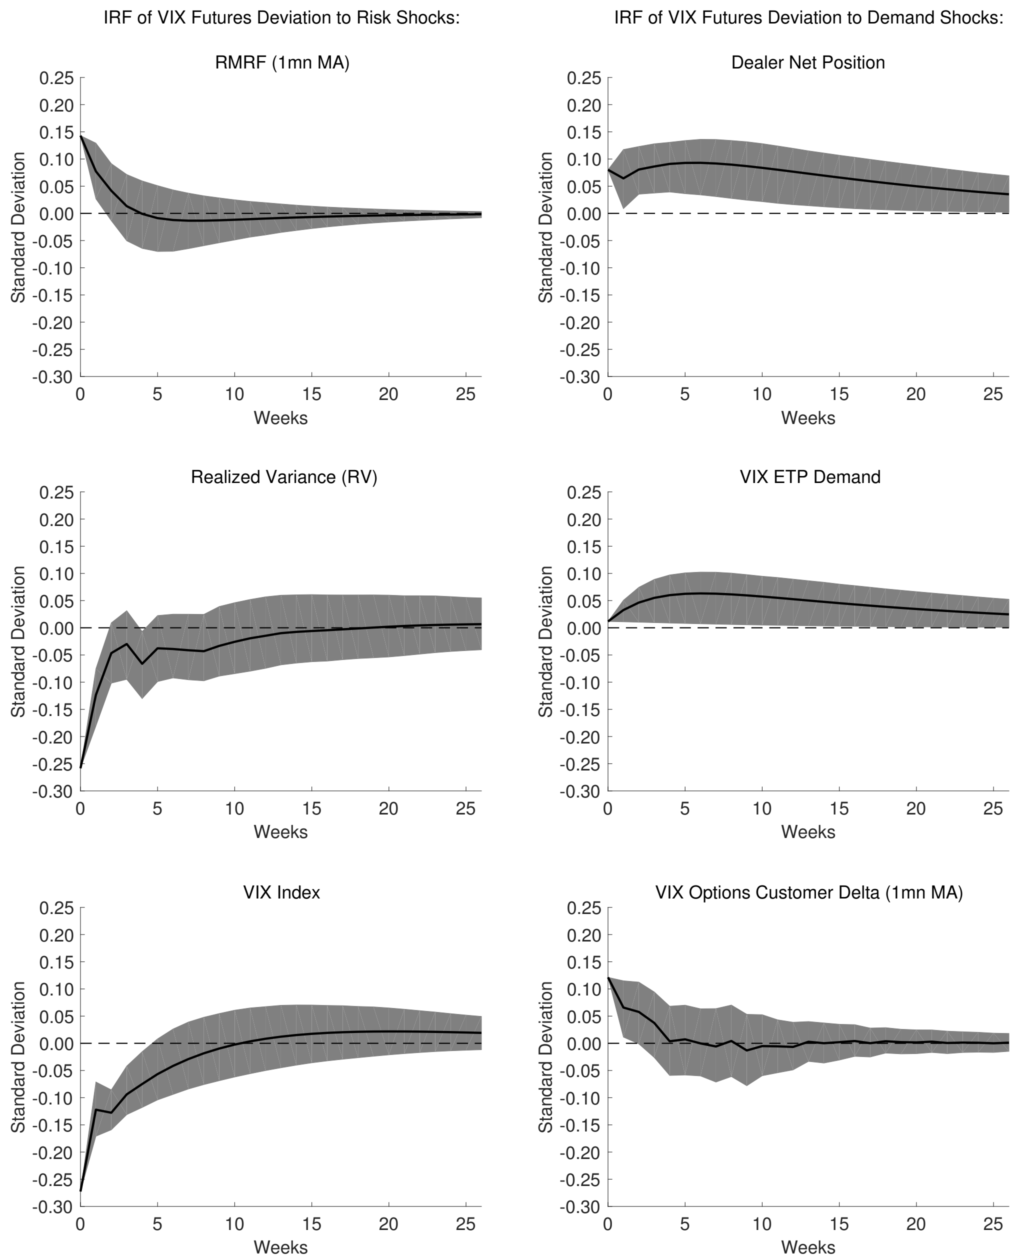

This figure plots impulse response functions from bivariate vector autoregressions to illustrate how the deviation measure reacts to risk and demand shocks. The left column reports the IRFs from bivariate VARs with risk variables. The right column reports the IRFs from bivariate VARs with demand variables. 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 Functions of Deviation and Dealer Position to VIX and Dealer Position Shocks Across Trivariate VAR Specifications

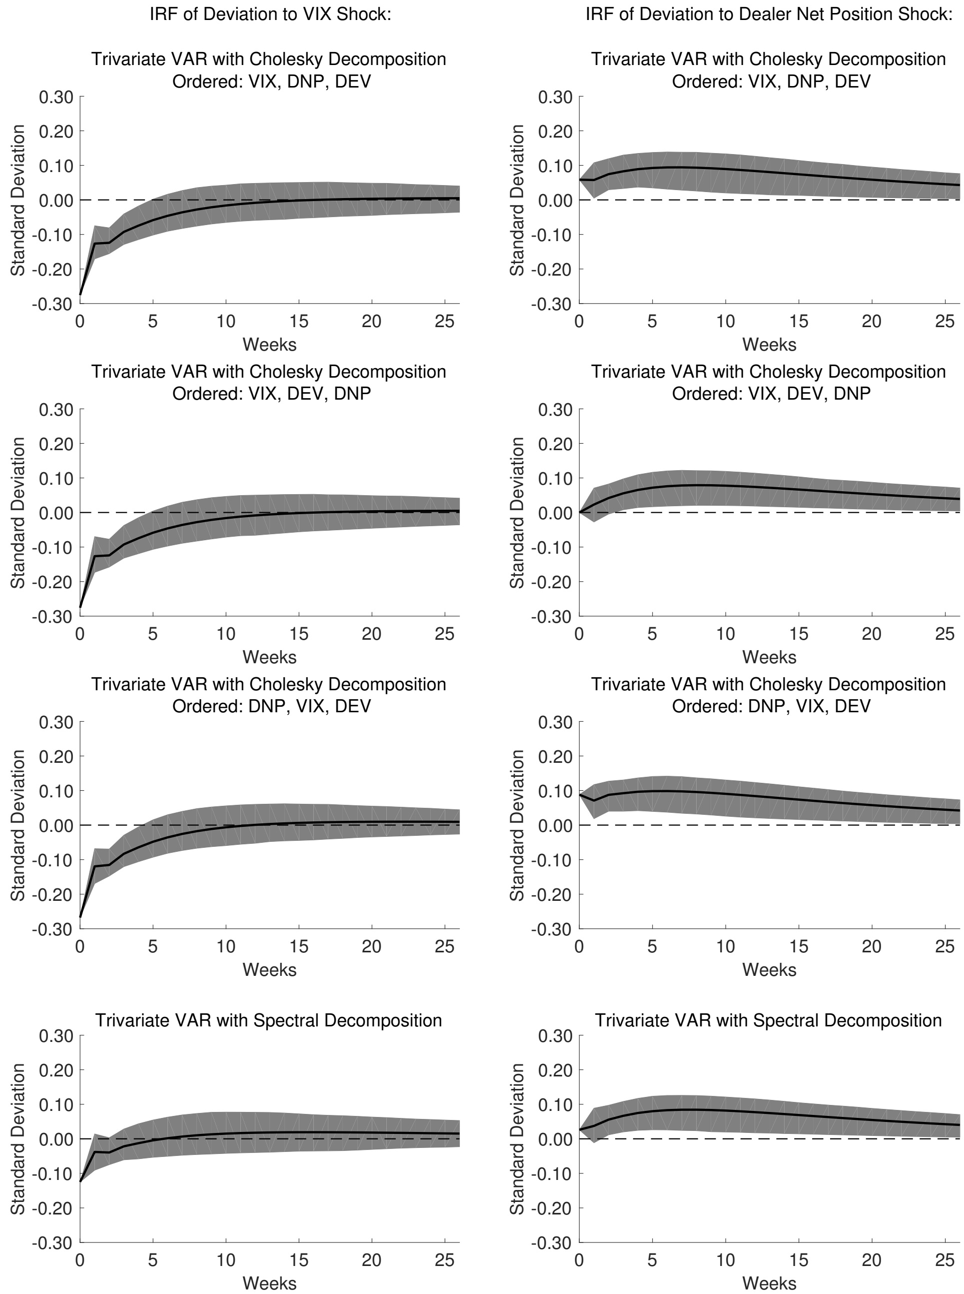

This figure plots the impulse response function of of the no-arbitrage deviation and dealer position to VIX and dealer position shocks across different trivariate VAR specifications. The first three rows report different orderings of the variables for a Cholesky decomposition. The fourth row reports the IRFs for a spectral decomposition. The variables in the vector autoregression are $y_t = [DEV_t\ VIX_t\ DNP_t]$ The VAR is estimated using weekly data from 2010 to 2018 with two lags. The IRFs take on similar shapes across specifications. 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.