DEPARTMENT OF ECONOMICS
AND BUSINESS ECONOMICS
AARHUS UNIVERSITY

CREATES
Center for Research in Econometric Analysis of Time Series

Betting on mean reversion in the VIX? Evidence from ETP flows

Ole Linnemann Nielsen and Anders Merrild Posselt

CREATES Research Paper 2022-06

Department of Economics and Business Economics
Aarhus University
Fuglesangs Allé 4
DK-8210 Aarhus V
Denmark

Email: oekonomi@au.dk
Tel: +45 9352 1936

Betting on mean reversion in the VIX? Evidence from ETP flows*

Ole Linnemann Nielsen[†] **Anders Merrild Posselt**[‡]

This version: September 1, 2021

* The authors were partially funded by a grant from the Danish Council for Independent Research (DFF - 0133-00151B). The authors are grateful for helpful comments and suggestions from Søren B. Brøgger, Kim Christensen, Charlotte Christiansen, Jonas N. Eriksen, Niels S. Grønborg, Michael Halling (discussant), Frank de Jong, Paul Karehnke, Cesario Mateus, Stig V. Møller, Anders B. Trolle, and Roine Vestman. We thank seminar participants at CREATES and Tilburg University and conference participants at the 2020 PhD Nordic Finance Network (BI Norwegian Business School).

† Jyske Bank, Vestergade 8-16, 8600 Silkeborg, Denmark. Email: oln@jyskebank.dk.

‡ CREATES, Department of Economics and Business Economics, Aarhus University, Fuglesangs Alle 4, 8210 Aarhus V, Denmark. Email: amp@econ.au.dk. Corresponding author.

Betting on mean reversion in the VIX?
Evidence from ETP flows

Abstract

We investigate flows of VIX ETPs with long volatility exposure. We find an inverse relation between flows and the level of the VIX, implying that investors sell VIX ETPs when the VIX is at elevated levels, consistent with investors incorporating the typical mean reverting behavior of volatility. We find no evidence supporting that investors consider exposure to risk factors when they evaluate VIX ETP performance. Finally, our results suggest that large outflows following increases in the VIX may be a partial explanation of the "low premium response puzzle" in the VIX premium.

Keywords: VIX ETPs; Flows; Asset Pricing Tests; VIX Premium

JEL Classification: G11; G12; G13; G14

This version: September 1, 2021.

I. Introduction

Since January 2009, volatility as an asset class has been available to all types of investors via VIX exchange traded products (ETPs). These are complex and risky products designed to provide either long or inverse exposure to market volatility. Since their launch, VIX ETPs have gained much traction. Figure I depicts the price and cumulative net flows of the *iPath S&P VIX Short Term Future ETN*(VXX), one of the largest and most liquid ETPs. As of February 2019, the cumulative net flow of VXX amounts to approximately \$8 bn. Despite the complex structure, requiring a certain level of sophistication to comprehend, VIX ETPs have reeled in a lot of interest from retail investors. Having access to holdings data, Todorov (2020) reports that the fraction of institutional holdings is less than 24%, on average, for the period 2009-2018.

In this paper, we are interested in ETPs with long volatility exposure. In the remainder of the paper, the term "VIX ETPs" simply refers to these. VIX ETPs are often promoted as tools for portfolio diversification, as they offer protection against down markets when volatility is high. However, insurance comes at a cost. The cost is realized through losses of the constant maturity strategy in VIX futures, which the products are tracking (see the left-hand axis of Figure I).[1]

The main contribution of this paper is to provide the first investigation of flows into VIX ETPs. Given the special characteristics and large popularity of VIX ETPs, we investigate how investors apply these products. Our study provides three main insights. First, by examining the relation between flows and the VIX, we find that increases in the VIX are typically accompanied by large outflows, which is consistent with investors incorporating the typical mean reversion of volatility in how they trade VIX ETPs. Second, by applying a revealed preferences approach as in Berk and van Binsbergen (2016), Barber et al. (2016), and Agarwal et al. (2018) we find no evidence of investors adjusting for exposure to risk

[1] See, e.g., Christensen et al. (2020) on the structure of VIX ETPs.

Figure I: Price development and cumulative net inflow of VXX

Figure I shows the price and cumulative net flow of VXX, from inception to February 2019. Over the lifetime of the product, the value of VXX has been severely eroded, and the issuer has made five 1-for-4 reverse splits. The depicted price development has been adjusted for these hence, the magnitude of the left-hand y-axis.

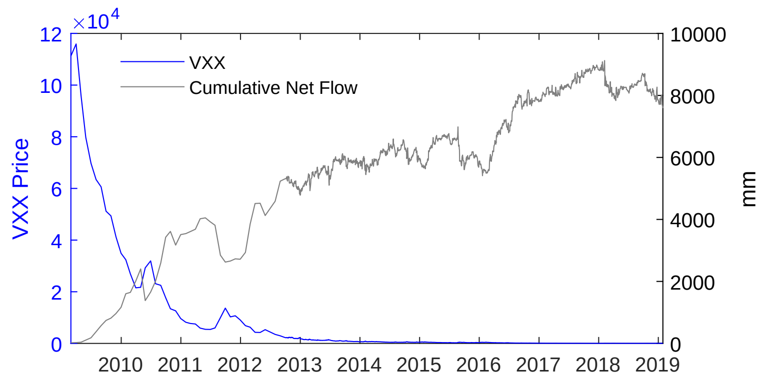

factors when they evaluate the performance of these products. Finally, we document that the low VIX premium response puzzle discovered by Cheng (2019) is likely related to price impact caused by the flow pattern of VIX ETPs.

VIX ETPs are closely linked to the VIX, which is known to be mean reverting by nature. Our starting hypothesis is that investors incorporate this mean reversion in how they trade VIX ETPs. Hence, following a VIX increase, we would expect to see outflows. Indeed, this is also the case. First, by regressing dollar flows, summed across ETPs, on the VIX and its lags, we obtain negative coefficients for the initial lags and positive but smaller in magnitude coefficients for longer lags. An event study of flows confirms this effect, as the largest increases in the VIX tend to be followed by very large outflows. We also find that the inverse relation between flows and the VIX varies with the level of persistence, defined by the speed-of-mean-reversion, in the VIX. When persistence is low, the inverse relation between flows and the VIX is slightly enhanced.

For our second finding, we consider the framework of Berk and van Binsbergen (2016), Barber et al. (2016), and Agarwal et al. (2018). These papers study the relation between flows and performance of mutual- and hedge funds, where performance is measured by abnormal returns (alpha) implied by different asset pricing models. Along with raw returns, we consider five different asset pricing models. The alphas for the different models are individually applied as explanatory variables in regressions that have the flows or the sign of the flows as the dependent variable. We hypothesize, that for holding VIX ETPs and thereby incurring negative expected returns, investors must price the systematic risk of VIX ETPs with some asset pricing model. The idea behind our hypothesis is that positive updates of alpha are followed by positive updates of flows. A positive model alpha will imply that VIX ETPs are cheap relative to their risk exposure, and investors will be inclined to buy the products. By comparing the ability of different models to predict flows, we can infer which type of risk, if any, is of concern to investors in VIX ETPs. From all the models we consider, raw returns are the only measure that significantly explains the flows of VIX ETPs. Thus we find no evidence of investors pricing any systematic risk in VIX ETPs. Put in another way, investors do not evaluate whether the insurance that VIX ETPs provide is considered expensive or cheap by the asset pricing models considered in our study.

To obtain our last finding, we estimate the VIX premium at different horizons. Given the flow pattern documented in our previous analyzes, we hypothesize that investor behavior causes issuers of the ETPs to reduce positions in VIX futures when the VIX increases. Consequently, this will put downward price pressure on VIX futures, contributing to the low VIX premium response puzzle documented in Cheng (2019). The relation between flows and VIX premiums is first studied via a simple OLS regression where we regress changes in the VIX premium on changes in aggregated dollar-flows. We find a positive relation between the flows of short-term products and VIX premiums, and the relation is decreasing on the horizon of the VIX premium. We investigate this positive relation further by means of a quantile regression. We find that the correlation between flows and VIX premiums is particularly high when the VIX premium is large in absolute terms. Finally, we consider bivariate dynamics using a vector autoregression of flows and the VIX premium. The impulse response functions show that a shock to flows impacts the VIX premium. All in all, our results indicate that the flow pattern in VIX ETPs may provide a partial explanation of the low premium response puzzle.

The prior literature on volatility assets can be divided into three different categories. The first category includes papers that examine the diversification benefits of volatility assets in broad investment portfolios. Recent studies such as Bordonado et al. (2017) and Berkowitz and DeLisle (2018) all reject the existence of any potential diversification benefits of VIX ETPs. Christensen et al. (2020), however, find that including VIX ETPs in a dynamic asset allocation strategy can have substantial economic value. The second category is concerned with the causality and price discovery between different markets of volatility assets. In particular, between the VIX futures market and the spot VIX (e.g., Shu and Zhang (2012), Bollen et al. (2017), and Fernandez-Perez et al. (2019)). The final category is concerned with the pricing of volatility assets (e.g., Zhang and Zhu (2006), Zhang and Huang (2010), and Gehricke and Zhang (2018)). To the best of our knowledge, none of the previous studies investigate how investors use volatility assets as investments. We do this by analyzing the flows of VIX ETPs.

This paper is also related to the literature on variance- and VIX premia. Bollerslev et al. (2009) and Bekaert and Hoerova (2014) calculate the variance risk premium as the risk-neutral minus physical expectation of the 30-day variance of S&P 500 returns. Both papers find that the variance risk premium positively predicts stock returns. Closer related to our work, Cheng (2019) defines the VIX premium as the risk-neutral minus the physical expectation of the future value of the VIX. Interestingly, both the variance- and VIX premia tend to become negative during periods of elevated market risk. Negative risk premia suggest that the demand for insurance decreases around periods of elevated market risk, which seems illogical. The findings in our study add support to the argument that this puzzle exists because of the systematic pricing impact of market participants and are not just due to measurement error.

The paper is organized into five additional sections. Section II briefly describes our data. Section III contains our first empirical results, where we investigate how flows relate to the VIX. Section IV examines the impact of risk factors, and Section V links flows to the VIX premium. Finally, Section VI contains the conclusions.

II. Data

The dependent variable of main interest is the daily net flows of VIX ETPs. We follow prior literature on fund flows (cf. Barber et al. (2016)) and define this as:

$$F_{it} = \frac{AUM_{it}}{AUM_{it-1}} - (1 + R_{it}). \tag{1}$$

Hence, it is the growth in assets under management (AUM) adjusted for the total return (R_{it}), assuming that all flows take place at the end of the day. Flows in dollar terms are defined as:

$$DF_{it} = AUM_{it} - (1 + R_{it}) \times AUM_{it-1}. \tag{2}$$

We obtain daily prices and AUM on all VIX ETPs from Bloomberg and daily close levels of the VIX from the Chicago Board Options Exchange (Cboe). Even though the first VIX ETPs were issued already in 2009, AUM is not available at a daily frequency before October 2012. As a consequence, the sample period spans from October 2012 to February 2019.

In this study, we are interested in the flow patterns of ETPs with a long exposure towards volatility. We consider both short-term (e.g., VXX) and mid-term products (e.g., VXZ) and also include products that use leverage (e.g., TVIX). Products with an inverse exposure towards volatility (e.g., XIV) or an average AUM of less than \$ 15 million are excluded from the sample.[2]

The final sample includes six different VIX ETPs. Table I, presents summary statistics of the ETPs in our sample. The average daily flow equals 0.49%, with a median of 0.11%. The median ETP has \$179.65 million (mm) in AUM, while the average is approximately twice as

[2] See Appendix A for a list of the VIX ETPs that we include in our sample.

large (\$354.24 mm). This suggests that there is a positive skew in product size. The average age is 5.45 years (median 5.45 years), and the average expense ratio is 1.08%. Finally, the average spread is 1.22 basis points (bp), with a median of only 0.12 bp. This indicates that there are large differences in liquidity across products.

Table I: VIX ETP summary statistics

This table presents summary statistics of characteristics for the six ETPs in our sample across 8001 ETP-day observations in the period October 2012 to February 2019. The data include only VIX ETPs with direct exposure to the VIX futures. Daily flow is calculated by Equation (1).

	Mean	Std. Dev	Median
Daily Flow	0.49%	6.26%	0.11%
Size (\$mil)	354.24	405.08	179.65
Age (years)	5.45	2.09	5.45
Yearly Expense Ratio	1.08%	0.33%	0.89%
Spread (bp)	1.22	5.92	0.12

III. Flows and the VIX

VIX ETPs are by construction closely linked to the VIX, incurring high positive returns when the VIX increases against negative returns when the VIX is low. As other volatility measures, the VIX is characterized by occasional swings from low to high levels, as depicted in Figure II. Despite varying degrees of persistence, the VIX eventually declines from the high levels, and the process appears to be mean reverting over the long run.[3]

Our first analysis investigates how the flows of VIX ETPs relate to the VIX. We make the empirical prediction that investors in VIX ETPs incorporate the mean reverting nature of the VIX in how they trade the products. At an increase in volatility, an investor who expects a reversal to the long term mean will also expect lower future returns of VIX ETPs and reduce her position in the instrument. Hence, following an increase in the VIX, we would expect to see net outflows.

Denote, $AggDF_t = \sum_{i=1}^{N} DF_{it}$, as the dollar flows aggregated across the different ETPs in

[3]See also Whaley (2009) for a description of the mean reverting nature of the VIX.

Figure II: The VIX

Figure II shows the level of the VIX during our sample period.

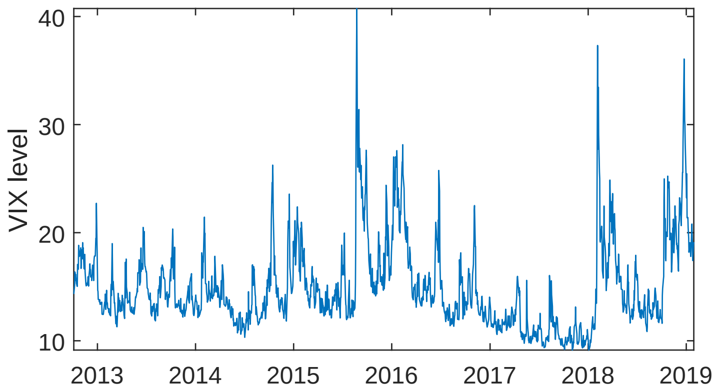

our study, at time t. We investigate our hypothesis by regressing the aggregated dollar flows on the VIX, lagged values of the VIX, and lagged flows:[4]

$$AggDF_t = a + \sum_{j=0}^{7} b_{j+1} \times VIX_{t-j} + \sum_{i=1}^{6} c_i \times AggDF_{t-i} + u_t. \quad (3)$$

Table II reports the estimated regression coefficients of the model in Equation (3). First, we note that the regression coefficient for the VIX at time t is positive, implying that an increase in the VIX tends to be accompanied by net inflows to VIX ETPs on the same day. Given the insurance-like characteristics of VIX ETPs, this is consistent with the flight-to-safety behavior that is typically seen in other markets at the signs of stress. For example, investors buy outright protection such as put options, causing the VIX to increase, or they buy safe-haven assets like treasuries and money-market instruments (cf. Longstaff (2004), Baele et al. (2019), and Adrian et al. (2019)). For the one to three days lag, the coefficients are all negative but only statistically significant at the one-day lag. A one-unit (one volatility point) increase in the VIX is associated with a net outflow of 24.49 mm at $t-1$. For longer lags (four to six days), the coefficient becomes positive again but much smaller in magnitude

[4] By means of an augmented Dickey-Fuller test of the VIX during our sample period, we reject the null hypothesis of a unit root at a 95% confidence level.

(12.81 mm at $t-4$ and 3.67 mm at $t-6$). In summary, the coefficients suggest that investors buy VIX ETPs on the day of an increase. In the few days following an increase, they sell before they again start to buy at longer horizons. This pattern seems consistent with investors expecting volatility to decrease after an initial increase. Then, as time has passed, volatility is likely to have declined and investors start to buy again as further decreases are less likely. Overall, the pattern that we observe in the regression coefficients of lagged VIX is in the direction we would expect if investors incorporate the typical VIX behavior in how they trade VIX ETPs.

Finally, we note that the coefficients for the lagged values of aggregated flows are statistically insignificant at all lags.

III.A. Flow pattern around the largest increases of the VIX

The above results suggest that investors sell VIX ETPs on the days following right after an increase in the VIX. To further illuminate this inverse relation, we now examine the flow pattern around the largest increases in the VIX by conducting an event study. Specifically, we define an event as a relative change of the VIX above its 95th percentile over the entire sample period. This gives us 78 event days in the sample. We consider an event window of -21 to +21 days relative to the VIX increase at day 0 (event day). Again, our variable of interest is the aggregated dollar flows. For each day t in the event window, we sum the flows over all previous days to obtain the cumulative flows as:

$$CumAggDF_t = \sum_{\tau=1}^{t} AggDF_\tau. \tag{4}$$

Figure III plots the average of $CumAggDF$ and the average level of the VIX across the 78 different events. From day -21 until day -11, there is a small upward trend in cumulative flows, where the level of the VIX is low. From day -10 until the event day, the VIX is slightly

Table II: AggDF and the VIX

This table presents the results for the regression as in Equation (3). The dependent variable is aggregated dollar flows, calculated as, $AggDF_t = \sum_{i=1}^{N} DF_{it}$. Explanatory variables are the VIX and lagged values of the VIX and aggregated dollar flows. T-statistics in parenthesis are calculated with Newey-West standard errors. ***, and ** represent significance at 1% and 5% respectively.

Intercept	63.30*** (4.87)
VIX_t	12.49** (2.19)
VIX_{t-1}	-24.49*** (-6.50)
VIX_{t-2}	-4.18 (-0.98)
VIX_{t-3}	-6.07 (-1.05)
VIX_{t-4}	12.81** (2.19)
VIX_{t-5}	1.92 (0.84)
VIX_{t-6}	3.67** (2.04)
AggDF_{t-1}	-0.03 (-0.94)
AggDF_{t-2}	-0.05 (-1.20)
AggDF_{t-3}	-0.03 (-0.71)
AggDF_{t-4}	0.04 (1.26)
AggDF_{t-5}	0.01 (0.23)
AggDF_{t-6}	-0.03 (-1.04)
Observations	1525
$R^2(\%)$	20.04

increasing, and the cumulative flows fluctuate around a fixed level. This indicates that the direction of cumulative investor flows becomes ambiguous as uncertainty increases. Then, after day 0, the event day, following a large decline, the cumulative flows become negative. The downward trend for flows is strongest for the first four days after the event and diminishes at day +8. In the remaining days of our event study, cumulative flows trend upwards. This aggregate trading pattern suggests that investors sell at elevated levels of the VIX. Then as the VIX stabilizes at lower levels, investors start to buy VIX ETPs again, which is consistent with the regression output in our previous analysis.

Figure III: Event study

Figure III depicts the average cumulative flows (left-hand y-axis) and the average level of the VIX index (right-hand y-axis) over the event window.

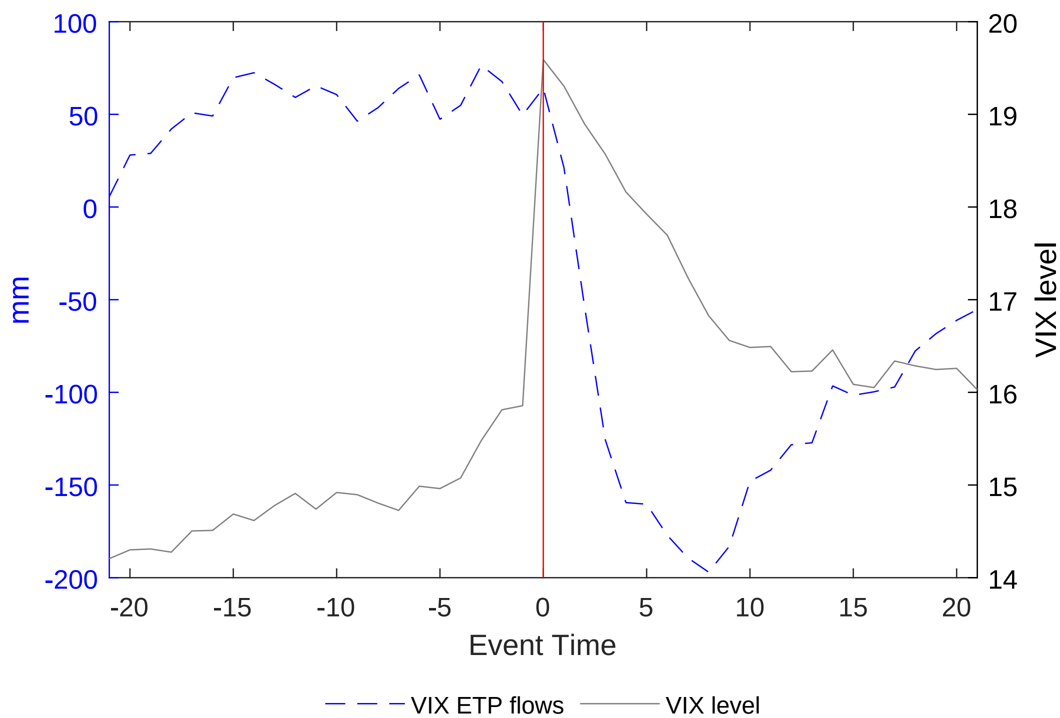

Table III reports the cumulative flows from Equation (4) with corresponding t-statistics . For the days -20 to -11, most of the cumulative flows are significantly positive. As we get closer to the event day, the VIX increases, and the cumulative flows become statistically insignificant from zero. After the event day, cumulative flows are, on average, negative. From day +3 to +9, the cumulative flows are significantly negative, which implies that there are large outflows in the days following a spike in the VIX. From day +10 and onward, the cumulative flows are no longer significantly below zero, which indicates that investors start to buy the VIX ETPs again as the VIX stabilizes at lower levels. Thus, outflows of VIX ETPs appear in large chunks after increases in the VIX. Inflows, on the other hand, occur continuously in lower magnitudes at low levels of the VIX.

We also perform separate event studies for each of the following product types: short-term, mid-term, leveraged, and un-leveraged. First, Figure IVa depicts the event study for short-term products. Our sample contains two leveraged products, and as both are short-term, we exclude them from the short-term group. We see that the pattern in flows is similar to the one in our initial event study, although there is no clear upward trend prior to the spike. After the event day, there are massive outflows, which decrease in magnitude as the VIX stabilizes at a lower plateau. After the stabilization of the VIX, outflows are eventually replaced by inflows. Second, Figure IVb focuses on mid-term products. For mid-term products, we see a noisy pattern in flows across our event window. In the days leading up to the spike, there is an upward trend in flows, although small in magnitude in terms of dollars. After the spike, we see further inflows, which are then followed by outflows. The event study for the mid-term products is the only case where cumulative flows are higher in the days right after the spike than the days just before the spike. However, if we look at the magnitude of cumulative dollar-flows, we see that they are much more modest for mid-term products than for short-term products. Third, Figure IVc depicts the evolution of flows for leveraged products. There is a positive trend in flows prior to the spike, whereas subsequent to the spike, there are large outflows for the first three days. The outflows then decrease in magnitude, and from day +7, there is a strong upward trend in cumulative flows. The reversal in flows, as the VIX stabilizes at lower levels, is also present for un-leveraged products, as seen in Figure IVd, although the reversal effect is much larger for leveraged products.

Table III: Event study

This table provides the results of the event study. The first column represents each day, t, in the event window. The second column shows the average cumulative flows. The third column shows the t-statistics of the cumulative flows. The fourth column shows the average value of the VIX on each day in the event window.

t	CumAggDF	t-stat	VIX level
-21	5.60	0.651	14.21
-20	28.04	2.076	14.30
-19	28.93	1.770	14.31
-18	42.06	2.081	14.28
-17	50.85	2.127	14.50
-16	49.10	1.775	14.51
-15	69.77	2.262	14.69
-14	72.45	2.348	14.62
-13	66.15	2.159	14.78
-12	59.13	1.755	14.91
-11	65.38	1.843	14.74
-10	60.65	1.555	14.92
-9	46.31	1.083	14.90
-8	53.57	1.192	14.81
-7	63.97	1.440	14.73
-6	71.30	1.548	14.99
-5	47.32	0.970	14.96
-4	54.88	1.107	15.08
-3	76.45	1.459	15.48
-2	67.66	1.267	15.81
-1	49.57	0.912	15.86
0	64.09	1.095	19.59
1	21.23	0.358	19.30
2	-53.92	-0.912	18.89
3	-125.66	-2.105	18.57
4	-159.50	-2.658	18.16
5	-160.35	-2.598	17.92
6	-177.15	-2.762	17.70
7	-189.35	-2.872	17.24
8	-197.00	-2.841	16.83
9	-183.11	-2.493	16.56
10	-148.01	-1.921	16.48
11	-142.00	-1.867	16.49
12	-128.26	-1.668	16.22
13	-127.22	-1.677	16.23
14	-96.53	-1.213	16.46
15	-101.57	-1.279	16.09
16	-99.72	-1.268	16.05
17	-97.11	-1.196	16.34
18	-77.63	-0.946	16.29
19	-68.36	-0.819	16.25
20	-61.15	-0.710	16.26
21	-55.31	-0.637	16.03

Figure IV: Event study - split on different products

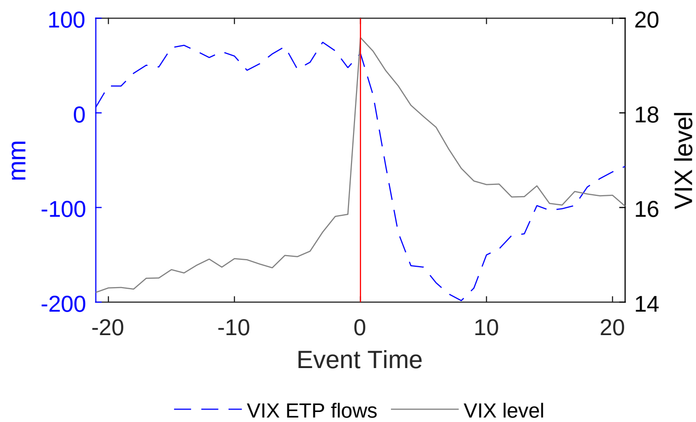

(a) Short-term products

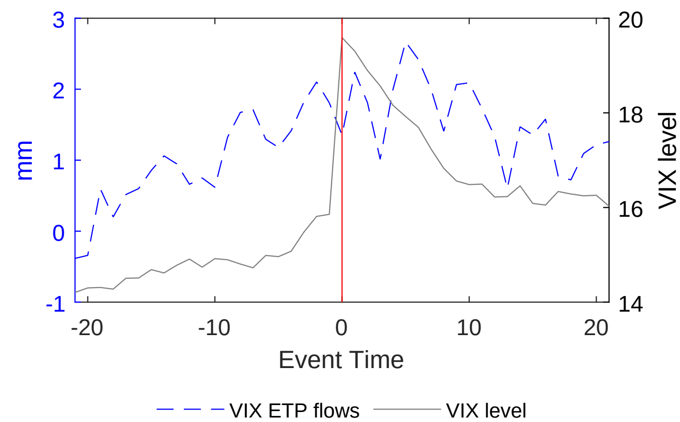

(b) Mid-term products

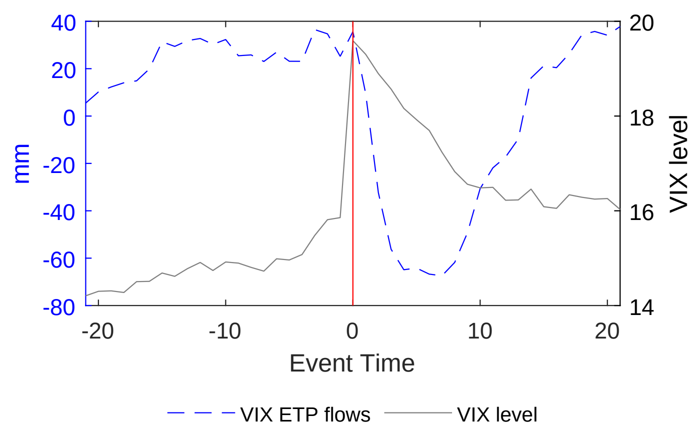

(c) Leveraged products

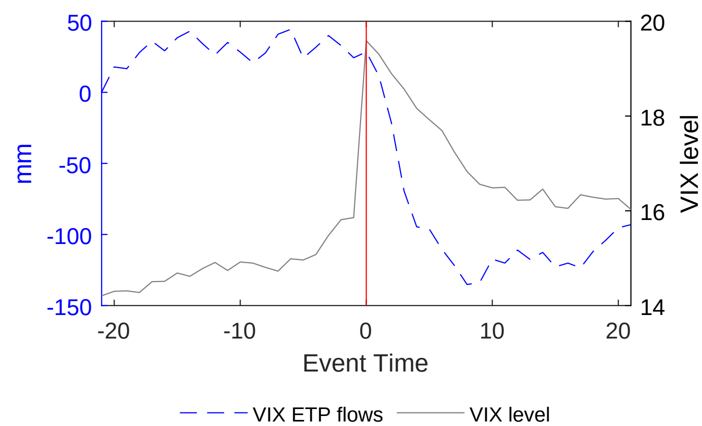

(d) Unleveraged products

Overall the event study shows a clear inverse relation between flows and the level of the VIX. Peak values in the VIX are followed by large outflows, suggesting that investors expect a decrease in volatility and the future returns of the VIX ETPs. Interestingly, the large outflows occur while the assets generate positive returns. VIX ETPs typically have their largest appreciation on the day of the increase in the VIX (the event day), but returns tend to stay positive for several days due to persistent backwardation in the VIX futures term structure. As an example, consider the VIX increase on August 24, 2015 ("Black Monday" on the Chinese stock market). On this day, VXX, the largest product, had a return of 17.7%. Over the following six days, the total outflows amounted to $719 mm, despite an average daily return of 4.2%.

III.B. Speed of mean reversion

Persistence in the VIX varies through time. From Figure II, we see that there are periods where the reversion towards lower levels after spikes happens fast and periods where high levels are more persistent. We now examine whether the level of persistence in the VIX impacts how investors trade VIX ETPs. It is likely that during periods where reversals from high levels happen at a faster pace, investors will sell more following an increase in the VIX. Vice versa, in periods where volatility is more persistent, investors are more uncertain about the direction of future volatility and will be less inclined to sell following an increase in the VIX. We hypothesize that in periods with low persistence, defined by the high speed of mean reversion, the negative relation between flows and lagged VIX is amplified.

We quantify the speed of mean reversion in the VIX by its half-life. Specifically, we define the half-life of the VIX as the expected number of days it takes for the VIX to reduce half of the distance to its long-run mean. In order to quantify the speed of mean reversion, we assume that the stochastic process of the level of the VIX follows a mean reverting Ornstein-Uhlenbeck (OU) process:

$$dVIX_t = \kappa(VIX^{LR} - VIX_t)dt + \sigma dW_t. \tag{5}$$

Here, $\kappa > 0$ is the speed of reversion for the OU process, VIX^{LR} is the long-run mean of the VIX, VIX_t is the current value of the VIX, $\sigma > 0$ is the instantaneous volatility, and W_t is a Brownian motion. The process in Equation (5) is mean reverting since negative deviations $((VIX^{LR} - VIX_t) > 0)$ from the long-run mean, on average, lead to upward revisions in the VIX, and vice versa.

In order to operationalize the process in Equation(5), we discretize it to an AR(1) process.

$$VIX_t = c + \phi VIX_{t-1} + b\epsilon_t, \quad |\phi| < 1. \tag{6}$$

For $c = \kappa VIX^{LR}\Delta t$, $\phi = (1 - \kappa\Delta t)$, $b = \sigma\sqrt{\Delta t}$, and $\epsilon \sim \mathcal{N}(0, 1)$, we get the Euler-Maruyama discretization of the OU process in Equation (6). Weak stationarity implies that $\mathbb{E}[VIX_t] = \mu = c/(1 - \phi)$ for all values of t. This allows us to define deviations from the stationary mean as:

$$\widehat{VIX}_t = \phi\widehat{VIX}_{t-1} + u_t, \tag{7}$$

where $\widehat{VIX}_t = VIX_t - \mu$ is the deviation from the long-run mean of the VIX. In order to calculate the expected number of days it takes for the VIX to reduce half of the distance to its long-run mean, we solve for k in the following expectation:

$$
\begin{aligned}
\mathbb{E}[\widehat{\mathrm{VIX}}_{t+k}] &= 0.5\widehat{\mathrm{VIX}}_t, \\
\Leftrightarrow \phi^k \widehat{\mathrm{VIX}}_t &= 0.5\widehat{\mathrm{VIX}}_t, \\
\Leftrightarrow k &= -\frac{\log(2)}{\log|\phi|}.
\end{aligned} \tag{8}
$$

Here, the expectation, $\mathbb{E}[\widehat{\mathrm{VIX}}_{t+k}]$, is obtained via recursive substitution. To estimate the half-life, k, we use a rolling window of data comprising the latest two years of data (2*252 days). Further, we use an log-transformation of the VIX. The time-series dynamics of the half-life of the VIX is presented in Figure V.

Figure V: Half-life of the VIX

Figure V shows the expected number of days, k, it takes for the VIX to reduce half of the distance to its long-run mean as defined in Equation (8). Grey areas correspond to the upper (lower) decile of days where the speed of mean reversion (half-life) is high (low).

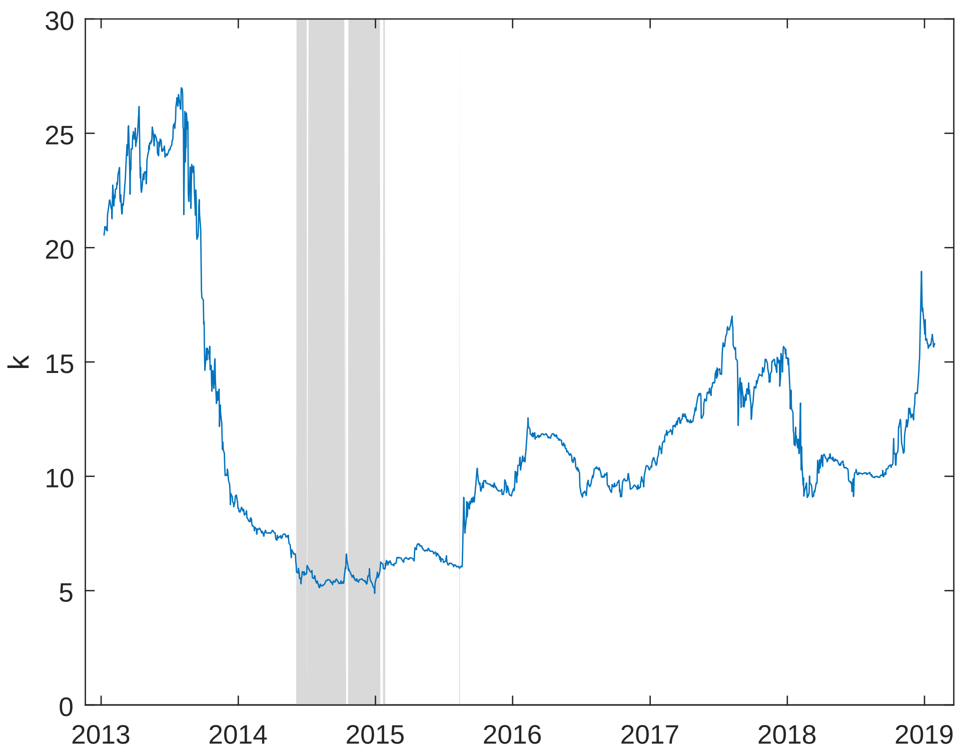

We re-run the regression in Equation (3) for days where the half-life of the VIX is below its 10th percentile. The output is reported in Table IV. As the negative coefficients for lags larger than one day were (and still is) insignificant, we only report the coefficient estimate for VIX_{t-1}. Comparing the results in Panel A and B, we find that when the speed of mean reversion is high, a one unit increase in the VIX leads to an aggregated dollar outflow of 25.21 mm on the following day. In periods with higher persistence in the VIX (lower speed of mean reversion), the outflows following a VIX increase are marginally lower (23.46 mm). In Panel C, we test whether this difference is significant by augmenting the regression with an interaction term between VIX_{t-1} and a dummy that takes on the value 1 for dates where the half-life is below its 10th percentile. Although it is statistically significant, the size of the coefficient suggests that outflows following an increase in the VIX at $t-1$ are marginally increased by only 1.14 mm in periods of high speed of mean reversion. Hence, we do detect some evidence that the level of persistence in the VIX impact how investors trade VIX ETPs, but the effect is rather small in economic terms.

IV. Do risk factors matter?

All VIX ETPs in our study produce long-term returns that are very negative. But despite negative returns, the products may still be attractive when adjusting for their systematic risk exposure. For instance, from a CAPM perspective, we would intuitively expect very low returns, given the very negative correlation that the assets have with the market portfolio. The negative returns may thus be regarded as an insurance premium. If so, a positive risk-adjusted return (alpha) would imply that the insurance embedded in the VIX ETP is cheap through the lenses of the CAPM model. In the previous section, we document an inverse relation between the VIX level and flows, suggesting that investors incorporate the typical mean reverting behavior of the VIX in their trading strategy. But it may also be the case that changes in alphas are the driver of flows, meaning that investors buy (sell) when the ETPs are cheap (expensive) relative to their embedded systematic risk exposure. In this

Table IV: Speed of mean reversion

This table shows the results from the panel regression as in Equation (3), where we have divided the sample into periods with high and low speeds of mean reversion, respectively. We define periods of high mean reversion as days where the VIX half-life is in the lower 10 percentile. Panel A presents the results for the period where the speed of mean reversion is high, and Panel B shows the results for the period where the speed of mean reversion is low. Panel C shows the results, where we have augmented the regression with an interaction term between the VIX at $t-1$ and a dummy, which is one if the speed of mean reversion is high. T-statistics in parenthesis are calculated with Newey-West standard errors. ***, and ** represent significance at 1% and 5% respectively.

Panel A: High speed of mean reversion	
VIX_{t-1}	-25.21***
	(-3.57)
Observations	152
$R^2(\%)$	30.55
Panel B: Low speed of mean reversion	
VIX_{t-1}	-23.46***
	(-4.58)
Observations	1373
$R^2(\%)$	17.90
Panel C: Entire period	
VIX_{t-1}	-23.77***
	(-6.10)
$\text{VIX}_{t-1}\times$ Speed of mean reversion	-1.14**
	(-2.41)
Observations	1525
$R^2(\%)$	17.94

section, we examine that hypothesis by taking a revealed preferences approach as in Berk and van Binsbergen (2016) (BvB), Barber et al. (2016) (BHO), and Agarwal et al. (2018) (AGR). That is, we compare the ability of different asset pricing models to explain the flows and thereby try to ascertain which risks factors, if any, are of concern to VIX ETP investors. The studies by BvB, BHO, and AGR suggest that mutual fund- and hedge fund investors are mainly concerned with market risk. If this is also the case for VIX ETP investors, we would expect that flows will react more strongly to CAPM alphas than to the alphas implied by more elaborate models.

IV.A. Asset pricing models

In this section, we are interested in the relation between VIX ETP flows and performance implied by different asset pricing models. By examining a set of candidate models, we can infer the risk model, which is closest to the model that investors use in making their investment decision. In addition to raw returns, we estimate alphas for five different asset pricing models. We consider the CAPM (Sharpe (1964) and Lintner (1965)), the downside CAPM (Hogan and Warren (1974) and Bawa and Lindenberg (1977)), the coskewness CAPM (Harvey and Siddique (2000), Mitton and Vorkink (2007), and Christoffersen et al. (2019)) and the four-factor (4F) model (Carhart (1997)). For the fifth model, we augment the CAPM model with the returns of a (long) straddle. The idea of including this additional factor is to capture a premium for a long volatility exposure. We estimate beta from the downside CAPM only for observations where the excess market return is below zero. Further, we follow Harvey and Siddique (2000), Mitton and Vorkink (2007), and Christoffersen et al. (2019) and measure coskewness beta with respect to squared market excess returns. We estimate alphas over the most recent 11 days (half a month) by subtracting the benchmark adjusted returns from the excess returns of the ETPs:

$$\hat{\alpha}_{M,it} = \frac{1}{11} \sum_{j=t-10}^{t} (r_{ij} - \sum_{n=1}^{N} \hat{\beta}_{M,n,it} \times f_{n,j}). \tag{9}$$

Here M is the asset pricing model, N is the number of factors, r_{ij} is the return of ETP i on day j in excess of the daily 1-month T-bill rate, $f_{n,j}$ is a risk factor, and $\hat{\beta}_{M,n,it}$ is an estimated beta for a risk factor in model M. We obtain daily factor returns from Kenneth French's webpage.[5] For the straddle returns, we use the series made publicly available by Travis L. Johnson.[6] For all asset pricing models, we estimate betas using a rolling window of 60 days. Finally, for the raw return measure (r^{raw}), we calculate the 11-days average of the

[5] https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/

[6] https://www.travislakejohnson.com/data.html

returns, not subtracting the risk-free rate.

In Panel A of Table V, we present summary statistics of the estimated VIX ETP betas. For the CAPM, the betas ($\hat{\beta}$) are, on average, negative, -4.60. This illustrates the very negative correlation that VIX ETPs have with equities. Given the negative CAPM betas, we expect the VIX ETPs to have negative returns. The estimated downside betas ($\hat{\beta}^{-}$) are, on average, even more negative, -5.51. This implies that the negative correlation becomes even stronger in down markets. For the 4F model, we find that the VIX ETPs on average load negatively, -0.40, on the momentum factor ($\hat{\beta}^{mom}$). Notably, we find that the load on the size factor ($\hat{\beta}^{smb}$) is, on average positive, 0.35. Harvey and Siddique (2000) find that the size factor, to some extent, proxies for conditional skewness which could explain the positive load, as the returns of VIX ETPs are, in general, positively skewed (e.g., Christensen et al. (2020)). A positive coskewness is also confirmed via the coskewness CAPM, where the load on squared excess market returns (β^{cosk}) on average equals 53.71. Not surprisingly, the loading on the straddle factor is positive, 0.22.

Panel B of Table V provides summary statistics for the performance of the VIX ETPs. The average daily raw return measure is -0.21%, while the median return is -0.29%. This reflects the positive skewness in the return distribution of these products. Similarly, the average daily alpha for the coskewness CAPM is negative, -0.31%. For the CAPM and 4F model, the average daily alpha is much less negative, although still significant, -0.01%, and -0.02%, respectively. For the downside CAPM, the daily alpha becomes positive, 0.05%, and for the straddle model, alpha is zero on average.

Table VI presents a correlation matrix between return measures from our different asset pricing models. We find that raw returns are modestly correlated with alphas from the five asset pricing models. The two lowest correlation coefficients are between raw returns and downside alpha (0.24) and raw returns and coskewness alpha (0.32). We find the largest correlation coefficient between the CAPM and the 4F alphas and the lowest correlations between raw returns and downside CAPM, coskewness, and straddle alphas.

Table V: Summary Statistics

This table presents summary statistics of characteristics, betas, and performance for the six ETPs in our sample across 8001 ETP-day observations in the period October 2012 to February 2019. The data include only VIX ETPs with direct exposure to the VIX futures. $\hat{\beta}$ is CAPM beta and $\hat{\beta}^{-}$ is downside CAPM beta. $\hat{\beta}^{mkt}$, $\hat{\beta}^{mom}$, $\hat{\beta}^{smb}$, and $\hat{\beta}^{hml}$ are factor betas from the Fama-French-Carhart 4F model. $\hat{\beta}^{cosk}$ is the load on the second factor from the coskewness CAPM model. $\hat{\beta}^{straddle}$ is the load on straddle returns from the straddle model. Note that beta and performance estimates are stated using daily values.

	Mean	Std. Dev	Median
Panel A: VIX ETP betas			
$\hat{\beta}$	-4.60	2.99	-3.85
$\hat{\beta}^{-}$	-5.51	4.04	-4.48
$\hat{\beta}^{mkt}$	-4.67	3.08	-3.81
$\hat{\beta}^{mom}$	-0.40	0.76	-0.29
$\hat{\beta}^{smb}$	0.35	0.81	0.22
$\hat{\beta}^{hml}$	-0.22	1.17	-0.19
$\hat{\beta}^{cosk}$	53.71	94.82	29.93
$\hat{\beta}^{straddle}$	0.11	0.08	0.10
Panel B: VIX ETP performance measures			
	Mean	Std. Dev	Median
r^{raw}	-0.21%	1.46%	-0.29%
$\hat{\alpha}^{capm}$	-0.01%	0.80%	-0.05%
$\hat{\alpha}^{downside}$	0.05%	0.90%	0.00%
$\hat{\alpha}^{4F}$	-0.02%	0.77%	-0.05%
$\hat{\alpha}^{cosk}$	-0.31%	0.78%	-0.22%
$\hat{\alpha}^{straddle}$	0.00%	0.01%	0.00%

IV.B. Estimation of the flow-alpha relation

We examine the relationship between performance and flows by estimating the model:

$$F_{it} = a + b \times \hat{\alpha}_{M,it-1} + c \times X_{it-1} + u_{it}, \tag{10}$$

where F_{it} is the daily ETP flow estimated via Equation (1), $\hat{\alpha}_{M,it-1}$ is the alpha estimated from model M by Equation (9), and X_{it-1} is a vector of control variables. As controls, we include the lagged ETP flow from day t-1, the day t-1 ETP bid-ask spread in basis points, and the t-1 ratio of price to net-asset-value (NAV). We calculate t-statistics using robust standard errors by double clustering by ETP and day.

Table VII reports the results of the panel regression from Equation (10). Columns (1)-

Table VI: Correlations of Performance measures

This table shows the correlation coefficients between performance measures from raw returns, CAPM, downside CAPM, Fama-French-Carhart 4F model, coskewness CAPM, and straddle model.

	r^{raw}	$\hat{\alpha}^{capm}$	$\hat{\alpha}^{downside}$	$\hat{\alpha}^{4F}$	$\hat{\alpha}^{cosk}$	$\hat{\alpha}^{straddle}$
r^{raw}	1.00	0.52	0.24	0.56	0.32	0.41
$\hat{\alpha}^{capm}$		1.00	0.86	0.95	0.79	0.84
$\hat{\alpha}^{downside}$			1.00	0.86	0.75	0.86
$\hat{\alpha}^{4F}$				1.00	0.75	0.78
$\hat{\alpha}^{cosk}$					1.00	0.76
$\hat{\alpha}^{straddle}$						1.00

(6) present results for the raw return measure, the CAPM, downside CAPM, 4F model, coskewness CAPM, and the straddle model, respectively. We only find a significant flow-performance relation for the raw return measure and the downside CAPM, with the downside CAPM estimate being significant only at the 10% level. For the raw return measure, we find that a one percentage point increase in a VIX ETP's raw returns is associated with a 0.72 percentage point decrease in flows on the following day. In Section III, we have documented that a rise in the VIX is followed by a decrease in aggregated dollar flows. Since increases in the VIX are associated with positive returns of VIX ETPs, the inverse relation between raw returns and flows comes as no surprise. For the downside CAPM, a one percentage point increase in alpha leads to a 0.49 percentage point increase in future flows. The intuition of the downside CAPM model is that investors are concerned about how an asset correlates with the market during downturns. The positive coefficient then implies that investors increase (decrease) their positions when the VIX ETPs offer cheap (expensive) protection against market downturns, as implied by a positive (negative) alpha. Given the insurance-like payoff of the products, this result is intuitively appealing.

For the included control variables, we interestingly find that the coefficient of Price-to-NAV is positive. In principle, this implies that investors tend to buy (sell) the ETPs when they are expensive (cheap) relative to the value of the underlying assets. However, in economic terms the size of the coefficient is very small.

Table VII: Product flow-performance regression

This table presents the results for the panel regression as in Equation (10). The dependent variable is daily percentage VIX ETP flow. Performance measures are the average daily raw return or risk-adjusted returns from the CAPM, downside CAPM, 4F model, coskewness CAPM, and straddle model over the last 11 days calculated as in Equation (9). Control variables are the lagged ETP flow, the ETP spread measured in basis points, and the price-to-NAV ratio. T-statistics in parenthesis are calculated with double clustered standard errors by product and day. ***, **, and * represent significance at 1%, 5%, and 10% respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
r^{raw}	-0.732*** (-4.403)					
$\hat{\alpha}^{capm}$		-0.163 (-0.713)				
$\hat{\alpha}^{Downside}$			0.212* (1.837)			
$\hat{\alpha}^{4F}$				-0.250 (-1.049)		
$\hat{\alpha}^{cosk}$					-0.228 (-0.297)	
$\hat{\alpha}^{straddle}$						-0.042 (-0.189)
Lagged flow	0.056 (1.367)	0.084* (1.773)	0.084* (1.837)	0.084* (1.753)	0.084* (1.790)	0.085* (1.799)
Spread	-0.002*** (-3.224)	-0.002*** (-2.784)	-0.002*** (-2.848)	-0.002*** (-2.760)	-0.002** (-2.500)	-0.002*** (2.864)
Price-to-NAV	0.004*** (2.811)	0.005*** (3.263)	0.005*** (3.387)	0.005*** (3.247)	0.005*** (2.660)	0.005*** (3.472)
Observations	8001	8001	8001	8001	8001	8001

IV.C. Are the results unique to the chosen performance horizon?

In Table VIII, we report the results of the panel regression from Equation (10) for the raw returns and downside CAPM model using performance horizons of 20, 40 days, and 60 days. Changing the performance horizon does not alter the results. Also for the other asset pricing models in our study (CAPM, 4F, coskewness CAPM, and straddle), we obtain almost identical results to those reported previously.

Table VIII: Robustness check - performance horizon

This table presents the results for the panel regression as in Equation (10). Dependent variable is daily percentage VIX ETP flow. Performance measures are the average daily raw return or risk- adjusted returns from the downside CAPM model calculated as in Equation (9) but with different performance horizons (20 days, 40 days, and 60 days). Control variables are the lagged ETP flow, the ETP spread measured in basis points, and the price-to-NAV ratio. T-statistics in parenthesis are calculated with double clustered standard errors by product and day. ***, and * represent significance at 1% and 10%, respectively.

	20 days horizon		40 days horizon		60 days horizon	
r^{raw}	-0.704***		-0.660***		-0.587***	
	(-4.189)		(-4.424)		(-4.178)	
$\hat{\alpha}^{Downside}$		0.148*		0.217*		0.266*
		(1.817)		(1.814)		(1.814)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8001	8001	8001	8001	8001	8001

IV.D. Sign-test

We also investigate the alpha-flow relation via the methodology of BvB who develop a sign-test to explain flows of mutual funds. The test relates the sign of a flow to the sign of performance, given as alpha, implied by an asset pricing model. If investors consider exposure to risk factors then intuitively positive updates of alpha lead to positive flows, and vice versa. The frequency with which positive (negative) alphas generate inflows (outflows) provides a ranking of competing asset pricing models. The model with the highest frequency is the model that best explains flows. Below we provide a description of our implementation of the test.

Let $\phi(\cdot)$ be a function that returns the sign of a real number, taking the value 1 for a positive number, -1 for a negative number, and 0 for zero. For every asset pricing model, we regress the signs of ETP flows on the signs of alpha from an asset pricing model and obtain the coefficient $B_{F\alpha}$, given by:

$$B_{F\alpha} = \frac{cov(\phi(F_{it}), \phi(\hat{\alpha}_{M,it-1}))}{var(\phi(\hat{\alpha}_{M,it-1}))}, \tag{11}$$

where M refers to the asset pricing model. To avoid look-ahead bias, we use lagged alphas and contemporary flows as in AGR. From Equation (11), we calculate the average probability of future flows being positive (negative) conditional on past alphas being positive (negative) as $(1+B_{F\alpha})/2$. An average probability above (below) 50% indicates that positive updates of performance predicts positive updates of flows. If future flows and past alphas are unrelated, then $B_{F\alpha} = 0$, and the average probability becomes 50%.

In order to compare models M and K, we use the following regression:

$$\phi(F_{it}) = \gamma_0 + \gamma_1 \times \left(\frac{\phi(\hat{\alpha}_{M,it-1})}{var(\phi(\hat{\alpha}_{M,it-1}))} - \frac{\phi(\hat{\alpha}_{K,it-1})}{var(\phi(\hat{\alpha}_{K,it-1}))} \right) + \zeta_{it}. \tag{12}$$

If γ_1 is positive, it implies that model M better explains subsequent flows than model K. This statement is conditional on the hypothesis that there is a positive relation between flows and alpha. For the regressions in Equation (11) and Equation (12), we calculate t-statistics with robust standard errors by double clustering by ETP and day.

Table IX presents results from Equations (11) and (12). The first column presents estimates of $B_{F\alpha}$. For all models, the coefficient is negative. That is, negative performance predicts inflows across all models and vice versa. The second column presents the average probability of the flows being positive (negative) conditional on past performance being positive (negative) for each asset pricing model. We find that the raw return measure has the lowest average probability, 43.93%, to predict future positive flows conditional on positive past returns. This is equivalent to an average probability of 56.07% to predict positive flows conditional on negative past returns. For the downside CAPM, the coskewness CAPM, and the straddle model we cannot reject that $(1 + B_{F_\alpha})/2$ is different from 50%, which suggests that these models have no predictive power.

If VIX ETP investors consider exposure to risk factors, then we would expect a significant positive relation between alpha and flows, which we do not find. Hence, these results do not

Table IX: Sign-test

This table provides the results of the sign-test as in Berk and van Binsbergen (2016). $\hat{B}_{F\alpha}$ is the coefficient estimate from regressing the signs of daily flows on the signs of daily performance as in Equation (11). The percentage of signed flows explained by signed performance is obtained as $(1+\hat{B}_{F\alpha})/2$. Tests for pairwise comparisons of models as in Equation (12) are provided in the last five columns. T-statistics given in parenthesis are calculated with double-clustered standard errors by product and day. *** and ** represent significance levels of 1% and 5%, respectively.

Model	$\hat{B}_{F\alpha}$	% Flow explained	CAPM	Downside CAPM	4F	Coskewness	Straddle
Raw return	-0.122*** (-2.919)	43.93%	-0.065** (-2.066)	-0.074** (-2.608)	-0.075** (-2.380)	-0.062** (-2.474)	-0.072** (-2.536)
CAPM	-0.038*** (-2.974)	48.08%		-0.087*** (-4.008)	-0.021 (-0.947)	-0.021 (-0.972)	-0.039 (-2.450)
Downside CAPM	-0.007 (-0.489)	49.64%			0.043 (1.948)	0.018 (0.768)	0.006 (0.347)
4F	-0.031** (-2.433)	48.45%				-0.009 (-0.552)	-0.023 (-1.497)
Coskewness	-0.023 (-1.362)	48.84%					-0.012** (-0.652)
Straddle	-0.012 (-0.743)	49.41%					

support that hypothesis. In the last five columns, we provide pairwise comparisons of the asset pricing models with estimates of γ_1 from Equation (12). Comparing raw returns against the other models, we see that γ_1 is estimated to be significantly below 0. Given the negative relation that we have estimated between flows and alpha, this result is of minor interest as it only suggest that negative raw returns is best at explaining subsequent positive flows.

V. Flows and the VIX premium puzzle

In this final empirical section, we link our findings to the study by Cheng (2019) (IHC). IHC defines the VIX premium at time t with horizon $T - t$ as:

$$VIXP_t^T \equiv E_t^Q[VIX_T] - E_t^P[VIX_T], \tag{13}$$

where the physical expectation is given by the estimate from some statistical model, and the risk neutral expectation is the VIX futures price with maturity T.[7] This premium can be interpreted as the expected dollar loss for a long VIX futures position with $1 notional value held through maturity. Interestingly, there is a negative or flat relation between changes in the VIX premium and changes in various risk measures (e.g., realized volatility, the VIX, and SPX Skew). Hence, when risk increases, the VIX premium tends to decrease or stay flat before increasing at a later point in time. This empirical pattern is in IHC labeled as the "low premium response" puzzle. Apparently, the low premium response is tradable. A short investor who sees her estimated premiums falling can close her position and sidestep ex-post low-profit high-risk situations.

We hypothesize that the low premium response puzzle can, at least partially, be explained by the large VIX ETP outflows that we have documented tend to follow after an increase in the VIX. Brøgger (2019) and Todorov (2020) show how the hedge demand of issuers of futures-based ETPs is increasing (decreasing) in inflows (outflows). Thus, large capital outflows decrease the hedging requirement for the issuer of the ETP, who will reduce the position in VIX futures. This reduction in the issuer's position will, all else equal, put downward pressure on the futures prices. Consider an example for illustration: Due to some perception of increased risk in the equity markets, the VIX increases. Following this increase, investors on a large scale reduce their positions in VIX ETPs. Consequently, the issuers of VIX ETPs reduce their inventory in VIX futures, which puts downward pressure on prices. The expected future VIX then increases less under the risk-neutral measure than under the physical measure, making the VIX premium, as defined by Equation (13), negative.

To explore our hypothesis, we first estimate VIX premiums at different horizons (one-month to seven-month). We follow the same procedure as IHC and estimate $E_t^P[VIX_T]$ with an ARMA(2,2) model, and we obtain $E_t^Q[VIX_T]$ as the price of a futures contract, which

[7] IHC use an ARMA(2,2) model for the physical expectation but makes a robustness check with several other models.

is rolled over on the last trading day of a month.[8] For our time T forecast of $E_t^P[VIX_T]$, we estimate the ARMA(2,2) model using all data (daily VIX close from 1990) through date t. Figure VI plots our estimates of the one-month, four-month, and seven-month VIX premium. For all horizons, the premiums are mainly positive. However, large downward movements occur at the end of August 2015 (China's renminbi devaluation), the end of June 2016 (Brexit referendum), and the beginning of February 2018 ("Volmageddon"). The downward movements are especially large for the one-month premium. The premiums are highly correlated but smaller in absolute value for larger horizons.

Figure VI: VIX Premiums

Figure VI shows our estimated VIX premiums defined as the difference between the expected future value of the VIX under the risk-neutral and physical measure, respectively. For estimating these, we follow Cheng (2019). That is, the expected value under the risk-neutral measure is estimated as the price of a futures contract, which is rolled on the last trading day of a month. The expected value under the physical measure is estimated via an ARMA (2,2) model, where we use all data (daily VIX close from 1990) through date t.

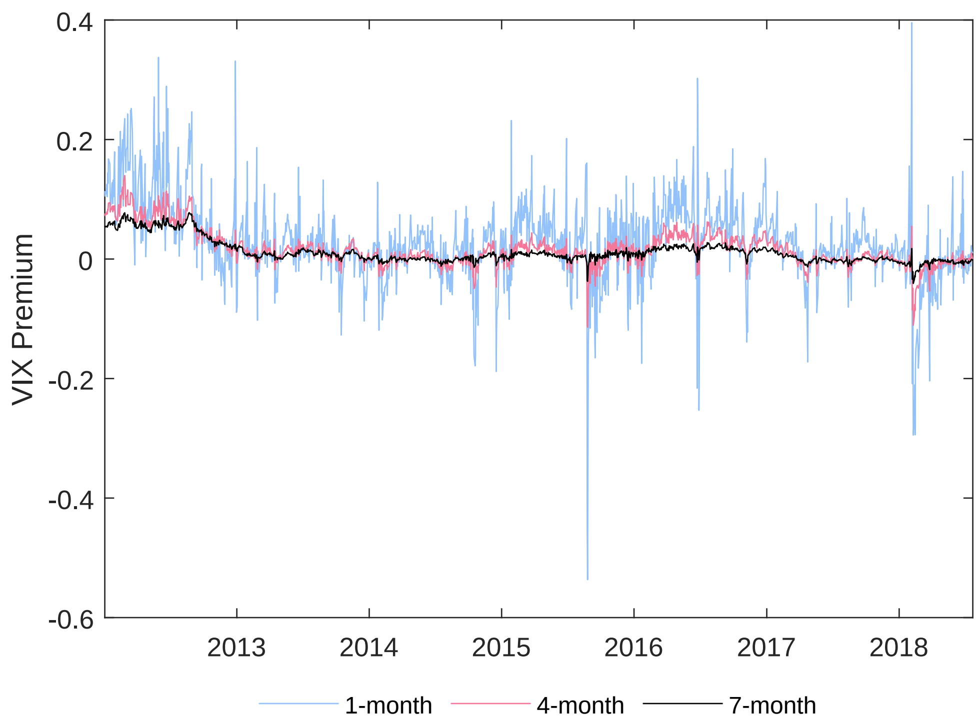

[8]The one-month premiums estimated by IHC along with a description of the estimation procedure are made available at: `http://www.dartmouth.edu/~icheng/`

V.A. Simple regression

As an initial examination of the relation between premiums and flows, we estimate the simple model of Equation (14):

$$\Delta VIXP_t^h = a + b \times \Delta AggDF_t + c \times \Delta X_t + u_t. \tag{14}$$

Here, $VIXP_t^h$ is the estimated horizon-h VIX premium at time t, $AggDF_t$ is the aggregate dollar-flow for all products in our sample as defined in Equation (2), and X_t is a vector of control variables. As controls, we include the default spread (between Moody's BAA and AAA corporate bond yields) and the term spread (between the ten-year T-bond and the three-month T-bill yields), both obtained from the website of the Federal Reserve Bank of St. Louis. We also include the P/E ratio for the S&P 500 index, the WTI-crude oil price, the gold price, and the USD/JPY FX rate, all obtained via Bloomberg. For several of our control variables, we fail to reject the null hypothesis of a unit root, which is why we apply the first difference transform in Equation (14).

The event study in Section III.A shows that flows in short-term products follow a different pattern than flows in mid-term products. Hence, if our hypothesis is true, we expect the relationship between the VIX premium and flows to be different for short-term and mid-term VIX ETPs. Therefore, we estimate Equation (14) for short-term and long-term products separately. For ease of interpretation, we standardize all variables to have zero mean and unit standard deviation.

Panel A in Table X shows the regression output for short-term products. The regression coefficients for the changes in flows are displayed along the columns for the different VIX premium horizons. T-statistics (reported in parenthesis) are calculated with Newey-West standard errors. We see that there is a positive relation between changes in flows in short-term products and changes in VIX premiums. For example, a one standard deviation increase in $\Delta AggDF$ leads to a 0.331 standard deviation increase in $\Delta VIXP$. The coefficients are highly significant at all horizons. However, we also see that the relationship is strongest for the first two VIX premiums and decreases for longer horizons. This is the pattern we expect to see if hedging activity of short-term products impacts VIX premiums. This is the case since hedging activity of short-term products occurs from trades in the futures contracts at the short end of the term structure. Hence, short-dated futures contracts are exposed to the largest pricing impact. Most of the VIX products in our sample are ETN's, which are not required to hold the underlying assets. This means that they are free to hedge their exposure as they desire (or not hedge at all). It is possible that ETN issuers are hedging some of the exposure using futures further out on the term structure, which could explain the positive relation between short-term product flows and longer horizon VIX premiums.

Panel B shows the regression output for mid-term products. The regression coefficients of $\Delta AggDF$ are insignificant at all horizons. This implies that the hedging activity of mid-term products cannot explain the low premium-response puzzle. This fits well with the findings in our event study in Section III.A that investors in mid-term products are less prone to sell during high levels of the VIX. Hence, the issuers will be less likely to reduce their hedge position in VIX futures when market risk is elevated.

V.B. Relation in the quantiles

From IHC, we know that the VIX premium is lowest when risk is high, which is also when large outflows tend to occur. Then, if the low premium response puzzle is explained by the large outflows following increases in the VIX, we would expect the regression coefficient of aggregated flows in Equation (14) to be higher at the lowest quantiles of the distribution of VIX premiums. We explore whether this is the case via a quantile regression (see Koenker and Bassett (1978)), with the criterion function to minimize given by:

$$Q_T(\beta_q) = \sum_{t:y_t \geq x_t'\beta}^{T} q|y_t - x_t'\beta_q| + \sum_{t:y_t < x_t'\beta}^{T} (1-q)|y_t - x_t'\beta_q|. \tag{15}$$

Table X: VIX Premium puzzle

This table shows the results of the regression in Equation (14). The dependent variable is changes in the horizon-h VIX premium. The explanatory variable of interest is changes in dollar-flow aggregated across products. Panel A contains the results where we only include flows from short-term products. Panel B contains results where we only include mid-term products. In both cases, control variables are the default spread(between Moody's BAA and AAA corporate bond yields), the term spread (between the ten-year T-bond and the three-month T-bill yields), the S&P 500 P/E ratio, the WTI-crude oil price, the gold price, and the USD/JPY FX rate, all in terms of differences. T-statistics given in parenthesis are calculated with Newey-West standard errors. *** represents significance at 1%.

Panel A: Short-term products							
	$\Delta VIXP_t^1$	$\Delta VIXP_t^2$	$\Delta VIXP_t^3$	$\Delta VIXP_t^4$	$\Delta VIXP_t^5$	$\Delta VIXP_t^6$	$\Delta VIXP_t^7$
$\Delta AggDF_t$	0.331***	0.338***	0.310***	0.279***	0.273***	0.296***	0.254***
	(5.161)	(4.973)	(5.271)	(4.850)	(4.905)	(4.894)	(4.978)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1352	1352	1352	1352	1352	1352	1350
$R^2(\%)$	29.64	29.54	27.15	25.37	26.49	29.01	28.60

Panel B: Mid-term products							
	$\Delta VIXP_t^1$	$\Delta VIXP_t^2$	$\Delta VIXP_t^3$	$\Delta VIXP_t^4$	$\Delta VIXP_t^5$	$\Delta VIXP_t^6$	$\Delta VIXP_t^7$
$\Delta AggDF_t$	-0.000	0.016	0.012	0.006	0.018	0.032	0.002
	(-0.014)	(0.255)	(0.232)	(0.129)	(0.428)	(0.663)	(0.267)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observation	1352	1352	1352	1352	1352	1352	1352
$R^2(\%)$	18.73	18.15	17.59	17.63	19.11	20.37	22.19

Here, we define $y = \Delta VIXP_t^h$ and $x_t'\beta = a + b \times \Delta AggDF_t + c \times \Delta X_t$, which is similar to Equation (14). In Figure VII, we plot the slope estimate for $\Delta AggDF$ for quantiles in the range $q \in (0.01 : 0.01 : 0.99)$ with corresponding bootstrapped 95% confidence bands in grey. We obtain slope estimates using one-month VIX premiums and aggregate flows for short-term products. For comparison, we also plot the OLS slope estimate from Equation (14) with 95% confidence bands. From Figure VII, we first note that there is a significant and positive regression coefficient across all quantiles. Secondly, we see a prominent U-shape in the relation between VIX premiums and short-term product flows. That is, the relation between flows and the VIX premium is strongest at extreme quantiles, particularly the lower quantiles which we know coincide with high levels of risk and large outflows.

Figure VII:
Quantile regression: One-month VIX premium and short-term VIX ETP flows

Figure VII shows slope estimates for $\Delta AggDF$ from Equation (15) for quantiles in the range $q \in (0.01 : 0.01 : 0.99)$ with corresponding bootstrapped 95% confidence bands in grey. We use 1,000 bootstrap resamples to obtain confidence bands. Slope estimates are obtained using one-month VIX premiums and aggregate flows for short-term VIX products. For comparison, we also plot the OLS slope estimate from Equation (14) with 95% confidence bands in pink.

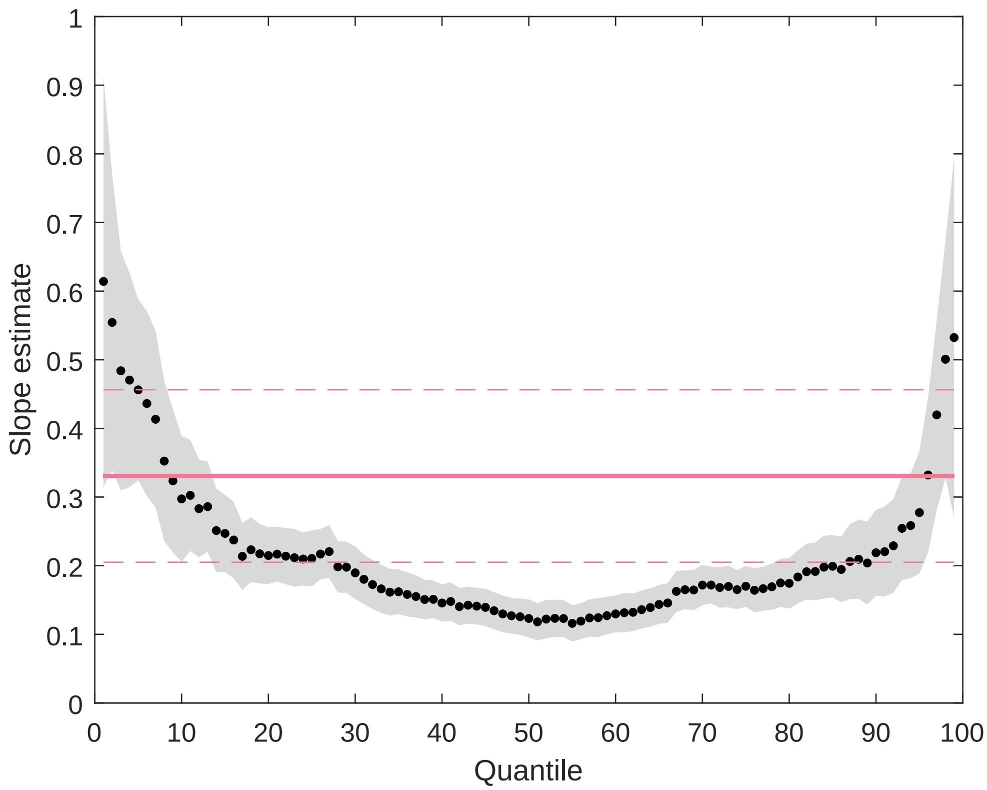

V.C. Bivariate dynamics

The estimates in Table X do not allow for the dynamics of premiums to feedback into flows. We allow for this by estimating a three-lag vector autoregression (VAR) system of aggregated flows and the one-month VIX premium, both in levels.[9] Both variables are again standardized and we order flows first to assess the contemporaneous impact of flows on the VIX premium. Figure VIII depicts the impulse response functions of the VAR system along with the bootstrapped 95% confidence bands. The shocks are orthogonalized, and the time horizon is days after a shock. The VIX premium increases on the impact of a shock to the aggregated flows. Despite a small increase at day 3 following the shock, the effect decreases monotonically and is practically negligible after 10 days. A shock to the VIX premium has no significant impact on flows at any horizon. In Appendix B, we provide plots of the impulse response functions where we have included the VIX in the VAR system and also where we have applied the reverse ordering. In both cases, we find that a shock to aggregated flows impacts the one-month VIX premium.

In sum, the empirical results of this section suggest that the flow pattern in VIX ETPs, at least to some degree, may explain why the VIX premium tend to become negative when risk is elevated.

A final note; a negative VIX premium positively predicts VIX futures returns. Hence, an implication of our hypothesis being true is also that VIX ETP investors predict the VIX to revert faster than it does ex-post. Put another way, when the VIX increases, it is much more persistent than investors expect, and they end up selling volatility too cheap. This also fits well with Fernandez-Perez et al. (2019), who find that the hedging demand by VIX ETPs increases deviations in the actual futures prices from the efficient futures prices.

[9] See Appendix B for details on lag order selection.

Figure VIII: AggDF and VIXP impulse responses

Figure VIII shows the responses down the rows to the shocks of AggDF and the one-month VIX premium. Both series are standardized to have zero mean and a unit standard deviation. The shocks are orthogonalized with AggDF ordered first. The dashed lines are 95% confidence bands based on bootstrapped standard errors with 1,000 bootstrap resamples.

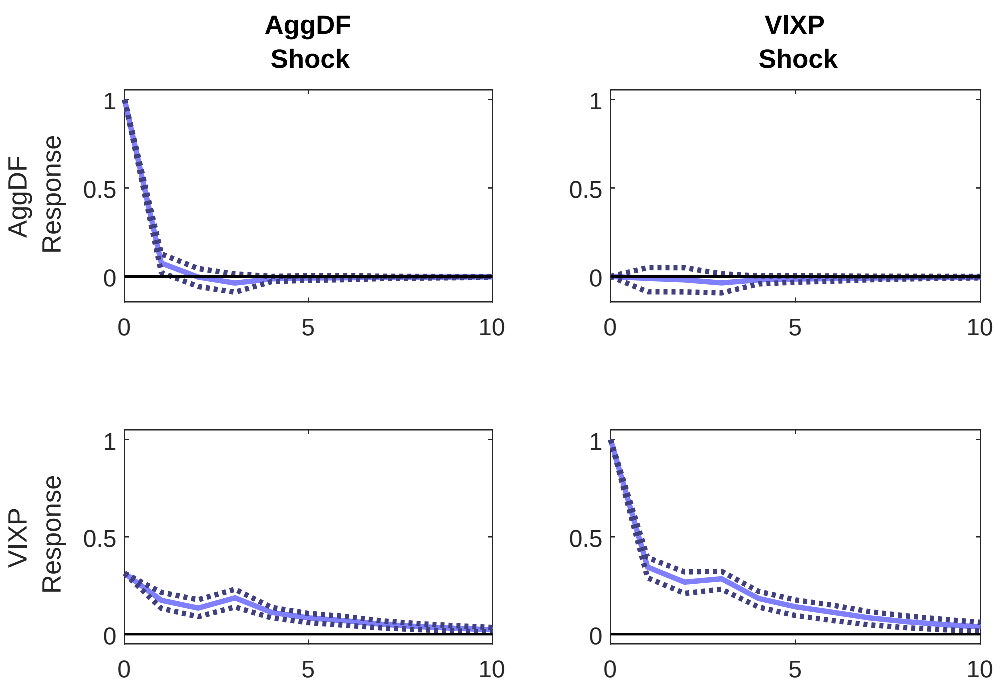

VI. Conclusion

VIX ETPs with long volatility exposure have become very popular. The main purpose of this paper is to examine how these products are applied by investors.

From a regression of aggregated dollar flows on the VIX and its lags, we find that an increase in the VIX is followed by outflows. This inverse relation is further examined in an event study of flows around the largest VIX increases, which shows very large outflows at elevated VIX levels. The documented flow pattern is consistent with investors incorporating the typical mean reversion of volatility in how they trade VIX ETPs. We estimate the speed of mean reversion as the half-life of the distance to a long-run level of the VIX, and we find that the inverse relation between flows and the VIX is slightly amplified in periods with a high speed of mean reversion.

Using the framework of Barber et al. (2016), Berk and van Binsbergen (2016), and Agarwal et al. (2018), we investigate whether investors adjust for exposure to risk factors when they invest in VIX ETPs. Across the different asset pricing models that we consider, none of the models explain flows better than the simple raw returns. Hence, we do not find clear evidence of investors adjusting for risk factors.

Finally, we provide a possible explanation for the low premium response puzzle documented in Cheng (2019). By regressing the changes of VIX premiums on changes in flows, we find a significant and positive relation. In line with our prediction, this only holds for short-term products, and the effect is decreasing for VIX premiums at longer horizons. From a quantile regression, we show that the positive relation between flows and the VIX premium increases at the extreme levels of the premium. We examine the bivariate dynamics by a vector autoregression of the aggregated dollar flows and the VIX premium. From the impulse response functions, a shock to the aggregated flows increases the VIX premium, but a shock to the VIX premium has no effect on flows. In sum, these results indicate that the flow pattern in VIX ETPS, at least to some degree, may explain why the VIX premium tends to become negative when risk is elevated.

References

Adrian, T., R. K. Crump, and E. Vogt (2019). Nonlinearity and flight-to-saftey in the risk-return trade-off for stocks and bonds. *Journal of Finance 74*(4), 1931–1973.

Agarwal, V., T. C. Green, and H. Ren (2018). Alpha or beta in the eye of the beholder: What drives hedge fund flows? *Journal of Financial Economics 127*(3), 417 – 434.

Baele, L., G. Bekaert, K. Inghelbrecht, and M. Wei (2019). Flights to saftey. *Review of Financial Studies 33*(2), 689–746.

Barber, B. M., X. Huang, and T. Odean (2016). Which factors matter to investors? evidence from mutual fund flows. *Review of Financial Studies 29*(10), 2600–2642.

Bawa, V. S. and E. B. Lindenberg (1977). Capital market equilibrium in a mean-lower partial moment framework. *Journal of Financial Economics 5*, 189–200.

Bekaert, G. and M. Hoerova (2014). The vix, the variance premium and stock market volatility. *Journal of Econometrics 183*(2), 181 – 192.

Berk, J. B. and J. H. van Binsbergen (2016). Assessing asset pricing models using revealed preference. *Journal of Financial Economics 119*, 1–23.

Berkowitz, J. P. and J. DeLisle (2018). Volatility as an asset class: Holding vix in a portfolio. *Journal of Alternative Investments 21*(2), 52–64.

Bollen, N. P. B., M. J. O'Neil, and R. E. Whaley (2017). Tail wags dog: Intraday price discovery in vix markets. *Journal of Futures Markets 37*(5), 431–451.

Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected stock returns and variance risk premia. *Review of financial studies 22*(11), 4463–4492.

Bordonado, C., P. Molnár, and S. R. Samdal (2017). Vix exchange traded products: Price discovery, hedging, and trading strategy. *Journal of Futures Markets 37*(2), 164–183.

Brøgger, S. B. (2019). The market impact of uninformed flows: Evidence from the vix futures market. *Working paper*.

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance 52*(1), 57–82.

Cheng, I.-H. (2019). The vix premium. *Review of Financial Studies 32*(1), 180–227.

Christensen, K., C. Christiansen, and A. M. Posselt (2020). The economic value of vix ETPs. *Journal of Empirical Finance 58*, 121 – 138.

Christoffersen, P., M. Fournier, K. Jacobs, and M. Karoui (2019). Option-based estimation of the price of co-skewness and co-kurtosis risk. *Journal of Financial and Quantitative Analysis, forthcoming*.

Fernandez-Perez, A., B. Frijns, A. Tourani-Rad, and R. I. Webb (2019). Does increased hedging lead to decreased price efficiency? the case of vix etps and vix futures. *Financial Review 54*, 477–500.

Gehricke, S. A. and J. E. Zhang (2018). Modeling vxx. *Journal of Futures Markets 38*(8), 958–976.

Harvey, C. R. and A. Siddique (2000). Conditional skewness in asset pricing tests. *Journal of Finance 55*(3), 1263–1295.

Hogan, W. W. and J. M. Warren (1974). Toward the development of an equilibrium capital-market model based on semivariance. *Journal of Financial and Quantitative Analysis 9*(1), 1–11.

Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica 46*(1), 33–50.

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics 47*(1), 13–37.

Longstaff, F. A. (2004). The flight-to-liquidity premium in u.s. treasury bond prices. *Journal of Business 77*(3), 511–526.

Mitton, T. and K. Vorkink (2007, 01). Equilibrium underdiversification and the preference for skewness. *The Review of Financial Studies 20*(4), 1255–1288.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance 19*(3), 425–442.

Shu, J. and J. E. Zhang (2012). Causality in the vix futures market. *Journal of Futures Markets 32*(1), 24–46.

Todorov, K. (2020). Passive funds actively affect prices: Evidence from the largest etf markets. *Working Paper*.

Whaley, R. E. (2009). Understanding the vix. *Journal of Portfolio Managment 35*(3), 98–105.

Zhang, J. E. and Y. Huang (2010). The cboe s&p 500 three-month variance futures. *Journal of Futures Markets 30*(1), 48–70.

Zhang, J. E. and Y. Zhu (2006). Vix futures. *Journal of Futures Markets 26*(6), 521–531.

A. VIX ETPs

Table XI: VIX ETP overview

This table provides an overview of the VIX ETPs (inverse products not included) that have been active during our sample period. The Average AUM is the average calculated over our sample period. All products, except EVIX, track either the S&P 500 VIX short- or mid-term futures indexes, indicated by either ST or MT. EVIX tracks VSTOXX Short-Term Futures Investable Index. The notation TR and ER denotes the total return and the excess return version of the indexes. The leverage factor of the ETP is applied to the daily return of the index that it tracks. A leverage ratio of 1 means that the ETP promises the daily rate of return on the underlying index. A leverage factor of 2 means that the ETP promises twice the return of the index. The expense ratio is an annual management fee and is charged on a daily basis. The last column indicates whether we have included the ETP in our sample.

Synbol	Name	Date of inception	Average AUM (million $)	ST/MT	Leverage factor	Expense ratio	Included
VXX	iPath S&P 500 VIX Short-Term Futures ETN	01/29/2009	1103.10	ST	1	0.89	Yes
VXZ	iPath S&P 500 VIX Mid-Term Futures ETN	01/29/2009	48.35	MT	1	0.89	Yes
VIXM	ProShares VIX Mid-Term Futures ETF	01/03/2011	41.72	MT	1	0.85	Yes
VIIZ	VelocityShares VIX Medium-Term ETN	11/29/2010	1.11	MT	1	0.89	No
EVIX	VelocityShares 1x Long VSTOXX Futures ETN	05/02/2017	10.23	ST	1	1.35	No
VMAX	REX VolMAXX Long VIX Futures Strategy ETF	05/03/2016	2.79	ST	1	1.25	No
VIXY	Proshares VIX Short-Term Futures ETF	01/03/2011	153.65	ST	1	0.85	Yes
VIIX	VelocityShares VIX Short-Term ETN	11/29/2010	11.94	ST	1	0.89	No
TVIX	VelocityShares Daily 2x VIX Short-Tern ETN	11/29/2010	340.88	ST	2	1.65	Yes
UVXY	ProShares Ultra VIX Short-Term Futures ETF	10/04/2011	421.03	ST	1.5	1.65	Yes
TVIZ	VelocityShares Daily 2x VIX Mid-Term ETN	11/29/2010	2.56	MT	2	1.65	No

B. Flows and the VIX Premium

Table XII: VAR order selection

This table reports the log-likelihood (Log LH), the Bayesian information criterion (BIC), and Akaike's information criterion (AIC) for VAR models of AggDF and the one-month VIX premium, of lag order 1-5.

	VAR(1)	VAR(2)	VAR(3)	VAR(4)	VAR(5)
Log LH	-3808.91	-3764.19	-3715.09	-3730.29	-3706.05
BIC	7661.51	7601.21	7561.24	7562.52	7572.28
AIC	7629.81	7548.38	7466.18	7488.57	7456.10

Figure IX: VIX Premiums

Figure IX shows the responses down the rows to the shocks of the VIX, AggDF, and the one-month VIX premium. All series are standardized to have zero mean and a unit standard deviation. The shocks are orthogonalized, with the VIX ordered first. The dashed lines are 95% confidence bands based on bootstrapped standard errors with 1,000 bootstrap resamples.

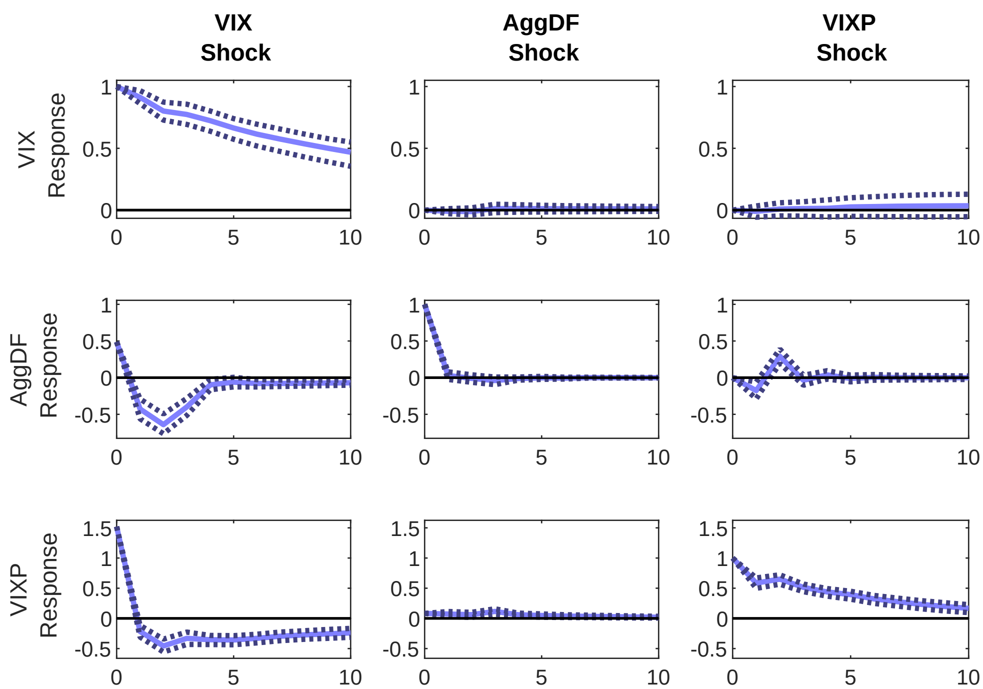

Figure X: VIXP and AggDF impulse responses

Figure X shows the responses down the rows to the shocks of AggDF and the one-month VIX premium. Both series are standardized to have zero mean and a unit standard deviation. The shocks are orthogonalized, with the VIX premium ordered first. The dashed lines are 95% confidence bands based on bootstrapped standard errors with 1,000 bootstrap resamples.

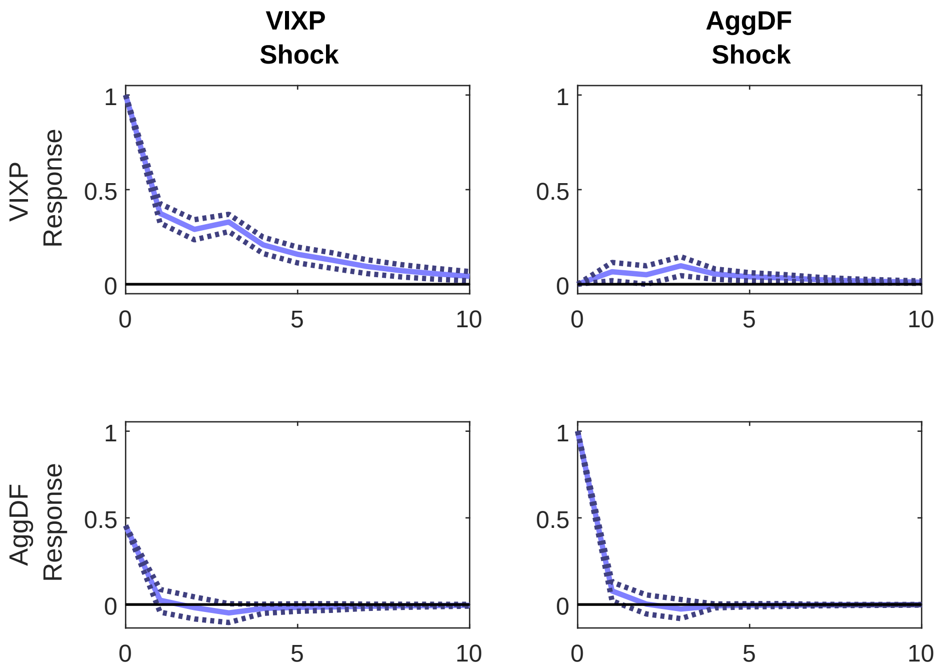

Research Papers

CREATES
Center for Research in Econometric Analysis of Time Series

2021-06:	Gloria González-Rivera, Carlos Vladimir Rodríguez-Caballero and Esther Ruiz Ortega: Expecting the unexpected: economic growth under stress
2021-07:	Matei Demetrescu and Robinson Kruse-Becher: Is U.S. real output growth really non-normal? Testing distributional assumptions in time-varying location-scale models
2021-08:	Luisa Corrado, Stefano Grassi and Aldo Paolillo: Modelling and Estimating Large Macroeconomic Shocks During the Pandemic
2021-09:	Leopoldo Catania, Alessandra Luati and Pierluigi Vallarino: Economic vulnerability is state dependent
2021-10:	Søren Johansen and Anders Ryghn Swensen: Adjustment coefficients and exact rational expectations in cointegrated vector autoregressive models
2021-11:	Bent Jesper Christensen, Mads Markvart Kjær and Bezirgen Veliyev: The incremental information in the yield curve about future interest rate risk
2021-12:	Mikkel Bennedsen, Asger Lunde, Neil Shephard and Almut E. D. Veraart: Inference and forecasting for continuous-time integer-valued trawl processes and their use in financial economics
2021-13:	Anthony D. Hall, Annastiina Silvennoinen and Timo Teräsvirta: Four Australian Banks and the Multivariate Time-Varying Smooth Transition Correlation GARCH model
2021-14:	Ulrich Hounyo and Kajal Lahiri: Estimating the Variance of a Combined Forecast: Bootstrap-Based Approach
2021-15:	Salman Huseynov: Long and short memory in dynamic term structure models
2022-01:	Jian Kang, Johan Stax Jakobsen, Annastiina Silvennoinen, Timo Teräsvirta and Glen Wade: A parsimonious test of constancy of a positive definite correlation matrix in a multivariate time-varying GARCH model
2022-02:	Javier Haulde and Morten Ørregaard Nielsen: Fractional integration and cointegration
2022-03:	Yue Xu: Spillovers of Senior Mutual Fund Managers' Capital Raising Ability
2022-04:	Morten Ørregaard Nielsen, Wonk-ki Seo and Dakyung Seong: Inference on the dimension of the nonstationary subspace in functional time series
2022-05:	Kristoffer Pons Bertelsen: The Prior Adaptive Group Lasso and the Factor Zoo
2022-06:	Ole Linnemann Nielsen and Anders Merrild Posselt: Betting on mean reversion in the VIX? Evidence from ETP flows