Volatility markets underreacted to the early stages of the COVID-19 pandemic

Ing-Haw Cheng[1]

April 2020

Abstract

In late February and early March 2020, VIX futures prices were too low and observably undervalued in real-time, even as COVID-19 pandemic risks were growing. An investor who traded based on real-time signals of undervaluation would have earned significant trading profits by taking a long position in VIX futures in late February and holding it over March as the VIX reached record highs. The underreaction of VIX futures prices to growing risks was a vivid example of broader patterns in the VIX futures market. A trading strategy based on the proposed valuation signal generates positive risk-adjusted returns.

[1] Tuck School of Business, Dartmouth College, 100 Tuck Mall, Hanover NH 03755, USA. Email: inghaw.cheng@dartmouth.edu. Web: http://www.dartmouth.edu/~icheng. Key data series are available on the author's website. I am grateful to Jonathan Lewellen and Juhani Linnainmaa for helpful discussions.

This article provides evidence that the VIX futures market underreacted to the growing risks of the COVID-19 pandemic during the pandemic's early stages. An investor who went long futures in late February or early March based on trading signals available in real-time would have made significant trading profits from anticipating the subsequent spike in stock market volatility and worsening of the pandemic.

A simple example illustrates the logic the article will follow. On March 2, 2020, the VIX stood at 33. By this date, coronavirus cases were spiking in Europe, and the U.S. had reported possible community spread as well as its first coronavirus-related death. The S&P 500 had fallen to just under 3100.

With the VIX at 33, the VIX futures contract expiring March 18 settled at a futures price of 26, suggesting that the market expected the VIX to fall. The VIX tends to move predictably back towards its long-term average, which is around 20; in recent years, the average VIX has been even lower. A futures price of 26 thus does not seem unreasonable.

But precisely because the VIX predictably moves back towards its long-term average, we can also ask: On March 2, what would a statistical model have forecasted for the VIX as of March 18? The answer, as I discuss below, was higher – a value just exceeding 30.

A futures price below the fair statistical forecast suggests that futures were too cheap and undervalued. Futures prices should (and typically do) exceed statistical forecasts because long VIX futures investors pay a premium over the fair statistical forecast to hedge uncertainty and market downturns.

VIX futures continued to remain undervalued in early March even as news about the pandemic grew much worse. For example, On March 12, a day after the WHO declared the coronavirus outbreak a pandemic, the VIX jumped from 54 to 75, a gain of 21 points. The March futures price, with just days until expiration, gained only 12 points to settle at 58. This price was 7 points below that day's statistical forecast of 65 for the VIX, one of the most substantial deficits since the 2008 financial crisis. The VIX increased to 83 a few days later and remained above 70 for several days afterward; the March futures contract ultimately settled at 70.

The undervaluation suggests that participants in volatility markets underreacted to the growing risks of a COVID-19 pandemic in its early stages. It is as if, just when the risk of the pandemic was growing, VIX futures prices were too sluggish to rise relative to a statistically fair forecast of the VIX. Of course, both futures prices and statistical forecasts turned out lower than the VIX on March 18 after the fact, but the key observation is that VIX futures were undervalued in real-time. A watchful trader could have seen signals of the undervaluation beforehand and made substantial profits in March by taking a long position in VIX futures.

The rest of the article uses the broader history of data in VIX futures to show that the underreaction of volatility markets to the growing risks of a pandemic in late February and early March 2020 was a vivid example of systematic patterns. Using the whole history of data is important to evaluate whether any signals provide genuine information about valuation.

Three findings support this conclusion. First, the futures price minus the fair statistical forecast, or “VIX premium,” provides a genuine signal of valuation because it predicts subsequent movements in futures prices with the expected magnitude. Next, a trading strategy that times the valuation signal produces significant risk-adjusted returns net of transaction costs that improves upon the performance of a strategy that always shorts volatility. Finally, premiums have systematically underreacted to increases in risk in that such increases tend to push premiums towards undervaluation. The COVID-19 episode increased the magnitude of these relationships.

The article builds on and extends results from Cheng (2019), which first considers the VIX premium and how the premium underreacts to risk. The article is also related to Lochstoer and Muir (2019), which formally models of how investors’ expectations about future volatility underreact to the news. Finally, the article builds on the growing interest in volatility derivatives and the variance risk premium (see, e.g., Carr and Lee, 2009 for a review; early works included Coval and Shumway, 2001, and Bakshi and Kapadia, 2003). Compared to the literature, the methods proposed in this article use relatively parsimonious techniques, leading to insights that are straightforward to implement and build upon in practice. The article concludes with avenues for practical application and development of the proposed trading strategy.

I. THE EARLY STAGES OF THE COVID-19 PANDEMIC

Table 1 lists the values of the S&P 500, VIX, VIX futures settlement prices, and fair statistical forecasts for each futures expiration date in late February and early March 2020, the key period for this section of the article. Data come from Bloomberg Professional.

On February 19, the S&P 500 closed at a record high, just 14 points shy of 3400. The VIX closed around 14, a value which was, as with most days over the bull market of the preceding years, well below its historical average of 20. These events occurred even though the first U.S. case of the COVID-19 coronavirus had been reported in the mainstream news and confirmed by the CDC on January 21. The news out of China had been dire up until that point; on February 12, the media reported 14,000 new cases in Hubei Province alone. By March 2, the news was getting worse: coronavirus cases were spiking in Europe, and the U.S. had reported possible community spread as well as its first coronavirus-related death. The S&P 500 had fallen to just under 3100. The VIX had risen to 33, a substantial increase from 14.

The news worsened from there: after March 2, Italy placed travel restrictions on the northern part of the country, the United States restricted travel from Europe, and the WHO officially declared the outbreak a pandemic. By March 12, the VIX had exploded to 75. A few days later, the United States declared a national emergency, the VIX closed at 83 and then stayed above 70 for several days. Eighty-three was an all-time high for the VIX, higher than the VIX reached even in the 2008 financial crisis.

As noted in the introduction, futures prices were effectively too sluggish to rise during these early stages of the outbreak compared to fair statistical forecasts of where the VIX would be on futures expiration dates. For example, with the VIX at 33 on March 2, the futures price for a contract expiring March 18 settled at 26, even though a statistical forecast would have placed the VIX higher at 30 for that expiration date. When the VIX closed at 75 on March 12, the March futures settled at 58, while a statistical forecast put the VIX at 65.

To produce these statistical forecasts, I estimate a standard statistical model that assumes the VIX follows an ARMA process at a daily frequency. An ARMA process captures two key features. The autoregressive (AR) component captures mean reversion, or the tendency of the

VIX to move towards its long-term average. The moving average (MA) component captures the possibility that recent unexpected VIX movements may have a direct impact on short-term fluctuations. (Hamilton, 1994 provides a textbook treatment.) Cheng (2019) estimates an ARMA(2,2) model using data on the VIX from its starting date in 1990 through the start of 2004 and finds that it fits well compared to several other forecast models. I update the methodology in that paper by estimating the ARMA model on every trading date using an expanding window that uses, for any given trading date, all available data of the VIX starting in 1990 up until the previous trading date. I then use estimates from the model combined with information up to the trading date to obtain statistical forecasts of the VIX as of each futures expiration date.

The key valuation metric is the futures price minus the statistical forecast, or "VIX premium":

$$VIXP_t = F_t^T - \widehat{VIX}_t^T, \tag{1}$$

where F_t^T is the date-t futures price for a contract on date expiring on date T and $\widehat{VIX}_t^T$ is the estimated model's forecast for the VIX for date T as of date t. The VIX premium equals the estimated expected premium a long investor pays on a VIX futures contract over the remaining life of the contract, before applying the contract multiplier. Equivalently, it equals the expected profit for a short VIX futures position.

Table 1 shows that, toward the end of February, VIX premiums turned negative even as market volatility increased due to worsening news about COVID-19. Premiums were negative not only for the March and April contracts shown in the Table, but also throughout the futures curve.

A negative premium is an odd occurrence. Conceptually, it should be the other way around: premiums should be positive, with the futures price exceeding the fair forecast. The reason is that a trader who goes long VIX futures should expect to pay a premium to hedge future uncertainty. A negative premium thus represents undervaluation. (A positive premium is not necessarily a signal of overvaluation precisely because the positive premium is the compensation a short futures investor may require for insuring the long investor against increased uncertainty.) Table 1 reports that premiums are positive on average in the history of VIX futures, although there have been substantial deviations in the past.

A watchful trader could have observed this undervaluation in real-time at the end of February and profited from holding a long futures position over March. For example, if a trader opened a long position in the April contract and held it over March, she would have seen April futures prices double from 23 to over 46, resulting in a return of 100% for a fully collateralized long futures position for the month (ignoring interest on margin).

Figure 1 illustrates the potential for profit from such a trading strategy. The chart plots the 1-month VIX premium in 2020 in the solid line. The 1-month premium is the premium referencing the contract expiring the next month; on the last day of the month, the premium “rolls” the reference contract forward. For example, the 1-month premium in February 2020 references the March 2020 contract; on the last day of February, the premium rolls and references the April contract. The figure also plots the change in the futures price over the next day for the 1-month reference contract.

The figure suggests that negative VIX premium estimates tended to anticipate daily increases in futures prices, consistent with genuine undervaluation. Furthermore, the few positive premium estimates in March tended to anticipate futures price decreases. As we discuss below, this is not an isolated incident, and using the VIX premium as a trading signal systematically provides useful information about what futures position to take.

Overall, in the early stages of the COVID-19 episode, VIX futures prices were too sluggish to rise relative to a statistically fair forecast of the VIX just when the risk of the pandemic was growing, leading to undervaluation. It is worth comparing early 2020 to the fall of 2008, one of the few comparable market downturns of similar magnitudes in recent history. Figure 2 plots the estimated 1-month VIX premium for the three months leading up to the COVID-19 pandemic (beginning January 1, 2020) together with the three months leading up to the worst part of the financial crisis (beginning August 1, 2008). In both episodes, VIX premiums turned negative as each situation worsened before beginning to recover. It is as if, on the eve of large market downturns, long investors found it cheap to hedge future uncertainty.

The rest of the article places the COVID-19 episode in the context of historical patterns in VIX futures prices.

II. HISTORICAL EVIDENCE

This section shows that the underreaction of VIX futures prices during the early stage of the COVID-19 pandemic was consistent with and increased the magnitude of pre-existing historical patterns. It builds on and puts in practical context several findings of Cheng (2019).

A. VALIDATING THE VIX PREMIUM AS A VALUATION MEASURE

If VIX premiums provide genuine information about valuation, then estimated premiums should *systematically* forecast VIX futures price movements. I test this idea using the following return forecast regression:

$$xr_t = a + b\,VIXR_{t-1} + e_t. \tag{2}$$

The term xr_t is the monthly excess return from a fully collateralized long position in a VIX futures contract; equivalently, xr_t is the un-levered return less a risk-free interest rate earned on margin. The term $VIXR_t$ converts the VIX premium in Equation 1 into a monthly expected excess return.

I take the perspective of an investor who takes a rolling position in the 1-month contract by holding, over every month *t*, the futures contract expiring in month *t*+1. The investor establishes the position for month *t* at the end of the month *t*-1 at price F_{t-1}^{t+1} and closes the position at the end of month *t* at price F_t^{t+1}. For example, in February 2020, the investor would hold the March 2020 contract; in March, the investor would hold the April contract. This "roll" ensures the position is typically invested in the most liquid contract across the term structure and avoids issues associated with the mid-month final settlement of VIX futures (Griffin and Shams, 2018).

Given this perspective, the excess return and expected return in Equation 2 equal:

$$xr_t = \frac{F_t^{t+1}}{F_{t-1}^{t+1}} - 1, \tag{3}$$

$$VIXR_{t-1} = \left(\frac{\widehat{VIX}_{t-1}^{t+1}}{F_{t-1}^{t+1}}\right)^{\frac{21}{n}} - 1, \tag{4}$$

where *n* is the number of trading days between the end of month *t-1* and the mid-month futures expiration date in month t+1. The scaling factor of 21/*n* re-scales the expected return to a 1-month horizon by accounting for the number of such days.

Figure 3 plots $VIXR_t$ through time. On average, it is negative, although there are substantial deviations. From Equation 4, the expected return *VIXR* is negative when the premium *VIXP* is positive, and vice versa. A futures investor who expects to pay a positive premium *VIXP* for hedging future uncertainty should expect to earn a negative expected return *VIXR*. If futures are undervalued so that the premium *VIXP* is negative, then the expected return *VIXR* is positive.

If VIX premium estimates are valid estimates of expected returns, then we expect estimates of *b* around 1. In contrast, if the estimates do not have statistical power to forecast subsequent returns, then we expect estimates of *b* that are statistically indistinguishable from zero or that are even negative. For example, if positive values of $VIXR_{t-1}$ were erroneous and true expected returns are negative, negative values of xr_t would tend to follow positive values of $VIXR_{t-1}$.

Table 3 reports estimates of *b* close to 1 and statistically distinct from zero. Columns 1 and 2 differ by including and excluding the first quarter of 2020, respectively. Samples in both columns begin in April 2004, when VIX futures started trading. Although *VIXR* has predictive power for returns in both columns, the larger estimates in column 1 indicate that positive *VIXR* predicted particularly large positive returns in 2020. The R-squared is also much larger in Column 1 than in Column 2, which was already large for this type of return forecast regression. The early period of 2020 would have thus been very profitable for an investor tracking the VIX premium.

Estimates from a daily frequency in Columns 3 and 4 paint a similar picture. I use a scale factor of 1/*n* in Equation 4 since we are predicting daily returns. I use $VIXR_{t-2}$ as a predictor instead of $VIXR_{t-1}$ as a further robustness to ensure that all information would be available in real-time to an investor.

Overall, VIX premiums are valid estimates of premiums in the sense that they tend to predict future returns with a coefficient near 1. (Note, however, that like any statistical measure, it contains noise.) To interpret this result in terms of undervaluation, note that a positive expected return $VIXR_{t-1}$ indicates that futures prices are undervalued as of date *t*-1. The estimates of *b* in Table 3 suggest that these periods tend to pre-sage increases in futures prices in month *t* as one would expect if futures prices were genuinely undervalued.

These insights hold even though the ARMA statistical model is unlikely to be the optimal VIX forecasting model or the VIX's exact process. This article uses the ARMA model as a baseline because it is well-known, parsimonious, and easy to implement in practice, requiring only the history of the VIX and minor computational resources to estimate. Cheng (2019) shows that, for the relevant horizons considered here, an ARMA(2,2) model fits the data reasonably well at the monthly horizon when compared with models with other lag structures or direct forecast models that allow for a heterogeneous autoregressive (HAR) structure or other predictors. Undoubtedly, such forecasts can be improved in practice. As noted above, any existing measurement error should push estimates of *b* towards zero or even into negative territory.

B. RISK-ADJUSTED RETURNS

If VIX premiums provide genuine information about valuation, and valuations fluctuate for reasons unrelated to risk, it should also be possible to generate positive risk-adjusted returns from using *VIXR* as a trading signal. This section shows that a baseline "threshold" trading strategy that goes long or short based on the sign of the premium provides risk-adjusted returns net of transaction costs.

The threshold strategy is a dynamic strategy that holds long 1-month futures over date *t* if *VIXR* is greater than zero on date *t*-2 and short futures if *VIXR* is less than zero, with any necessary buy, sell, or roll transactions occurring on date *t*-1. For example, if $VIXR_{t-3} < 0$ but $VIXR_{t-2} > 0$, the strategy would hold a short position over date *t*-1 and then transact into a long position at the end of *t*-1 before holding that position over date *t*. The strategy rolls a short or long position in the reference contract forward at the end of each month. The strategy is monitored daily, and the trading signal dated *t*-2 is available in real-time on date *t*-1.

In the following analysis, I assume all transactions occur at the relevant bid or ask prices. The bid-ask spread is important because the strategy transacts at least once a month (during the roll) and because spreads can widen in times of major market movements.

This threshold strategy has a risk-adjusted return of 3.5% per month (t-statistic: 2.3) and CAPM beta of 0.2 over the history of VIX futures. Column 1 of Table 4 Panel A reports these estimates from a standard CAPM performance evaluation regression at the daily frequency. For the market return, I use the total return of the S&P 500. For the risk-free return, I use the return to a 1-month Treasury bill. A 3.5% monthly risk-adjusted return is large; by comparison, equity momentum strategies produce CAPM risk-adjusted returns on the order of 1% per month.

The estimates in Column 1 do not distinguish whether the threshold strategy produces risk-adjusted returns by successfully timing long/short positions or by shorting volatility. In general, shorting volatility generates positive CAPM risk-adjusted returns (Coval and Shumway, 2001; Bakshi and Kapadia, 2003). To better distinguish between the two sources of returns, Column 2 reports that an "always-short" strategy that only rolls short futures positions had a risk-adjusted return of 0.9% (t-statistic: 0.9) per month and a CAPM beta of 2.4 over the same time period. The risk-adjusted return is not statistically distinguishable from zero for reasons that will become clear shortly. The positive beta of shorting volatility is a result of the negative correlation between volatility movements and market returns (French, Schwert, and Stambaugh, 1987). Comparing columns 1 and 2 shows that the threshold strategy has a larger risk-adjusted return and lower market exposure than the always-short strategy.

To make clearer what is happening, columns 3 and 4 report estimates of the following variation of the CAPM regression:

$$xr_t = (a_0 + a_1 \mathbf{1}[VIXR_{t-2} > 0]) + (b_0 + b_1 \mathbf{1}[VIXR_{t-2} > 0])\left(r_{M,t} - r_{f,t}\right) + e_t. \qquad (5)$$

This regression decomposes the standard CAPM performance regression into two pieces: the risk-adjusted return and beta when the signal *VIXR* calls for a short position (a_0 and b_0), and the differential return and beta when *VIXR* calls for a long position (a_1 and b_1).

As the estimates in column 3 indicate, the threshold strategy earns positive risk-adjusted returns when it is short and statistically similar risk-adjusted returns when the strategy is long. The risk-adjusted return a_0 when the strategy is short equals 2.5% per month. The differential risk-adjusted long return a_1 equals an additional 0.1% per month, although the standard error is large and so that risk-adjusted long returns are statistically indistinguishable from short returns. The large standard error of a_1 is because returns are volatile following undervaluation, a topic we dive into more in the next section. The CAPM beta flips based on whether the strategy is long or short: the beta when the strategy is short is 2.6 and falls by 4.8 to -2.2 when the strategy is long.

Note that the risk-adjusted return in column 1 of 3.5% per month is higher than the estimates implied by column 3 because the standard CAPM performance regression in Column 1 fails to account for this significant time-variation in betas.

In contrast, an always-short strategy incurs significant risk-adjusted losses whenever the *VIXR* signal calls for a long position. Column 4 reports the estimates of Equation 5 for this strategy. It earns a risk-adjusted return a_0 of 2.9% per month when the signal *VIXR* calls for a short position. However, when *VIXR* calls for a long position, the estimate of a_1 indicates that the strategy earns risk-adjusted returns that are 7.6% lower for a risk-adjusted loss of 4.6% during these periods. This pattern helps explain why the always-short strategy did not earn a significant risk-adjusted return in the standard CAPM regression of Column 2. The reason is that the strategy was short when expected returns, as measured by *VIXR*, were positive.

Comparing the estimates in Panel A with the analogous estimates in Panel B that do not include 2020 shows that the COVID-19 event increases the estimates of the risk-adjusted returns for the threshold strategy. An always-short strategy suffered significant losses in 2020, even net of market exposure.

Overall, the estimates in Table 4 indicate that timing VIX premiums produces risk-adjusted returns and that the early stages of the COVID-19 pandemic were a particularly profitable period for such timing strategies. The conclusion of this article discusses how future research and development might tune and optimize such strategies.

C. SYSTEMATIC UNDERREACTION TO RISK

In the early stages of the COVID-19 pandemic, VIX futures prices became undervalued just when the risk of the pandemic was growing; as noted earlier, a similar pattern occurred during the 2008 financial crisis. Aside from just these two episodes, increases in risk systematically tend to move premiums toward undervaluation.

Table 5 examines how premiums *react* to forward-looking measures of risk and how premiums *forecast* future realized risk. If increases in risk systematically push premiums towards undervaluation, we should see that *VIXR* increases when forward-looking measures of risk rise and that higher *VIXR* forecasts higher subsequent realized risk.

Panel A starts with how premiums react to risk and reports estimates from the regression:

$$VIXR_t = a + b_0\, \sigma_t + \textstyle\sum_{s=1}^{3} b_s\, \sigma_{t-s} + \sum_{s=1}^{3} c_s\, VIXR_{t-s} + e_t. \qquad (6)$$

In column 1, the risk measure σ_t is the VIX itself, which measures forward-looking volatility in the S&P 500. In Column 2, the risk measure σ_t is the CBOE VVIX ("VIX of VIX") index, which measures the forward-looking volatility of VIX futures prices. The lag structure in Equation 6 allows for dynamics and accounts for any time-series predictability in the VIX premium.

The table reports a positive coefficient on b_0, indicating that increases in risk tend to push premiums towards undervaluation and positive expected returns. The coefficient b_1 is negative, and the magnitude indicates that premiums tend to correct about a month after the undervaluation. Columns 3 and 4 repeat this exercise but exclude 2020. Estimates are smaller, indicating that the COVID-19 episode only increased the magnitude of the estimated relationship.

We then turn this exercise around and further ask whether undervaluation predicts higher *subsequent* realized risk. For example, undervaluation at the end of February 2020 preceded a month of extraordinary volatility, and we can ask whether this pattern is systematic. Panel B reports estimates of the following volatility forecast regression:

$$\sigma_t = a + \textstyle\sum_{s=1}^{3} b_s\, VIXR_{t-s} + \sum_{s=1}^{3} c_s\, \sigma_{t-s} + e_t. \qquad (7)$$

The variable σ_t is the standard deviation of daily log returns in month *t* for the S&P 500 (Column 1) or fully collateralized 1-month VIX futures (Column 2). The lag structure in Equation 7 accounts for the predictability of volatility.

Columns 1 and 2 of Panel B show that the 1-month lag on $VIXR_{t-1}$ positively predicts volatility σ_t for both VIX futures and the S&P 500. This positive predictive relationship indicates that periods of positive expected returns – that is, periods of negative premiums and undervaluation – tend to pre-sage higher volatility in both VIX futures as well as the broader market. As before, Columns 3 and 4 indicate that the COVID-19 episode only increased the magnitude of the estimated relationship.

Overall, the estimates in Table 5 indicate that VIX futures underreact to risk and are undervalued just when risk going forward is high. It is as if a long futures investor pays a small expected premium – or, if anything, receives a premium – for hedging uncertainty just when markets (the stock market or VIX futures) are about to be very volatile.

D. RELATIONSHIP TO BACKWARDATION AND CONTANGO

The VIX premium is related to the slope of the futures curve as summarized by the futures basis, or futures price minus spot VIX. Fama (1984) and subsequent papers note that one can decompose the futures basis as:

$$F_t^T - VIX_t = (F_t^T - E_t[VIX_T]) + (E_t[VIX_T] - VIX_t). \tag{8}$$

The left-hand side is the futures basis. The first term on the right-hand side is the VIX premium, and the second term is the expected change in the VIX.

A common refrain is that short positions in VIX futures tend to be profitable when the futures curve is in contango (positive basis) and not profitable when the curve is backwardated (negative basis). Equation 8 makes clear why this can be the case: all else equal, a higher premium *VIXP* pushes the basis towards contango, and a negative premium *VIXP* pushes the basis towards backwardation. At either monthly or daily frequencies, the correlation between the 1-month VIX premium and the 1-month futures basis is roughly 0.8.

Table 6 reports estimates of the return forecast regression in Equation 2 using the futures basis as a predictor. Estimates from the full sample (columns 1 and 3) indicate that the basis has provided a signal for returns over the history of VIX futures. For example, backwardation preceded high futures returns in the early stages of the COVID-19 pandemic.

However, columns 2 and 4 indicate that the futures basis had weakened as a statistical predictor of futures returns by the end of 2019, as the estimates of *b* when excluding 2020 are not statistically distinguishable from zero. Equation 8 makes clear that the basis might be a noisy signal of premiums because it also contains information about expected VIX movements. Thus, contango can signal a large premium, but it can also signal a small premium with a large expected upward movement in the VIX. Conversely, backwardation can signal a negative premium, but it can also signal a positive premium with a large expected drop in the VIX.

By filtering out information about expected VIX movements from the basis, Table 3 showed that the premium was providing a signal of subsequent futures returns even at the end of 2019.

III. CONCLUSION

The VIX futures market is a key market for bets on volatility, surpassing even the SPX variance swap market in vega notional for short-term contracts (Mixon and Onur, 2014). In such an actively traded market, the underreaction documented in this paper surprising. Lochstoer and Muir (2020) write down a formal model of underreaction to volatility news and show that its implications are consistent with evidence from Cheng (2019) about variance risk premiums as well as the ambiguous relationship between volatility and equity returns (Moreira and Muir, 2017). Cheng (2019) additionally suggests that time-varying hedging demand contributes to underreaction.

Whatever the academic explanation, however, this article highlights that the underreaction of the VIX futures market to the increasing risks of a COVID-19 pandemic in late February and early March was a vivid example of a broader anomaly, with important implications for practice. The baseline trading strategy in this article can form the basis for strategies that either profit from going long futures during periods of undervaluation or that mitigate the risk of losses to short positions by exiting shorts when futures are undervalued and risk is high. Practitioners can tune the baseline trading strategy proposed in this article along several dimensions, including optimizing the VIX forecast model, trading frequency, leverage, adapting the trading strategy around the magnitude of the trading signal, or potentially going to cash in certain circumstances. These issues are promising areas for future practical application and development.

REFERENCES

Bakshi, Gurdip, and Nikunj Kapadia, 2003. Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies* 16, 527-566.

Carr, Peter, and Roger Lee, 2009, Volatility derivatives. *Annual Review of Financial Economics* 1, 319-339.

Cheng, Ing-Haw, 2019. The VIX Premium. *Review of Financial Studies* 32, 180-227.

Coval, Joshua D., and Tyler Shumway, 2001. Expected option returns. *Journal of Finance* 56, 983-1009.

Fama, Eugene F., 1984. Forward and spot exchange rates. *Journal of Monetary Economics* 14, 319-318.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987. Expected stock returns and volatility. *Journal of Financial Economics* 19, 3-29.

Griffin, John, and Amin Shams, 2018. Manipulation in the VIX? *Review of Financial Studies* 31, 1377-1417.

Hamilton, James, 1994. *Time Series Analysis*. Princeton University Press.

Lochstoer, Lars, and Tyler Muir, 2019. Volatility expectations and returns. Working paper, University of California at Los Angeles.

Mixon, Scott, and Esen Onur, 2014, Volatility derivatives in practice: activity and impact, Working paper, Commodity Futures Trading Commission.

Moreira, Alan, and Tyler Muir, 2017. Volatility managed portfolios. *Journal of Finance* 72, 1611-1644.

Newey, Whitney K., and Kenneth D. West, 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703-708.

Table 1 – SPX, VIX, VIX futures, and VIX forecasts in February and March 2020

This table reports the values of the S&P 500, VIX, VIX futures daily settlement prices, and VIX forecasts. Timeline notes come from the Think Global Health project (an initiative of the Council on Foreign Relations) and a search of articles on the New York Times and Wall Street Journal.

			March 18 Futures			April 15 Futures			
Date	**S&P 500**	**VIX**	**Price**	**Fcast.**	**VIXP**	**Price**	**Fcast.**	**VIXP**	**Notes**
Feb 12 (Wed)	3379.45	13.74	15.32	15.26	0.06	16.13	15.80	0.33	14,000 new cases in Hubei, China
Feb 13 (Thu)	3373.94	14.15	15.57	15.33	0.24	16.33	15.86	0.47	
Feb 14 (Fri)	3380.16	13.68	15.43	14.99	0.44	16.13	15.56	0.56	First deaths in Europe
Feb 18 (Tue)	3370.29	14.83	15.82	15.48	0.35	16.52	15.98	0.54	
Feb 19 (Wed)	3386.15	14.38	15.38	15.22	0.16	16.33	15.76	0.56	S&P 500 record high
Feb 20 (Thu)	3373.23	15.56	16.08	15.80	0.27	16.83	16.26	0.57	27 countries, 75,000 cases, 2,100 deaths
Feb 21 (Fri)	3337.75	17.08	16.92	16.65	0.27	17.33	16.99	0.33	Iran and South Korea surge; 34 U.S. cases
Feb 24 (Mon)	3225.89	25.03	20.08	21.11	-1.03	19.38	20.82	-1.44	Italy surges
Feb 25 (Tue)	3128.21	27.85	22.23	23.31	-1.09	20.88	22.72	-1.84	CDC issues U.S. warning; European spread
Feb 26 (Wed)	3116.39	27.56	22.33	23.91	-1.59	20.98	23.23	-2.26	First possible U.S. community spread Germany warns
Feb 27 (Thu)	2978.76	39.16	26.27	31.00	-4.72	23.52	29.30	-5.77	European cases surge
Feb 28 (Fri)	2954.22	40.11	26.33	32.95	-6.63	23.02	30.97	-7.94	56 countries, 84,000 cases, 2,900 deaths
Mar 2 (Mon)	3090.23	33.42	26.27	30.28	-4.01	23.33	28.69	-5.36	First U.S. deaths over the prior weekend Several states declare emergencies
Mar 3 (Tue)	3003.37	36.82	29.17	32.71	-3.53	25.52	30.75	-5.22	Federal Reserve 50bp emergency cut
Mar 4 (Wed)	3130.12	31.99	27.42	30.55	-3.12	24.63	28.92	-4.29	California declares emergency
Mar 5 (Thu)	3023.94	39.62	31.88	35.19	-3.32	27.52	32.83	-5.31	New Jersey, Maryland declare emergencies
Mar 6 (Fri)	2972.37	41.94	35.78	37.50	-1.72	30.33	34.77	-4.45	90 countries, 100,000 cases, 3,400 deaths
Mar 9 (Mon)	2746.56	54.46	44.38	46.21	-1.84	36.22	42.02	-5.80	National lockdown in Italy
Mar 10 (Tue)	2882.23	47.30	41.83	43.54	-1.72	34.78	39.92	-5.14	23 U.S. states with emergencies
Mar 11 (Wed)	2741.38	53.90	46.35	48.69	-2.34	38.58	44.03	-5.45	WHO declares pandemic
Mar 12 (Thu)	2480.64	75.47	58.30	65.19	-6.89	45.83	57.05	-11.2	Worst day for U.S. stock market since 1987
Mar 13 (Fri)	2711.02	57.83	53.42	55.75	-2.32	43.90	50.50	-6.60	U.S. declares a national emergency 121 countries, 143,000 cases worldwide

Table 2 – Summary statistics

This table reports the average and standard deviation of *VIXP* (Equation 1) through the end of March 2020 at the daily frequency. Values reference the *n*-month ahead contract over the month. For example, the 1-month premium in February 2020 references the March 2020 contract; on the last day of February, the premium rolls and references the April contract. Units are VIX points. The $n = 1$ the series begins in 2004 when VIX futures start trading and references the next available contract if a 1-month contract is not available. The series for $n \geq 2$ begins in 2007 when a complete term structure is available every month.

	Mean	**St. Dev.**	**T**
1-month	0.67	1.91	4,031
2-month	1.29	2.69	3,334
3-month	1.57	2.90	3,334
4-month	1.81	3.09	3,334
5-month	2.01	3.26	3,334

Table 3 – Predicting VIX futures returns

This table reports estimates of Equation 2 from the text. The regression forecasts the excess returns of fully collateralized rolling 1-month long VIX futures positions using VIXR. Columns 1 and 2 report estimates at the monthly frequency where returns are measured over month t and VIXR is measured at the end of month t-1. The full sample in column 1 includes months from April 2004 to March 2020, while Column 2 does not include 2020. Columns 3 and 4 report estimates at the daily frequency where returns are measured over date t and VIXR is measured at the end of date t-2. Returns are expressed in monthly percentage points in Columns 1 and 2 and daily percentage points in Columns 3 and 4. The table reports Newey and West (1987) standard errors with 3 lags at the monthly frequency and 22 lags at the daily frequency. Bold-faced coefficients are more than two standard errors away from zero.

Dep. Var.: Futures return, t	Monthly		Daily	
	Full samp.	Ex. 2020	Full samp.	Ex. 2020
	(1)	(2)	(3)	(4)
b : Slope on VIXR	**1.315**	**0.918**	**1.106**	**0.889**
	(0.430)	(0.296)	(0.291)	(0.212)
a : Constant	0.375	-1.435	0.034	-0.034
	(2.169)	(1.680)	(0.092)	(0.072)
T	192	189	4029	3967
R^2	0.139	0.076	0.007	0.005

Table 4 – Risk-adjusted returns

Panel A reports estimates of CAPM performance regressions and estimates of Equation 5 for the full sample from April 2004 through March 2020 for the threshold strategy and an always-short futures strategy. Columns 1 and 2 report estimates of standard CAPM regressions of excess returns as dependent variables on the excess market return. Columns 3 and 4 report estimates of Equation 5 of the text, which decomposes the CAPM return into periods when VIXR calls for a short or long position. Panel B reports estimates for the period excluding 2020. Units are in monthly percentage points. The table reports Newey and West (1987) standard errors with 22 lags. Bold-faced coefficients are more than two standard errors away from zero.

Panel A: Full sample

Futures return, t	CAPM		Decomposition	
	Threshold	Short	Threshold	Short
	(1)	(2)	(3)	(4)
a_0 : Alpha	**3.49**	0.88	**2.51**	**2.97**
	(1.53)	(0.96)	(0.96)	(0.95)
a_1 : Differential alpha			0.11	**-7.61**
			(2.57)	(2.76)
b_0 : Beta	0.18	**2.37**	**2.63**	**2.63**
	(0.31)	(0.22)	(0.19)	(0.19)
b_1 : Differential beta			**-4.75**	-0.50
			(0.39)	(0.33)
T	4029	4029	4029	4029
R^2	0.49	0.49	0.50	0.50

Panel B: Excluding 2020

Futures return, t	CAPM		Decomposition	
	Threshold	Short	Threshold	Short
	(1)	(2)	(3)	(4)
a_0 : Alpha	**2.78**	1.37	**2.50**	**2.98**
	(1.35)	(0.89)	(0.92)	(0.91)
a_1 : Differential alpha			-1.53	**-5.92**
			(2.25)	(2.26)
b_0 : Beta	0.43	**2.47**	**2.73**	**2.73**
	(0.30)	(0.25)	(0.16)	(0.16)
b_1 : Differential beta			**-4.91**	-0.55
			(0.45)	(0.41)
T	3967	3967	3967	3967
R^2	0.50	0.50	0.51	0.51

Table 5 – Underreaction to risk

Panel A reports estimates of Equation 6 from the text. The regression examines the monthly relationship between VIX premiums and forward-looking risk measures, controlling for lags of the premium. In Columns 1 and 3, the risk measure is the VIX, while in Columns 2 and 4, the risk measure is the VVIX. Panel B reports estimates of Equation 7 from the text. The regression forecasts monthly return volatility of either the S&P 500 (Columns 1 and 3) or VIX futures (Columns 2 and 4) using lags of VIXR and controlling for lags of volatility. Return volatility is calculated as the standard deviation of daily log returns each month in annualized percentage points; VIX and VVIX are also expressed in annualized percentage points. VIXR is expressed in monthly percentage points as in Equation 4. For brevity, the table reports only the *b* coefficients for each regression. The table reports Newey and West (1987) standard errors with 3 lags at the monthly frequency. Bold-faced coefficients are more than two standard errors away from zero.

Panel A. Reaction to risk

Dep. Var.: VIXR, t	Full sample		Ex. 2020	
	VIX	VVIX	VIX	VVIX
	(1)	(2)	(3)	(4)
b_0: Risk, t	**0.55**	**0.17**	**0.52**	**0.15**
	(0.07)	(0.03)	(0.07)	(0.03)
b_1: Risk, t-1	**-0.49**	**-0.10**	**-0.45**	**-0.10**
	(0.09)	(0.03)	(0.08)	(0.03)
b_2: Risk, t-2	-0.06	-0.05	-0.07	-0.05
	(0.06)	(0.03)	(0.06)	(0.03)
b_3: Risk, t-3	**-0.13**	-0.03	-0.11	-0.02
	(0.06)	(0.02)	(0.06)	(0.02)
T	190	166	187	163
R^2	0.71	0.53	0.71	0.55

Panel B. Predicting subsequent risk

Dep. Var.: Volatility, t	Full sample		Ex. 2020	
	S&P 500	VIX Fut.	S&P 500	VIX Fut.
	(1)	(2)	(3)	(4)
b_1: VIXR, t-1	**1.02**	**2.13**	**0.47**	**1.37**
	(0.49)	(0.77)	(0.23)	(0.51)
b_2: VIXR, t-2	**-0.52**	**-1.60**	**-0.36**	**-1.28**
	(0.18)	(0.47)	(0.15)	(0.43)
b_3: VIXR, t-3	-0.33	-0.41	-0.11	-0.14
	(0.24)	(0.53)	(0.12)	(0.42)
T	190	190	187	187
R^2	0.56	0.16	0.59	0.11

Table 6 – Relationship to backwardation and contango

This table reports estimates of Equation 2 from the text using the futures basis as a predictor. The exercise otherwise is identical to that of Table 3. The table reports Newey and West (1987) standard errors with 3 lags at the monthly frequency and 22 lags at the daily frequency. Bold-faced coefficients are more than two standard errors away from zero.

Dep. Var.: Futures return, t	Monthly		Daily	
	Full	Ex. 2020	Full	Ex. 2020
	(1)	(2)	(3)	(4)
b : Slope on basis	**0.976**	0.584	**0.552**	0.307
	(0.376)	(0.343)	(0.263)	(0.179)
a : Constant	1.867	-0.830	0.042	-0.072
	(2.663)	(2.470)	(0.126)	(0.094)
T	192	189	4029	3967
R^2	0.152	0.053	0.004	0.001

Figure 1 – VIX premiums in 2020

This figure plots the 1-month VIX premium in Equation 1, starting in February 2020. The premium in month t references the contract expiring in month t+1. The dots plot the futures price change over the next day.

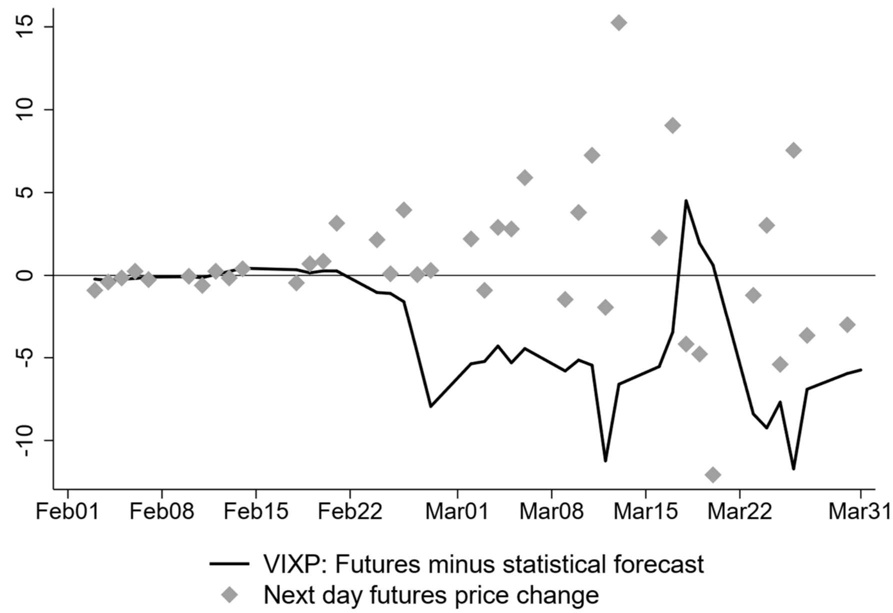

Figure 2 – Comparing the 2020 COVID-19 pandemic with the 2008 financial crisis

This figure plots the 1-month VIX premium in Equation 1 for the 3 months starting January 1, 2020, and the 3 months starting August 1, 2008.

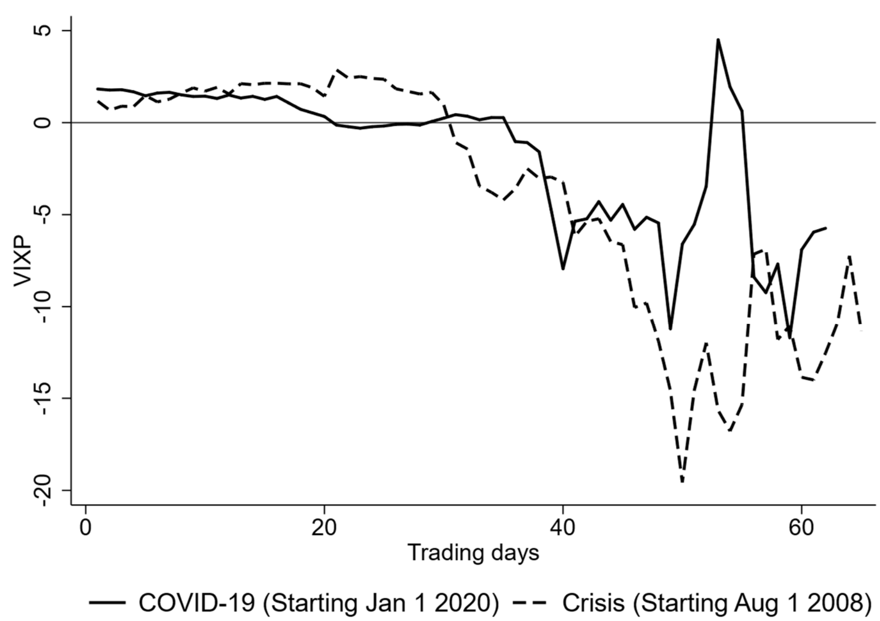

Figure 3 – VIXR through time

This figure plots the 1-month VIXR in Equation 4 through time. VIXR for month t is the expected return of the contract expiring in month t+1, calculated as of the end of month t-1, expressed in monthly percentage points.