Tail risk hedging with VIX Calls

Benjamin Siranosian
bsiranos@stanford.edu

June 2021

1 Introduction & Motivation

Investors would like to protect (hedge) against market downturns, and have many options for purchasing portfolio "insurance" to accomplish this. However, insurance can be expensive, especially in times of crisis. Ideally, proper protection would be purchased inexpensively before the event the investor is wishing to hedge against.

Here, I evaluate a strategy of systematic hedging against market downturns. The instrument was chosen to be out of the money call options on the volatility index (VIX). Often referred to as the "investor's fear gauge," the VIX is a calculated metric that reflects how much investors are willing to pay for portfolio insurance. VIX call options incrase in price with increased market volatility, which often accompanies severe market downturns. Out of the money options often expire worthless, and present a significant drag on portfolio performance in times of low volatility. However, when volatility spikes, these options can become worth 100x their initial value or more. As I will show here, the costs of systematic VIX call buying can be low enough, and the payoff large enough, that both overall and risk-adjusted returns of a portfolio holding a stock index can be improved. The benefit extends to a leveraged portfolio holding a mix of stocks and bonds. While I rely heavily on backtesting to show these results, I try to refrain from overfitting to the few market crashes in recent history, instead trying to select reliable principles to base the strategy on.

2 Inspiration

2.1 Why hedge?

Why should an investor hedge against market downturns? One reason may be to guarantee portfolio losses do not exceed a fixed amount over a given period of time. Another reason may be to trade some absolute returns for a lower volatility portfolio. Finally, hedging can allow an investor to tune the main part of their portfolio more aggressively, allocating a larger fraction to risky investments or applying leverage. As can be seen in the figure below comparing a hedged and un-hedged portfolio, the mean return of the hedged portfolio may be lower, but the expected value of the returns is higher. This is due in large part to the elimination of the left tail of the returns of the un-hedged portfolio.

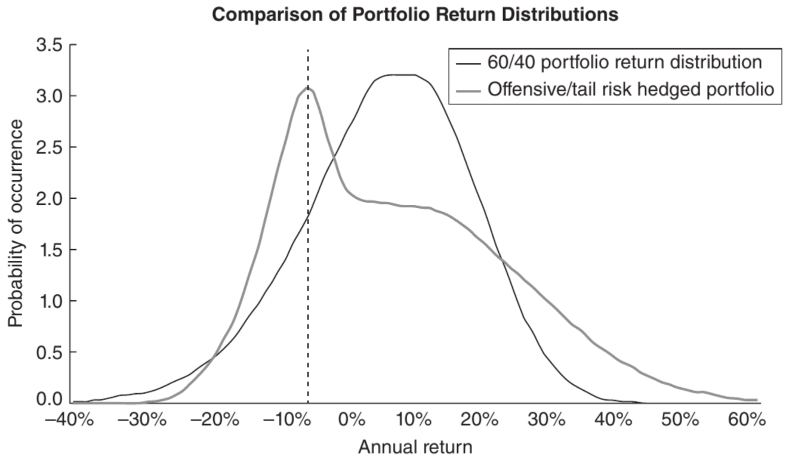

Figure 1: Return distribution of hedged (gray) and un-hedged (black) portfolios. The mean return of the hedged portfolio is lower, indicated by the dotted line, but the expected value of the returns is higher. Figure reproduced from [1]

2.2 VXTH: The VIX Tail Hedge index

I'm obviously not the first investor to think of a strategy like this. The Chicago Board Options Exchange maintains an index which tracks the returns of a strategy that systematically buys VIX calls, which was the inspiration for this research. This strategy places a variable allocation of it's portfolio into VIX calls, according to the allocation table below. Options with 30 days to expiration and a delta (an option "greek" describing the first-order sensitivity to price of the underlying) are purchased each month, and allowed to expire worthless or in the money. While this strategy did exceptionally well in 2020 (Figure 2), I believe it has several shortcomings that can be improved upon.

VIX Price	Hedge allocation
X $\leq$ 15	0%
15 > X $\leq$ 30	1%
30 > X $\leq$ 50	0.5%
X > 50	0%

2.3 Strategy decisions

There are a number of tradeable products on the market which benefit from volatility increases: VXX, UVXY, etc. Most of these products buy and roll VIX futures contracts to maintain their volatility exposure. They have also been excellent at erasing investor's capital since inception - simply look at the all time chart for one of these products to understand the point here. While the theta decay of options may be similar to the cost of holding these products, adding these products to a portfolio in a backtest does not improve returns.

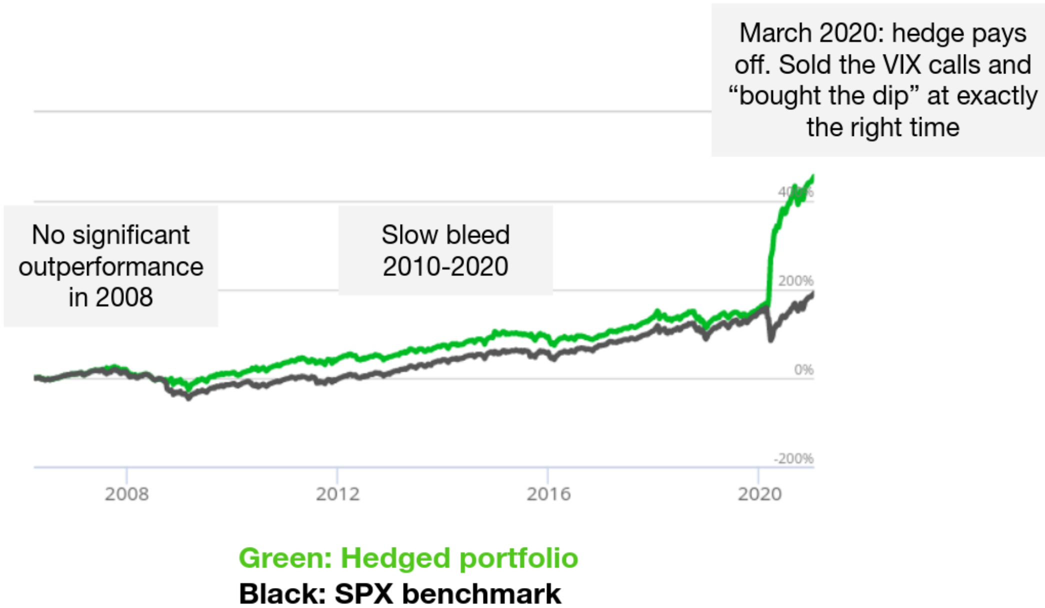

Figure 2: Equity curves of VXTH and the SPX index from 2006 to 2020. Key points are annotated.

2.3.1 Why use VIX calls?

- Convexity in a crisis: when volatility spikes, these options can easily return 10-100x the inital investment.
- Liquidity: VIX options are very liquid, with tens of thousands of contracts available at the bid and ask. In a crisis when liquidity in other markets can dry up, demand for these options stays high.
- Scalability: The options used in this strategy typically trade for $50-$100. With an allocation of 0.5%, that means a strategy buying a single option can be used on a portfolio as small as $10,000.

2.4 Areas of the strategy to improve

I focused on the following areas to improve the strategy implemented by VXTH.

- Trade signals: The relatively simple allocation of VXTH based on the VIX price could be improved, or at least tested to ensure it captures market conditions.
- Option delta: Buying options with a lower delta means they are more likely to expire worthless, but will have more convexity when volatility spikes.
- Hedge allocation: Should a larger or smaller fraction of the portfolio be dedicated to the hedge?

- Portfolio holding options with multiple expirations: Instead of a hedge containing a single option, it might be better to hold options with multiple dates to expiration.
- Monetization rules: The strategy above could, in theory, hold a contract through a volatility spike that resulted in greatly increased value, only to have the contract later expire worthless. I would like to develop a fixed return that the options should be sold at, with the profits re-invested in the rest of the portfolio.

3 Data Modeling

3.1 Stock and option data

Stock prices were downloaded using the ”Tidyquant” package in the R programming language. Option price data was obtained from Stanford's subscription to IvyDB. This dataset ended in 2019, so additional option data was purchased from historicaloptiondata.com. The two datasets were merged after verifying the prices in an overlapping month were consistent and wrangling the formats. Options on the VIX did not exist until 2006, so the period of 2006-2020 was considered for all backtests. There were few expirations and strikes available when VIX options began trading, so the backtests are limited in the earlier segments by this.

3.2 Backtesting code

I wrote a custom backtesting program in R. One of the major challenges was enabling fast lookup of option prices. I settled on a data structure with each day being an element in a list, and the option prices were a 3-D matrix with put/call, strike and expiration as dimensions. A separate list was maintained for the bid, ask and delta of each option.

The main program is basically a large loop for each trading day. Multiple portfolios can be defined, with different percentage allocations to different assets defined by a matrix. Each trading day, the values of the different portfolios are updated, and a decision to trade is made. If an option trade is made, it is added to the respective portfolios. I implemented logging of all trades and portfolio values for each day, which enabled better evaluation of the different strategies and debugging of the program.

Backtesting was done using the mid price for entries and exits, and assuming no transaction costs. This is one drawback of the model and may lead to real-world returns being less than backtests indicate. However, I did simulate a ”worst case scenario” where contracts were purchased at the ask and sold at the bid. While this portfolio did not preform as well, it was still better than a portfolio of pure stocks in both absolute and risk-adjusted returns. VIX option contracts can also be filled at the midpoint of the bid/ask (sometimes, not always, from experience).

Transaction costs are minimal for this strategy because it does not trade frequently, but can still represent a drag on real-world performance. Options contracts at my broker cost $0.65 to buy and sell, which may represent 1%-2% of the price of a contract. Stock trades are commission free and fees are marginal, theses and should not represent any appreciable costs.

3.3 Evaluation of results

I evaluated the portfolio performance on the following metrics: Compound annual growth rate (CAGR), Sharpe ratio (risk-adjusted returns) and maximum drawdown. To prevent overfitting, I refrained from optimizing the model to maximize a single value, instead preferring to understand how the parameters I was changing affected returns.

4 Backtest results

4.1 VXTH replication

First, I attempted to replicate the returns of VXTH. This was both an evaluation of my backtesting program and the strategy. Overall the returns of VXTH and my replicated strategy were similar, but there was divergence in a few key areas, which I believe are due to timing luck in VXTH. The index owned options that expired on the day when VIX had the highest value in the 2020 COVID crash, for example.

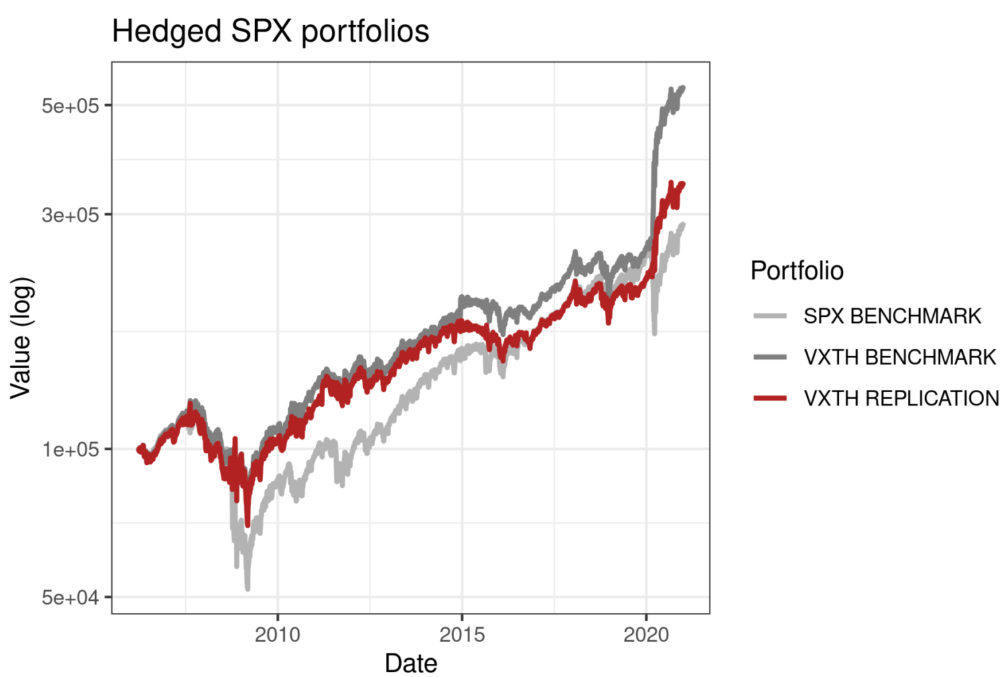

Figure 3: Equity curves comparing the SPX index (light gray), VXTH (dark gray) and the backtest-produced VXTH replication (red). While the two hedged portfolios have a similar "shape" of returns, differences in 2015 and 2020 set them apart. These differences are due to timing luck of the specific options purchased.

	SPX	VXTH (benchmark)	VXTH (replicated)
CAGR	7.49	12.2	8.27
Sharpe (Annualized)	0.49	0.67	0.62 %
Max Drawdown	52.5%	37.4%	35.1%

4.2 Trade signals

I next moved to test each of the points above, and evaluate their effect on the returns of my strategy. First, I evaluated the history of VIX prices, with the varying thresholds suggested by VXTH. We can see that VIX only spikes to the highest level on two occasions - the 2008 financial crisis and the 2020 COVID crash. By evaluating the transition matrix of VIX closing prices, we can see that there are never days when VIX jumps between more than one level. In particular, there is never a day when VIX goes from range 0 to range 2 or higher, without first spending a day in range 1. This means the strategy would not suffer by not hedging when volatility was at its lowest level.

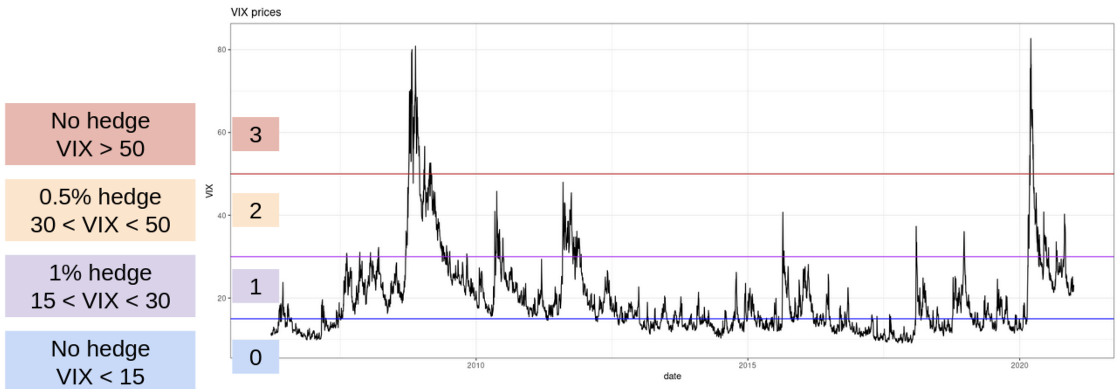

Figure 4: History of VIX prices from 2006-2020. Thresholds are determined by the allocation table above.

	0	**1**	**2**	**3**
0	1339	106	0	0
1	105	1756	35	0
2	0	35	262	9
3	0	0	9	65

Figure 5: Transition matrix of VIX closing prices, using the levels defined in the figure above. Each entry represents the number of days VIX changes from one level to the next, or stays the same (along the diagonal).

4.3 Option Delta

Buying lower delta options means you have more convexity - less frequent, but higher payoffs. I tested a range of option deltas from 50 to 5, and settled on 10 as a sweet spot. For a 90 day, 10 delta VIX call, 97.4% of crontracts will expire worthless. However, 2.9% of these contracts will experience a multiple of 50x or greater during their lifetime. During 2008 and 2020, these contracts experienced a multiple of 163x and 127x, respectively. Changing from 30 to 10 delta had a positive effect on absolute and negative effect on risk-adjusted returns in my backtest.

	SPX	SPX + Hedge (30 delta)	SPX + Hedge (10 delta)
CAGR	7.49	8.27	8.90
Sharpe (Annualized)	0.49	0.62	0.58

4.4 Hedge allocation

By lowering the hedge allocation, we can prevent money from being wasted on options that expire worthless. However, we still want the hedge to pay off when the time is right. I found that reversing the allocation of the hedge from what VXTH proposed was beneficial - less money was wasted on low-volatility times, and the payoff was greater in the big spikes.

	SPX	SPX + Hedge	SPX + Hedge (new allocations)
CAGR	7.49	8.27	9.81
Sharpe (Annualized)	0.49	0.62	0.58

4.5 Holding multiple options

A portfolio holding options with 30, 60 and 90 days to expiration was tested in an attempt to reduce the effect of timing luck. The total allocation to the hedge remained the same according to the thresholds. When the most recent option expired, a new 90-day option was purchased. In this way, a "ladder" of calls covered all events in the future 90 days. This change was beneficial to both absolute and risk-adjusted returns.

	SPX	SPX + Hedge (one call)	SPX + Hedge (ladder)
CAGR	7.49	8.27	9.89
Sharpe (Annualized)	0.49	0.62	0.64

4.6 Monetization rules

Rather than letting options expire, I tested if they should be sold at a fixed threshold. Again looking at the returns of all 90-day 10 delta VIX calls, I found that selling at 75x multiple would maximize returns. In backtesting, a 100x multiple provided the highest returns. It's likely that something in the 75-100x range would work well. This method reduced absolute but increased risk adjusted returns.

	SPX	SPX + Hedge	SPX + Hedge (100x monetization)
CAGR	7.49	8.27	8.14
Sharpe (Annualized)	0.49	0.62	0.75

4.7 Putting it all together

Finally, I tested portfolios that incorporated all of the modifications above. Together, all the changes had a positve impact on both absolute and risk-adjusted returns. I'm not sure how the effect of all of these changes became greater than the sum of the parts, and I'm still wary of timing luck in the model.

	SPX	SPX + Hedge	SPX + Hedge (all modifications)
CAGR	7.49	8.27	11.41
Sharpe (Annualized)	0.49	0.62	0.65

5 Extension to a leveraged portfolio

Having a systematic hedge may allow for a portfolio to use more leverage than would be appropriate otherwise. To simulate easy leverage a retrail trader would have access to, I constructed portfolios holding 100% UPRO (a product designed to have 3x daily SPX returns) or 55%/45% UPRO/TMF (3x daily returns of U.S. Treasury 20+ Year Bond Index). While this portfolio did exceptionally well in it's own right during the backtest period, adding the hedge still had a notable increase in performance. I haven't optimized specifically for this portfolio yet, it might be that the additional volatility of the leveraged portfolio necessitates an increased allocation to the hedge (Figure 6).

	SPX	SPX + Hedge	SPX + Hedge modified	leveraged + hedge
CAGR	7.49	8.27	11.41	23.7
Sharpe (Annualized)	0.49	0.62	0.65	0.88

6 Trading the strategy in practice

As this is a relatively simple strategy aimed at applicability to retail traders, I wanted to apply it to my own portfolio. I found that TD Ameritrade has a web API that clients can access, and there is a python package that allows for easy interface with this API (https://github.com/alexgolec/tda-api). Using these tools, I built a python program that could monitor market conditions, find the right options to buy, and scale the number of contracts depending on the value of your portfolio. I stopped short of implementing live trading with this program. Since the strategy trades so infrequently, the excess work to include execution and risk management was not worth it at the moment. The necessary trades are printed out and the user can execute them. I have put my money where my mouth is and I'm running this strategy with a fraction of my portfolio.

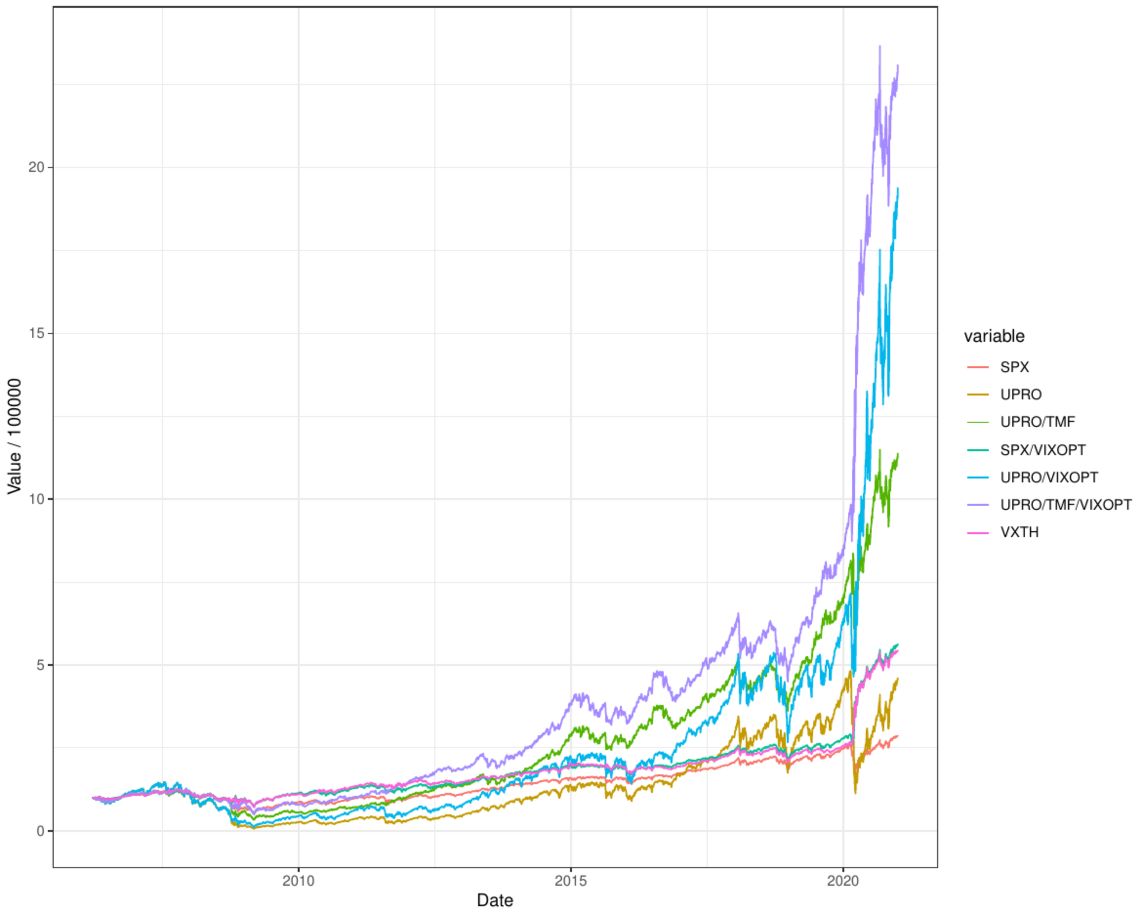

Figure 6: Equity curves many different portfolios. In purple, the leveraged stock/bond portfolio with the hedge does the best by far.

7 Discussion

The main questions I wanted to answer with this project were: how can systematic purchasing of insurance, which should have a negative expected value, result in excess returns? Shouldn't this free lunch not exist for the taking? Does this represent a fundamental disconnect in the way markets price this risk? It may be that VIX calls are systematically underpriced. If so, who is selling them to me, and why haven't they figured it out yet? Maybe after the 2020 crash, these calls will be higher priced, and the strategy will simply be too much of a drag on future portfolio returns.

While I can't say I have an answer to all these questions, I at least understand the mechanics of such a strategy better. Because it represents such a "long game" where it will not pay off except in events that happen once every 10-20 years, these risks may simply not factor into the trading model of the hedge funds running these markets. By waiting, patiently, possibly for decades, and by being prepared when the time comes, perhaps the lunch can be mine for the taking.

While I have evaluated the merits and different aspects of this strategy, my approach has many drawbacks. I relied heavily on backtesting, which is obviously backward-looking and says nothing of how this strategy will do in the future. While I attempted to avoid overfitting to the two significant market crashes in my dataset, a more robust backtest would include many more market crash periods. I could test the strategy on markets from other points in history, or in other countries. While this would involve calculating a measure of volatility and prices of derivatives, it would not be impossible. I also only explored hedging with a single instrument, while the universe of potential ideas to hedge market downturns is quite extensive. One of these may be a better hedge for my use case. Transaction costs and losses due to the bid/ask spread were not includded in my backtests, although I did address this in the Backtesting code section.

References

[1] Vineer Bhansali. *TAIL RISK HEDGING: Creating Robust Portfolios for Volatile Markets.* McGraw-Hill Education, 1st edition edition.