

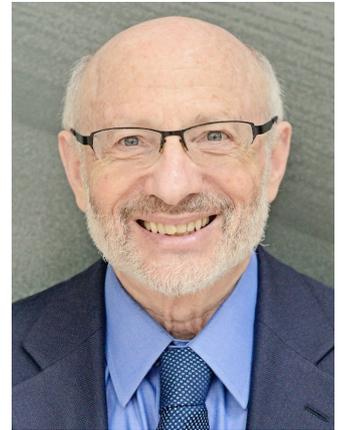
Empirical Support for Tail Risk Strategies

July 13, 2020

by Michael Edesess

The conventional wisdom is that put options are too expensive to use as “tail risk” insurance against extreme losses in equities. But reports earlier this year of their spectacular success in a fund advised by Nassim Nicholas Taleb prompted me to evaluate whether investors should use them.

After the S&P 500 index dived by 34% from February 19 to March 23, stories emerged of spectacular returns for so-called tail risk or “Black Swan” funds. The *Financial Times* reported that Universa Investments, advised by Black Swan author Nassim Taleb, “made a return for investors of more than 4,000 percent this year.” There ensued a Twitter battle between Taleb and Cliff Asness of hedge fund AQR, precipitated by AQR’s issue of – as Taleb tweeted – “2 flawed reports saying tail risk hedging doesn’t work.” Asness, of course, tweeted back in insulting Twitterdom fashion.



The evidence, however, shows that Taleb was right and Asness was wrong.

How to measure the performance of tail risk strategies

We don’t know the fine details of Universa’s strategies, but we can reasonably infer that they involve purchase of deep-out-of-the-money put options on an equity index. Their purpose is to buy a form of insurance that pays off if the market drops. Thus, they hedge against market declines.

As such, the performance of these funds should not be measured on the tail risk funds in isolation. They are clearly meant to enhance the performance of a whole portfolio of which they are only a part.

Thus, measuring return on investment the way the *Financial Times* reported doesn’t make sense, as former risk manager Aaron Brown pointed out in a May 29 *Bloomberg* article. “Suppose you pay \$100 per month for homeowner’s insurance on a house valued at \$250,000,” he says. If, one fine month, your house burns down you will receive \$250,000 – a one-month return, it would seem, of 249,900% on your monthly premium of \$100.

This isn’t very meaningful. First of all, in the same month, you lost \$250,000 on your house that burned down. The two should be considered in combination, not in isolation – the negative value of your lost house and the positive value of your insurance payoff.

Also, it’s most certainly not the return you get in the long run. To calculate that, you have to take into account all the monthly premiums you paid over the years, not just the premium in the month you got a payoff.

Part of the value of insurance is that it enables a positive form of moral hazard. If you have insurance, you can take more risk. If you couldn’t get homeowner’s insurance, you might have to be satisfied with a less valuable house, so that if it burns down you won’t lose as much.

By the same token, if you would otherwise reduce the risk of your portfolio by allocating to a 60/40 stock/bond mix, adding a tail risk strategy to mitigate the downside risk may make you feel comfortable investing 100 percent in equities instead. This is one of the chief benefits touted by the managers of tail risk funds.

A tail risk strategy, unlike a 60/40 mix, produces returns that are not symmetric. They curtail the downside but do not curtail the upside in a symmetric fashion. Therefore, conventional ways of measuring risk-adjusted return – such as a Sharpe ratio – do not apply well to the measurement of tail risk performance.

A historical study of the performance of a tail risk fund

Using historical option price data obtained from the Chicago Board Options Exchange (CBOE) from January 3, 2006 to

May 22, 2020, I evaluated the performance of a tail risk strategy similar to that provided by Universa.

Universa recommends that its clients invest 3.33% of their portfolios' funds in its tail risk fund, and the rest in risky assets. In its own calculations of combined fund performance Universa assumes that the client invests the rest of their portfolio in an SPX position – that is, the ETF with the ticker SPY.

Given these assumptions, Universa calculates¹ that the combination of 3.33% Universa Tail Hedge fund and 96.67% SPX returned +0.4% in the month of March 2020 (a month in which 100% SPX returned –12.4%), and 11.5% annualized since Universa's inception in March 2008 (compared to 7.9% for the SPX).

Are these numbers plausible? I ran my backtest to see.

Some decisions have to be made to specify the details of the backtest.

First, if the portfolio is allocated 3.33% to the tail risk fund and 96.67% to SPX, how often is that allocation rebalanced?

If it were rebalanced, say, quarterly, then since most deep-out-of-the-money options expire worthless or drop to low prices quickly if the market rises, it could result in losing as much as four times 3.33% every year, as the portfolio is rebalanced quarterly by replenishing the 3.33% lost in the previous quarter from the remaining SPX position.

This is unlikely to be what a tail risk fund or its clients would want to do. So I will assume the allocation is rebalanced annually, on the first day of each year. This limits the annual loss to the tail risk fund to the 3.33% invested in it.

I then assume that on that same first day of each year, the tail risk fund uses its 3.33% to purchase a put option that is approximately 60% out of the money and expires in approximately two years (in December of the following year). These are thinly-traded options, so the strategy may be liquidity-constrained, but the option purchased in 2020 would have achieved returns similar to those reported for Universa – so it is reasonable to assume it is using a similar strategy.

I ran the backtest on this strategy. Of chief interest is how the results compare with the results of the 60/40 stock-bond mix that the investor would otherwise invest in for risk mitigation. As a proxy for the bond portfolio I use the iShares Core U.S. Aggregate Bond ETF (AGG).

Month-to-month results are shown in Table 1 (11/2007-1/2020 means from November 2007 to January 2020, inclusive).

Table 1. Tail risk-hedged portfolio compared with 60/40 SPX/AGG and 100% SPX

Time period	Tail risk-hedged portfolio	60/40 SPX-AGG mix	100% SPX
Peak-to-peak: 11/2007-1/2020	8.5%	7.2%	8.5%
Full period: 1/2006-5/2020 ²	9.4%	7.0%	8.2%
Universa inception to April 7 Interim Letter: 4/2008-3/2020	11.0%	6.7%	7.9%
Without 2020 drawdown: 1/2006-12/2019	8.9%	7.4%	9.1%
Drawdown 1: 11/2007-2/2009	-31.8%	-33.3%	-52.2%
Drawdown 2: 2/2020-3/2020	13.5%	-12.3%	-20.9%
Drawdown 3: 5/2010-6/2010	-8.7%	-7.5%	-14.2%

These results support Universa's claims. Over the period from Universa's inception, when it claims to have achieved 11.5% to the SPX's 7.9%, my backtest returns 11.0%. (Of course, I don't expect to match Universa's performance figures because the strategy I am using is only an assumed approximation to Universa's.)

In every period except the small drawdown of mid-2010 the tail risk strategy outperforms the 60/40 mix, and it also outperforms 100% SPX in most periods too. In the deep bear market of November 2007 through February 2009 the tail risk

strategy lost only a little less than the 60/40 mix; thus, it protected the portfolio on the downside at least as well.

These results, performed over a limited period of time, lend credence to the proposition that a 100% equity portfolio combined with a tail risk strategy mitigates risk better than a constant 60/40 stock-bond mix, and performs better.

The April AQR research study – a disconcerting portent

Taleb's tweet was aimed in part at an [April 23, 2020 AQR research paper](#) that examined tail risk strategies compared to other risk-mitigating strategies, and concluded that the tail risk strategy is inferior to the ones that AQR practices.

Let us look at that research study. It begins with a contrived example that is so odd that it makes it nearly impossible to lend any credibility to the rest of the paper.

The paper has a message it wants to deliver to us, which is stated in the executive summary:

“Investors have a natural urge to protect their portfolios from sudden crashes like the one we’ve seen recently. We argue that they should instead focus on bad outcomes that unfold over longer periods, as those tend to be more detrimental to the long-term goal of wealth accumulation.”

They then construct an example of “Why the length of a drawdown matters” to illustrate their point:

“Take a highly simplified example: Suppose you manage a portfolio with an investment horizon of 10 years. This imaginary portfolio has a perfectly steady 6% annual return. But, there’s a catch: at some point in those 10 years, the portfolio stops making money and instead suffers a 20% loss. The silver lining, though, is that you get to choose how long that loss lasts: either A) over one month or B) over one year. For simplicity, assume that before and after the loss period, the portfolio goes right back to making 6% per year. Which do you choose?”

What they’re saying is that in nine out of 10 years you get a 6% return. But in the 10th year you either get A) a 20% loss in the first month, followed by going right back to making 6% a year for the rest of the year or B) a 20% loss for the year.

If you get a 20% loss in the first month followed by going right back to making an annualized 6% in the next 11 months, as in scenario A, your return will be –15.6% for the year. On the other hand, if you get a 20% loss over the whole year as in scenario B, then your return for the year will be –20%.

So what they’re saying is, would you rather get 6% a year for nine years and –15.6% for the 10th year, or would you rather get 6% a year for nine years and –20% for the 10th year?

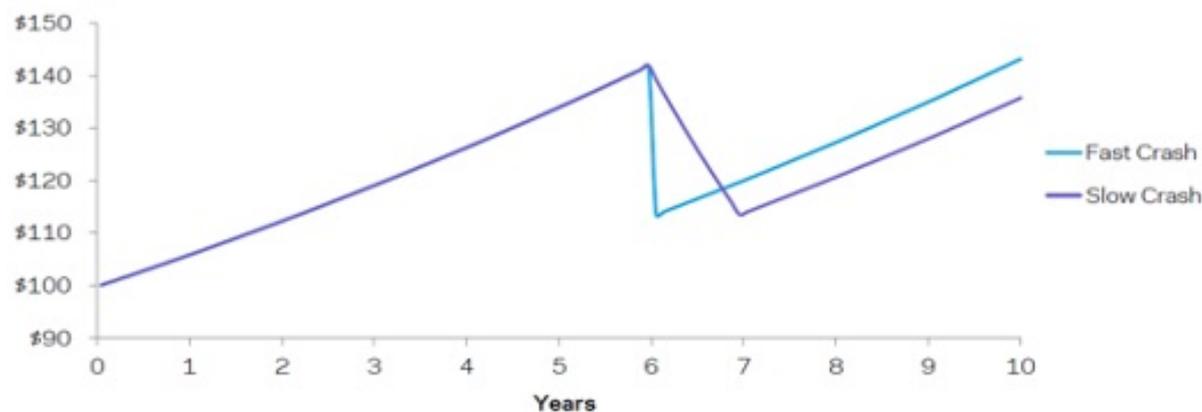
Huh?

Well, which would you choose? And what on earth does this illustrate?

Lest we misunderstand what they are saying, they provide a graph:

Exhibit 1: Even for the Same Magnitude Loss, Slower Losses May Be Worse

Two Hypothetical 20% Losses



Source: AQR. Both series have 6% annual return, except for the "Fast Crash" which loses 20% in one month and the "Slow Loss" which loses 20% over one year (at all other points, the series each realize 6% annual returns). Both losses start at the beginning of Year 6 (though could start at the beginning of any year shown). For illustrative purposes only. Hypothetical data has inherent limitations, some of which are disclosed in the Appendix.

The graph plainly shows that the interpretation I put on it is correct: In one scenario, a 20% loss occurs at the beginning of year six after which growth at the annualized rate of 6% continues for the remainder of the year, resulting in a loss of 15.6% for year six; in the other scenario, the loss in year six is 20%.

The paper credits a list of people for "helpful comments." The list includes Asness and other AQR luminaries. None of them, apparently, were disturbed by the vacuity of this example. Furthermore, several articles written by finance industry professionals, including two by Brown, were later published that cited the AQR paper. None of them mentioned it either.

This is an example of why the top-performing hedge fund company, Renaissance Technologies, run by and staffed by mathematicians, will never hire anyone with a finance background

The rest of the April AQR research study

In the rest of the AQR paper, one can sense the authors struggling to construct and present examples that will shore up the point they wish to make, which is that while put options can protect you against short-term drops in the market, they won't protect you in the long term as well as the risk mitigating strategies that AQR offers.

The problem is, while they provide performance data for a hypothetical options-based portfolio, they don't specify the options-based portfolio in any detail.

Compare the disclosure in AQR's paper with the disclosure I provided above about the backtest that I performed. Based on my disclosure, it would be possible for a reader to replicate my study, and either confirm or refute its results.

No such replication is possible based on the information in AQR's 19-page paper. It doesn't explain how they constructed the options-based portfolio that they studied, either in the text of the paper or in its appendix. In other words, they imply that you must take their word for their results; it's not possible for you to check them.

Ordinary, healthy skepticism, which should be the norm for reading and studying research such as this, would not be assuaged by this presentation. And in light of the simplistic and mathematically challenged example given at the beginning of the paper, skepticism reigns dominant – especially given that AQR has every incentive to find results that favor the strategies that they offer, like risk parity and trend following, and to denigrate competing strategies.

The July AQR paper

AQR followed up on July 8 with another paper, "Tail Risk Hedging: Contrasting Put and Trend Strategies." Its mission is again to compare tail risk strategies using puts with the trend strategy that AQR employs.

Once again the conclusion is that the tail risk strategies are inferior. This time the paper does specify what put options are used in the simulations. But although the paper is liberally footnoted, I was unable to find anywhere in it a specification of what constituted the rest of the assumed portfolio. Remember that in my simulation – as per Universa's recommendation – the put options constituted 3.33% of the portfolio, and the SPX the other 96.7%. No similar specification can be found in the AQR paper. It is impossible to confirm or refute its findings without further information.

The paper opens by noting that, “Many investors fear sharp market declines.” Given this introduction, the option strategy that it chooses to highlight is a strange one, namely: purchase of one-month 5% out of the money put options. This is a very expensive way to protect against sharp and deep market declines. AQR’s results show this, as expected.

However, the paper does then pose the question, “What about robustness to other put strategies?” To answer that question, it presents performance for strategies using options that are as much as but no more than 20% out of the money. This is still not in the tail risk range. A –20% return for a single year is only a little more than a one and a half standard deviation event. It would not have produced a return from Feb 19-Mar 23 as extreme as the one attributed to the Universa fund, and is therefore not likely to represent the tail risk strategy that it uses. I must, however, emphasize again that I do not have any inside information about what Universa is doing and can only attempt to simulate it.

Does this show that tail risk-hedging strategies are the best way to get the highest return with lowered risk?

The data I presented above does show that over the period from 2006 to the present the specific options-based risk hedging strategy that I studied was superior to a 60/40 stock-bond portfolio.

That said, there are strong reasons for caution.

The foremost reason is the usual – historically successful investment strategies, once made known publicly and widely circulated, tend to diminish in effectiveness over time, sometimes very quickly. (This caution, in fact, has applied well to the strategies that AQR has favored.)

The strategy that I simulated, using actual historical data, involved purchasing options that may have been thinly traded. If many more investment managers decided to adopt those strategies, they would affect the prices of the options in a manner that would make the strategy less effective. This is a danger – not just a risk but almost a certainty – that afflicts almost every investment strategy that has been historically successful.

Nevertheless, there are reasons why some investors, depending on their objectives and their attitudes toward risk, may wish to consider an options-based risk-hedging strategy. If they are especially sensitive to steep downside movements in the stock market, because it makes them feel that there is no way to be sure that it won’t result in a near-total loss of capital, then they may wish to purchase the insurance that an options-based hedging strategy provides.

So long as the strategy doesn’t become too widely practiced, an options-based risk-hedging strategy has some appeal; certainly more so than the strategies that AQR offers for risk mitigation, and for better reason.

In the Twitter battle, Taleb is right and Asness is wrong. The AQR paper is, as Taleb tweeted, flawed in its argument that tail risk hedging doesn’t work.

Economist and mathematician Michael Edesess is adjunct associate professor and visiting faculty at the Hong Kong University of Science and Technology, chief investment strategist of Compendium Finance, adviser to mobile financial planning software company Plynty, and a research associate of the Edhec-Risk Institute. In 2007, he authored a book about the investment services industry titled The Big Investment Lie, published by Berrett-Koehler. His new book, The Three Simple Rules of Investing, co-authored with Kwok L. Tsui, Carol Fabbri and George Peacock, was published by Berrett-Koehler in June 2014.

¹ Universa Interim Decennial Letter April 7, 2020.

² Partial month, to May 22.