



SAFE HAVEN INVESTING - PART ONE

NOT ALL RISK MITIGATION IS CREATED EQUAL

October 2017

There is a movement today among pension funds toward systemic risk mitigation—or “safe haven”—strategies. This makes great sense as a potential solution to the widespread underfunding problem. Many pensions still haven’t fully recovered from the crash of 2008, and can’t afford another. Moreover, truly effective risk mitigation must lead to an incrementally higher long-run geometric return, or compound annual growth rate (CAGR); and a higher CAGR is the way to raise a pension’s funding level over time.

Just how does risk mitigation raise the CAGR? Modern Portfolio Theory tells us that diversification and rebalancing can lower a portfolio’s risk and raise its geometric return, and with the right covariance structure it is counterintuitively possible for its geometric return to exceed that of any of its components.

This is investing’s own theory of relativity: There is no one value that we can assign to an investment, specifically a risk mitigation investment; rather, its value is unique to a given portfolio, and is

MARK SPITZNAGEL

President & Chief Investment Officer
Universa Investments L.P.

Mark founded Universa Investments L.P. in January 2007 and has developed its unique focus on risk mitigation in the context of achieving long-term improvements to portfolio construction. His investment career has spanned over 20 years as a derivatives trader, during which he has cultivated his approach to safe haven strategies, specifically bespoke tail hedging. Mark received an M.S. in Mathematics from the Courant Institute of Mathematical Sciences at New York University and a B.A. from Kalamazoo College.



relative to how it interacts with and thus mitigates risk in that portfolio. This sound theory behind true risk mitigation shows that a portfolio's geometric return can truly be greater than the sum of its parts.

In practice, however, such attempts at risk mitigation through diversification tend instead to lower CAGRs (in the name of higher Sharpe ratios); while measured risk may be lowered, it tends to be accompanied by lower geometric returns as well. One is then forced to apply leverage to raise the CAGR back up, which just adds back a different kind of risk by magnifying the portfolio's sensitivity to errors in one's spurious correlation estimates. Diversification is unfortunately not "the only free lunch in finance" that it has been made out to be. So much risk mitigation is simply about moving from concentration (or typically *beta*) risk to levered model risk.

True risk mitigation shouldn't require financial engineering and leverage in order to both lower risk and raise CAGRs. After all, lower risk and higher CAGRs should go hand in hand! It is well known that steep portfolio losses (or "crashes") crush long-run CAGRs. It just takes too long to recover from a much lower starting point—lose 50% and you need to make 100% to get back to even. I call this cost that transforms, in this case, a portfolio's +25% average arithmetic return into a 0% CAGR (and hence leaves the portfolio with zero profit) the "volatility tax": it is a hidden, deceptive fee levied on investors by the negative compounding of the markets' swings.

The destructiveness of the volatility tax to a portfolio explains in a nutshell Warren Buffett's cardinal rule "don't lose money."


Achieving higher sustained CAGRs through volatility tax savings is the name of the game in risk mitigation; it is precisely how risk mitigation adds value. All such strategies aim to do it, but not all are created equal. They all ultimately require a tradeoff between the degree of loss protection provided (the amount of the portfolio's negative compounding that is avoided) versus the degree of opportunity cost paid by the allocation of capital to that protection rather than to the rest of the portfolio (or the amount that the portfolio's arithmetic average return is lowered). These are the two sides of the safe haven coin, and we can only measure each side vis-à-vis the other. Evaluating this highly nonlinear tradeoff is tricky, and is fraught with mathematical mistakes, as the effect on the volatility tax is often indirect or invisible. The best risk mitigation solution can be a counterintuitive one.

The hope for this piece, then, and the ones that follow in this series, is to help make these solutions more intuitive, and perhaps even change our idea about what risk mitigation is and what it can do. And analysis being performed can mostly be replicated and verified in a spreadsheet in ten minutes or so. (Some of the results are so odd that verification may be in order.)

We will thus focus on a straightforward criterion: higher portfolio-level compound annual growth rates from lower risk (or specifically from paying less volatility tax). We will use this criterion to evaluate cartoon versions of the three canonical prototypes of safe haven strategies out there, where each exhibits a very distinct protection-cost tradeoff. They are depicted in Figure 1 on the next page.

Each of the three cartoon safe haven prototypes has its simple dynamics bucketed by four corresponding ranges of annual total returns in the SPX (a natural proxy for the

FIGURE 1



systemic risk we’re trying to mitigate). Think of these as contractual contingent payouts, with no noise or counterparty risk. They represent actual safe haven strategies only by virtue of the “principle of charity”, as they are in most cases idealized, optimistic interpretations of such strategies (for instance, classical diversification doesn’t come close to mitigating risk like any of these); this of course makes their usefulness even greater in this thought experiment exploring the limits of even what the best risk mitigation can do.

STORE-OF-VALUE

The “store-of-value” prototype on the left makes a fixed 2% real return (or annuity) each year, regardless of SPX returns; it provides great diversification, with a zero correlation in a crash. This might be short-term US Treasuries (being very generous), or even Swiss franc.


ALPHA

The “alpha” prototype in the center makes a 20% nominal annual return in the crash bucket (when the SPX is down 15% or more for the year), 10% in the second bucket (when the SPX is down less than 15%), and 5% in the other two buckets; it provides a nice negative correlation in a crash, and is always positive-carry. This looks somewhat like the intended performance (and even the historical performance of the best of survivors, at least for a while) of systematic trend-following CTA strategies, “contrarian global macro” and “long volatility” strategies, or even gold.

INSURANCE

The “insurance” prototype on the right makes an explosive profit of 900% in the crash bucket, and

FIGURE 2



loses 100% in every other year (whenever the SPX isn't down by over 15%); it is highly nonlinear or "convex" to crashes (a "9-to-1 longshot"). This looks a lot like a tail risk hedging strategy (at least when done right, though most such funds seek profiles much more like the alpha prototype), and this extreme asymmetry is the touchstone for what I do as a practitioner.

Over the past 20 calendar years (an arbitrarily selected round number), the stand-alone arithmetic average returns of this store-of-value, alpha, and insurance payoff profile have been about +4% and precisely +7% and 0%, respectively. (There are two years in the crash bucket, or 10% of the data—not exactly "black swans".)

Which of these three strategies would have most

effectively mitigated the systemic risk in a portfolio and thus added the most value by improving its CAGR, historically? Let's see what the empirically correct answer is by testing three portfolios where each strategy was paired with an SPX position. We used a weighting of 90% SPX + 10% safe haven prototype in the first two cases, and 97% SPX + 3% in the insurance case. Changing the 10% allocation sizes would not have materially changed the results, and the much smaller allocation size of the insurance prototype is due to its extreme convexity. The higher a strategy's "crash-bang"-for-the-buck, the less capital it requires to move the needle and the more capital is available for the rest of the portfolio, in this case for the SPX. All are rebalanced annually, and of course the insurance allocation is replenished each year that the SPX isn't down over 15%.

Figure 2 depicts historical performance profiles of each

hypothetical portfolio over the past 20 calendar years, bucketed again by corresponding annual SPX total returns.

The blue bars are the average annual portfolio returns for that bucket, next to the SPX alone in gray, and the line plots are the ranges of annual returns.

Keep in mind that Figure 2 does not account for the volatility tax savings that continue to compound even during non-crash years—as volatility tax savings are essentially invested in the market. We have erroneously compartmentalized each volatility tax savings into a single year, and have thus hidden those savings that accumulate beyond that year. We have appropriately left it to overall, unconditional CAGR outperformance to fully capture this hidden accounting and show the true risk mitigation value added.

The portfolio with the store-of-value prototype showed some, but not much, risk mitigation in the crash bucket, though the portfolio CAGR was actually lower than SPX alone by 17 basis points. The opportunity cost of the safe haven versus owning more SPX shows itself quite clearly.

The obvious pick for most would have been the portfolio with the alpha prototype, with its 7% average return and impressive negative correlation in the crash bucket. Adding it to the SPX portfolio lowered the arithmetic return of that portfolio (since its arithmetic mean is lower than that of the SPX), but in turn also raised the CAGR of that portfolio by 18 basis points. It thus created a modest cost savings on the volatility tax. But that savings was surprisingly low, and the portfolio still realized heavy 20%+ crash losses.

How did the insurance prototype compare, with its meager 0% average return? With only a 3% allocation, the crash bucket SPX losses were almost entirely offset and hence the portfolio CAGR outperformed that of the SPX alone by 67 basis points (almost four times the outperformance of the alpha portfolio). The insurance added the most risk mitigation value to the portfolio by saving so much of the volatility tax.

To put this in perspective, in order for a 3% allocation to a store-of-value strategy to similarly raise the portfolio's CAGR by 67 basis points, that store-of-value strategy would require a fixed almost 30% nominal annual return (which would of course attract all the capital in the world).

While we just saw in the alpha prototype example that the expected geometric return of a portfolio can be greater than that of any of its component parts, it is nonetheless hard to believe the severity in the insurance prototype case. It simply runs contrary to the common perception of this type of payoff as expensive, as well as the conventional wisdom that for a risk mitigation strategy to add value it must have a positive expected return. What at first appears to gratuitously lower the arithmetic return of the portfolio (and drag on the portfolio as a line item in 9 out of 10 years) turns out to be a CAGR boon.

We can crank up the alpha allocation size further, up to about 30%, and increase that portfolio's outperformance a bit, but it still never gets anywhere close to the 3% insurance allocation's level of outperformance.

Moreover, during this time period the insurance portfolio outperformed both the HFRI hedge fund

index and a 60% SPX + 40% Treasury bond portfolio—including over the past 5 years and, remarkably, over the majority of the years in frequency.

Importantly, these results are extremely robust to the time period selected, as they don't materially change whether testing over the past 10, 20, or even all the way back over the past 100 years (though the alpha strategy's CAGR outperformance to the SPX completely disappears as we lengthen the time period).

To be clear, one needn't have assumed an explicit 97% SPX + 3% insurance portfolio in order to have benefitted from the risk mitigation as we have seen here. We have simply shown how various types of risk mitigation allocations proportionally transformed, and as a result outperformed—or not—the returns from one's systemic exposure (or *beta*). For instance, one could have moved from a 50% systemic portfolio exposure, or *beta*, to a 48.5% exposure + 1.5% insurance prototype allocation and realized the same incremental risk mitigation value added proportional to that 50% systemic exposure.

The value of a payoff with crash convexity as extreme as the insurance prototype is almost entirely derived from the perspective and context of the end user, from what it specifically does to that end user's portfolio. As a standalone strategy, the insurance prototype's value was an arithmetic average return of 0% and geometric return of -100% (both by design); it had no long-term value. As a 3% allocation added to a portfolio composed entirely of the SPX, for instance, the insurance prototype's value was an

incrementally higher portfolio CAGR, higher than allocations to either the store-of-value or the alpha prototypes tested (as well as just about any other alternative out there); this portfolio's geometric return was greater than the sum of its parts.

This is how risk mitigation adds value to a portfolio, and is its ultimate goal. Effectively achieving this goal—through an effective savings in volatility tax, or minimizing negative compounding—means achieving an optimal protection-cost tradeoff, and this tradeoff seems to thus greatly favor maximal convexity. The implications for how pension funds might best approach underfunding problems through risk mitigation are huge.

IMPORTANT DISCLOSURES

This document is not intended to be investment advice, and does not offer to provide investment advice or sell or solicit any offer to buy securities. Universa does not give any advice or make any representations through this document as to whether any security or investment is suitable to you or will be profitable. The discussion contained herein reflects Universa's opinion only. Universa believes that the information on which this document is based is reliable, but Universa does not guarantee its accuracy. Universa is under no obligation to correct or update this document.

Neither Universa nor any of its partners, officers, employees or agents will be liable or responsible for any loss or damage that you may incur from any cause relating to your use of these materials, whether or not the circumstances giving rise to such cause may have been within Universa's or any other such person's control. In no event will Universa or any other person be liable to you for any direct, special, indirect, consequential, incidental damages or any other damages of any kind even if such person understands that these damages might occur.

The information shown in Figures 1 and 2 is purely illustrative and meant to demonstrate at a conceptual level the differences among different types of risk mitigation investment strategies. None of the information shown portrays actual or hypothetical returns of any portfolio that Universa manages.
