



SAFE HAVEN INVESTING - PART THREE

THOSE WONDERFUL TENBAGGERS

December 2017

Parts One and Two of this *Safe Haven Investing* series were an objective, transparent analysis of what makes risk mitigation effective, and more specifically, when and how it adds value to a portfolio. What we've learned so far is that the shape of the protection payoff profile—or the returns of the risk mitigation strategy conditional on concurrent returns of the variables that it's protecting, such as the SPX in this case—is the most important determining feature of an effective risk mitigation strategy, even more important than its unconditional realized return or even the likelihood of steep portfolio losses. (That is, predicting anything other than the payoff profile was largely unnecessary and even counterproductive.)

Every step of the way we came back to the same optimal payoff profile that added the most value to a portfolio with systemic risk—the highly asymmetric, convex-shaped payoff of the cartoon insurance safe haven prototype. And perhaps it was the degree and consistency of that dominance that was the most counterintuitive result, and further runs contrary to how nearly every allocator or portfolio manager understands risk mitigation.

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.



So we have seen how risk mitigation can go right and how it can go wrong. The question at hand is: What is our margin of error between the two?

The defining parameter of how risk mitigation went right in the insurance prototype has been the degree of asymmetry between the amount of return generated in a crash (as defined by the SPX down by 15% or more over a year) and the loss the rest of the time. Up to this point, we have been using the “tenbagger” as our standard crash return and have been keeping the annual non-crash loss amount fixed at -100%, which means an asymmetric payoff of a 10-to-1 annual crash profit versus annual non-crash loss. “Tenbagger” is of course a reference to Peter Lynch’s writing on “Those Wonderful Tenbaggers”—a term, borrowed from baseball, that he used for “a stock in which you’ve made ten times your money.”

We have intentionally kept the assumptions in these cartoons basic to be both illustrative and able to best

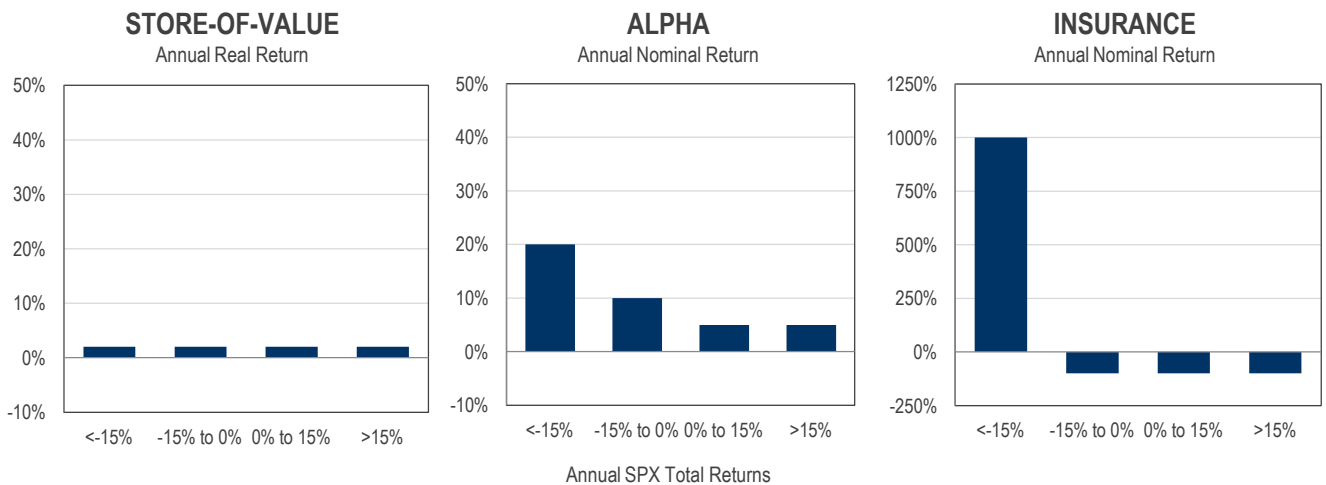
identify and isolate risk mitigation effects (as well as to allow anyone to replicate the tests with little time or effort). We will now begin to examine sensitivities to the basic assumptions in more detail here, specifically those of the insurance safe haven. In particular, we will examine the shape of the payoff as driven by the chosen crash payoff, and the chosen size of the allocation to the insurance payoff in the portfolio.

Figure 1 shows, once again, the simply-defined dynamics of those three idealized, cartoon safe haven prototypes—the “store-of-value”, “alpha”, and “insurance” safe havens.

And recall that we paired each prototype with an SPX position to form very basic test portfolios.

How did we choose the tenbagger crash payoff size for our previous testing? This payoff was chosen simply

FIGURE 1



because it provided an exact 0% arithmetic average return over the time periods tested (and it was slightly adjusted *ex post*, depending on the specific time period, to keep that average return at 0%); so this payoff thus presumed perfectly efficient and predictive forward-looking markets. (This is probably a fairly good assumption, though only investors with significant expertise could ever expect to consistently position for and realize such a tenbagger payoff profile—or better—in the derivatives markets.)


How sensitive are our results to that assumption of a 0% arithmetic return? What happens to these results if we move that around? Specifically, we can just look at the impact across a range of insurance crash payoffs (or, equivalently, a range of stand-alone insurance payoff arithmetic average returns) on the total portfolio geometric returns, or more specifically the compound annual growth rate (CAGR) outperformance of the 97% SPX + 3% insurance portfolio over the SPX alone (as we did in *Parts*

One and Two). Recall again that a higher long-run CAGR is precisely how effective risk mitigation manifests itself in adding value to a portfolio; it is all about avoiding the destructive “volatility tax” that is paid by a portfolio through the punitive negative compounding effect of large drawdowns.

The results are depicted in Figure 2 below. Each x-axis value of crash payoff (i.e., the return that the insurance prototype makes in a year when the SPX is down more than 15%) represents the payoff profile, for instance 1000% is equivalent to the tenbagger or 10-to-1 (or 1000%-to-100%) crash payoff, as represented in Figure 1.

The degree of convexity that was needed to add value as a risk mitigation strategy over any specific time period depended mostly on the frequency of the systemic losses (or the “fatness of the left tail” of the SPX return distribution) during that period. The more frequent the

FIGURE 2



losses, the greater the accumulated negative compounding effect, the greater the volatility tax charged to the portfolio's CAGR, but also the more frequent the insurance crash return profits—so the less crash payoff was needed to mitigate that volatility tax.

The insurance payoff required a minimum of about an eightbagger crash payoff (an “8-to-1 longshot”) in order for it to add risk mitigation value to the portfolio through all of the three timeframes measured (with the 100-year timeframe requiring the highest payoff). This corresponded to an annual arithmetic average return for the stand-alone insurance payoff of about -20% (versus 0% for the tenbagger). At the other extreme, only about a sixbagger was required over the past 10 years, which corresponded to an annual arithmetic average return for the payoff of about -30%.


(The next time someone says that such an insurance profile is expensive, you can respond that you could have had an average of a 30% loss, including a crash, and your portfolio would have been no worse off.)

This is the investing theory of relativity at work in the insurance prototype. Its stand-alone 0% arithmetic average return mapped to significant portfolio CAGR outperformance across time periods, all because of the profile of that 0% arithmetic average return relative to that of the SPX.

In *Part Two* we showed how the historical CAGR outperformance for all of the safe haven prototypes was much greater with respect to a 60% SPX + 40% bonds portfolio (rather than to just the SPX alone), and this is a more realistic as well as appropriate comparison as it provides a more apples-to-apples comparison in terms of the risks in each portfolio (and is closer to a more

FIGURE 3

97% SPX + 3% INSURANCE VERSUS 60% SPX + 40% BONDS



conventional allocation). This is also apparent as we observe the range of crash payoffs against the resulting CAGR outperformance of the 97% SPX + 3% insurance portfolio over a 60% SPX + 40% bonds portfolio, in Figure 3 on the previous page.


In this case, the insurance payoff required something like a sixbagger in a crash (a “6-to-1 longshot”) in order for it to add risk mitigation value to the portfolio through all of the three timeframes measured (with the 20-year timeframe requiring the highest payoff). This corresponded to an annual arithmetic average return for the stand-alone insurance payoff of about -25%. At the other extreme, only about a fourbagger was required over the past 100 years, which corresponded to an annual arithmetic average return for the payoff of about -55%.)

There is a rather linear relationship between the crash payoff (or the stand-alone arithmetic average return) and portfolio CAGR outperformance, and one could argue all day whether it is a particularly flat linear function or not. (The flatter it is, or the lower its slope, the less sensitive the outperformance is to the payoff, and the less possibility of over-fitting or “data-mining” in our tests.) What matters to me, as a practitioner, is that the function provides a very generous margin of cushion for adding risk mitigation value to a portfolio. (I would warn again, however, that it is very likely an insufficient cushion for a naïve replicator of the insurance safe haven prototype.)

Also of significance is the difference between the flatness of that outperformance function of each of the three safe haven prototypes, depicted below in Figure 4.

Over that last 20 years, both the store-of-value and the

FIGURE 4



alpha prototypes have been almost the same in terms of the risk mitigation value that they have provided to an SPX portfolio across their ranges of stand-alone arithmetic average returns. Moreover, their outperformance was much more sensitive to that arithmetic average return than was that of the insurance prototype.


The insurance prototype's value added was much more robust to its assumed return parameter, as well as more elevated for any return assumptions except exceedingly and unrealistically high returns (approaching a stand-alone arithmetic average return of 20%). Any presumption of market efficiency (or 0% arithmetic average returns) was highly damning to the store-of-value and the alpha prototypes as worthwhile strategies; not so to the insurance prototype.

The arithmetic average returns for the store-of-value and the alpha prototypes were adjusted by simply adjusting the returns equally in each of their respective SPX buckets

(from Figure 1). For instance, an alpha prototype with a 12% arithmetic average return over the past 20 years (rather than 7%, as in Figure 1) had a 25% crash return and 10% returns whenever the SPX was positive. (Incredibly, 10% allocated to this payoff still added only about the same value to the SPX portfolio as did 3% allocated to our original insurance prototype with a 0% arithmetic return.)

A wonderful consequence of the wonderful tenbagger, or the extreme degree of convexity of the insurance payoff, is the very small allocation size required of that payoff in order to move the risk mitigation needle. The fine-tuning of this sizing, both of the insurance prototype as well as of the other two prototypes, is the other important parameter that we want to understand in terms of the sensitivity of our results. To do this, again we will stress that sizing input for the SPX + safe haven portfolios and see the resulting changes in the portfolios' respective CAGRs. The insurance prototype is depicted in Figure 5 below.

FIGURE 5



As we saw in *Part Two*, in the words of the 16th century Swiss physician Paracelsus: “The right dose differentiates a poison from a remedy.” This necessarily small 3% optimal dose of the insurance payoff—thanks to its very large “crash-bang”-for-the-buck—is such an important part of what makes it consistently add value to a portfolio whose risk it is mitigating. Up the dose too much, and it starts to subtract value.

It is commonly known in portfolio theory that the expected geometric return of a portfolio can be greater than that of any of its component parts (depending on the structure of the return covariance matrix and rebalancing). This effect of a mere 3% allocation simply takes that counterintuitive insight to its extreme. (Recall from *Part One* how a 3% allocation to an insurance prototype with a stand-alone 0% arithmetic average return created a portfolio CAGR outperformance over the past 20 years that was equivalent to the same 3% allocation size to an annual fixed almost 30% nominal return store-of-value prototype—and this bears repeating only because it might be one of the most counterintuitive observations that you’ll ever read in finance.)

We want to add just enough risk mitigation to “clip” the negative compounding of the left tail, and no more.

Conditional on the SPX being down over 15%—as we have defined the crash bucket for the safe haven prototypes—the SPX was down on average about 30% across the three time periods we looked at, and we probably would have guessed at a range of around -20% to -40% without even looking at the data. With a roughly 1000% return for the insurance prototype in that crash

bucket, we would then have assumed that the required allocation to clip that tail would be about 3%. (A 1000% return on a 3% allocation equals an incremental 30% return to the portfolio, thus cancelling the loss.) So we could have arrived at our 3% allocation rather logically and, as Figure 5 makes clear, whether it was 2% or 4% (corresponding to our -20% to -40% loss guesstimate) wouldn’t have changed the results materially. The point here is there was no precise *ex post* fit on the insurance allocation size in order to get the CAGR effect that we wanted. A napkin calculation would have done the same.


The insurance prototype focuses on its strengths, and can get out of the way and leave what it’s not particularly good at for other areas of the portfolio. It doesn’t try to be all things to all investors.

The portfolio’s small required allocation for risk mitigation leaves more capital to focus on non-crash returns, in this case the SPX. This is tough for most people to appreciate, as our mental accounting tends to prefer that each portfolio line item accomplish all tasks on its own (and we have a hard time dealing with a negative number). But we can see that such clear segmentation, when done right, clearly leads to more effective risk mitigation, and consequently higher portfolio compound returns.

This contrasts quite sharply with the other two safe haven prototypes, or specifically the alpha prototype which maxed out at a 35% allocation and whose impact on mitigating the negative compounding was such that moving to that optimal allocation level raised the portfolio CAGR only slightly. See Figure 6 on the next page.

FIGURE 6

CAGR Outperformance of Portfolio Versus SPX (since 1997)



Recall from *Part One* that the performance of the store-of-value and the alpha prototypes in adding risk mitigation value to a portfolio relative to that of the insurance prototype is largely insensitive to these allocation parameters. (We stuck to a 10% allocation to these two prototypes simply because adding too much more seemed unreasonable and unrealistic.) Moreover, their payoff profiles already represent quite rosy scenarios (notwithstanding their stand-alone arithmetic average returns required to add any portfolio value).

The important point, though, is that we have clearly not overfit the allocation sizing, as no reasonable change would have changed the relative merits of the three safe haven prototypes in this analysis.

Honest and effective risk mitigation needs to be robust to the realization of that risk. In fact,

robustness just might be the most important attribute of effective risk mitigation. One cannot rely on a black box that worked in the past based on precise, empirically dialed-in parameters, such as sizing or timing.

Effective risk mitigation needs to be able to add value within a broad spectrum of very general and logical parameters. Observing how well our safe haven prototypes, particularly the insurance prototype, have held up under these requirements has led us once again, in a highly transparent fashion, to the ways that risk mitigation can go right and how it can go wrong, and the margin of error between the two.

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 through 6 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.
