



SAFE HAVEN INVESTING - PART TWO

NOT ALL RISK IS CREATED EQUAL

November 2017

In *Part One* of this *Safe Haven Investing* series (*Not All Risk Mitigation is Created Equal*), we saw how adding a simple payoff with extreme crash convexity to a portfolio of equities was the optimal approach to risk mitigation and, consequently, higher portfolio compound annual growth rates (CAGRs). This is precisely how risk mitigation adds value. We continue here with that thought experiment, which was something of a mathematical puzzle, in order to gain further intuition about the effectiveness of different risk mitigation solutions.

While we previously evaluated risk mitigation retrospectively by unconditionally looking at different hypothetical portfolio returns relative to historical SPX total returns, we will now look at risk more prospectively. Specifically, does the risk being mitigated vary with time or regime; can we identify that *ex ante*; and does our ability or inability to identify it change the relative effectiveness of the three safe haven prototypes from *Part One*?

We will approach this using different recognizable valuation periods for equity markets, likely the best forward-looking indicator of risk. As I have shown in the past, there has been a highly significant historical connection between equity valuations and risk. It is so tempting to overthink the causes

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.



and catalysts of market crashes, but crashes have simply been the logical and natural consequence of high market valuations. (Of course, the causes of high valuations are a whole other economic matter.)

This empirical regularity is depicted in Figure 1, where SPX total return five-year “maximum drawdowns” (or the maximum cumulatively negative total return of the SPX during a five-year time period) are bucketed against the level of the cyclically adjusted price-to-earnings ratio (CAPE) at the start of each rolling five-year period, going back a century. (We used the SPX once again as a proxy for the systemic risk we’re trying to mitigate, and the CAPE because it is a fairly rigorous and generally accepted valuation measure—but we could have used the Tobin’s Q ratio instead with the same results. We used 100 years of annual data throughout, rolled monthly, in order to capture enough CAPE data points to sufficiently span the full range of historical valuation cycles.) The bars represent the median and 20th percentile maximum drawdowns in

each quartile bucket, with their 95% confidence intervals also included.

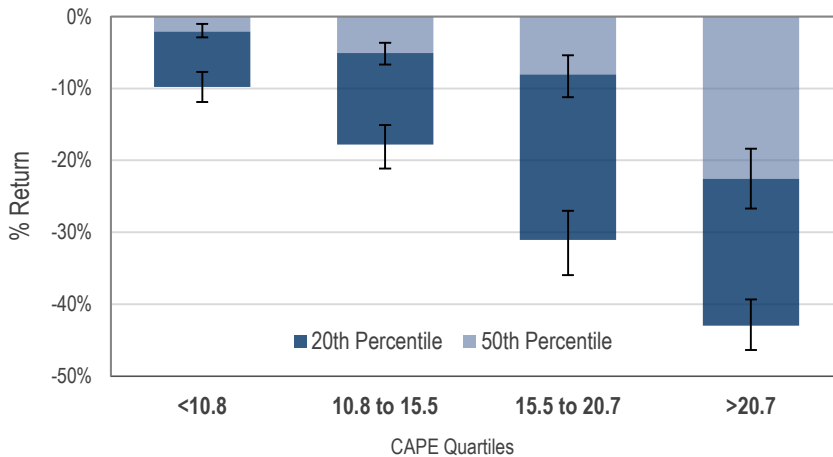
With 95% statistical significance, the higher the valuations in the stock market, the greater has been the risk (or rather the certainty) of steep market losses to come. When valuations are high (as they are today, with the CAPE measuring above 30), subsequent crashes cease to be “black swans”.

Crashes are rather the markets’ insurmountable homeostatic mechanism at work. (In addition to higher drawdowns, subsequent average returns have also been much lower with higher valuations, with 95% confidence.)

We know what these large drawdowns do to the rate at which a portfolio’s capital compounds. For instance, lose 50% and then make 100% and, despite a +25% average

FIGURE 1

Median and 20th Percentile of S&P 500 5-Year Total Return Maximum Drawdowns (since 1917)



arithmetic return, you have made zero profit and a 0% CAGR over the period; this is the result of what I called “paying the volatility tax.” Minimizing this negative compounding, or paying less volatility tax, which results in higher sustained CAGRs, is the name of the game in risk mitigation—and successful investing. So, if effective risk mitigation ultimately lessens a portfolio’s exposure to such drawdowns, it stands to reason that risk mitigation would be most effective when such drawdowns are most prevalent. Therefore, the effectiveness of a risk mitigation strategy—meaning its ability to raise the CAGR of a portfolio of equities—ought to be highly dependent on the valuation environment. And if so, doesn’t this imply an ambitious requirement of market timing skill in effective risk mitigation?

To test this expectation, we invoke again our three idealized, cartoon safe haven prototypes from *Part One*—the “store-of-value”, “alpha”, and “insurance” safe havens. This time we’ll be looking back 100 years per the CAPE data from Figure 1. During this period, the stand-alone

average arithmetic returns of this store-of-value, alpha, and insurance payoff profile have been just over +5% and +7% and precisely 0%, respectively. (9% of the annual data are in the crash bucket during this time period, and the insurance payoff in that bucket has been set to a “tenbagger” so that its average arithmetic return over the entire time period is precisely 0%.)

The three safe haven prototypes’ simply-defined dynamics are depicted for reference in Figure 2 below.

Next, we paired each safe haven with an SPX position, using a weighting of 90% SPX + 10% safe haven in the first two cases, and 97% SPX + 3% in the insurance case. (As in *Part One*, and as we’ll demonstrate later, the results are not sensitive to this 10% parameter.)

To find dependence, we simply binned each of these three hypothetical safe haven portfolios’ historical five-year CAGR outperformance (versus the SPX) by the CAPE level at the start of each corresponding five-year period,

FIGURE 2





FIGURE 3

5-Year CAGR Outperformance of Portfolio Versus SPX (since 1917)



binning by CAPE quartiles as before. The results are depicted in Figure 3 above.

Valuation levels have had very consistent implications across the different types of safe haven profiles. Most, if not all, of the safe haven portfolios' outperformance over the SPX came during the periods in the highest quartile valuation bucket. Moreover, this was particularly pronounced in the insurance prototype portfolio, where its top CAPE quartile bucket had the highest outperformance of any, by far. (As was previously mentioned in *Part One*, over the past century the alpha portfolio's CAGR actually failed to outperform that of the SPX alone, and thus the insurance prototype was the only one of the three that effectively mitigated systemic risk.)

Risk mitigation looks so much less productive in low-valuation environments that it appears to come down to a

market-timing exercise, where risk mitigation should be avoided whenever markets aren't overvalued. (Cheap markets certainly do already come with their own inherent risk mitigation in the form of a "margin of safety".) But in fact, as we saw in *Part One*, the benefits from avoiding negative compounding accumulate over time, as the savings in volatility tax continue to compound even as the market climbs—after all, the volatility tax savings are essentially invested in the market. This is not fully accounted for when bucketing 5-year outperformance data, but it is a very powerful effect nonetheless; fortunately, unconditional CAGR outperformance fully captures this hidden accounting and shows the true risk mitigation value added.

As the cliché goes, "Offense wins games. Defense wins championships." (Clichés tend to become clichés when they are true.)


Using CAGR outperformance as the scorecard by which we judge the effectiveness of risk mitigation, we can say that the insurance safe haven prototype has been the historical undisputed winner both unconditionally and during the five years following high valuation regimes. In practice, however, any risk mitigation solution will be judged relative to not only no risk mitigation at all (the SPX alone, in this case)—which makes it more of a timing decision—but rather to the conventional wisdom, best practice risk mitigation solution of a diversified, “balanced” portfolio. The most commonly assumed balancing technique in practice is a 60% equities + 40% bonds allocation (in this case we will use SPX and the 10-year US Treasury Note, rolled annually—though these results are insensitive to that choice). This ought to get us closer to an apples-to-apples comparison when measuring CAGR outperformance.

However, going back 100 years, the hypothetical store-of-value, alpha, and insurance portfolios unconditionally outperformed this balanced portfolio by 1.11%, 1.35%, and 2.34% on an annualized basis, respectively—and the insurance portfolio beat the other three here with 80+% statistical significance. (They all outperformed going back 20, 10, and 5 years as well.)

Clearly, as we should have expected, conventional wisdom has not given rise to the best risk mitigation solution.

What’s more interesting is how that outperformance was distributed according to starting valuation levels. Figure 4 depicts 100 years of five-year CAGR outperformance of

FIGURE 4



each of the three hypothetical portfolios relative to the balanced 60% stocks + 40% bonds portfolio, again bucketed by starting CAPE quartiles.

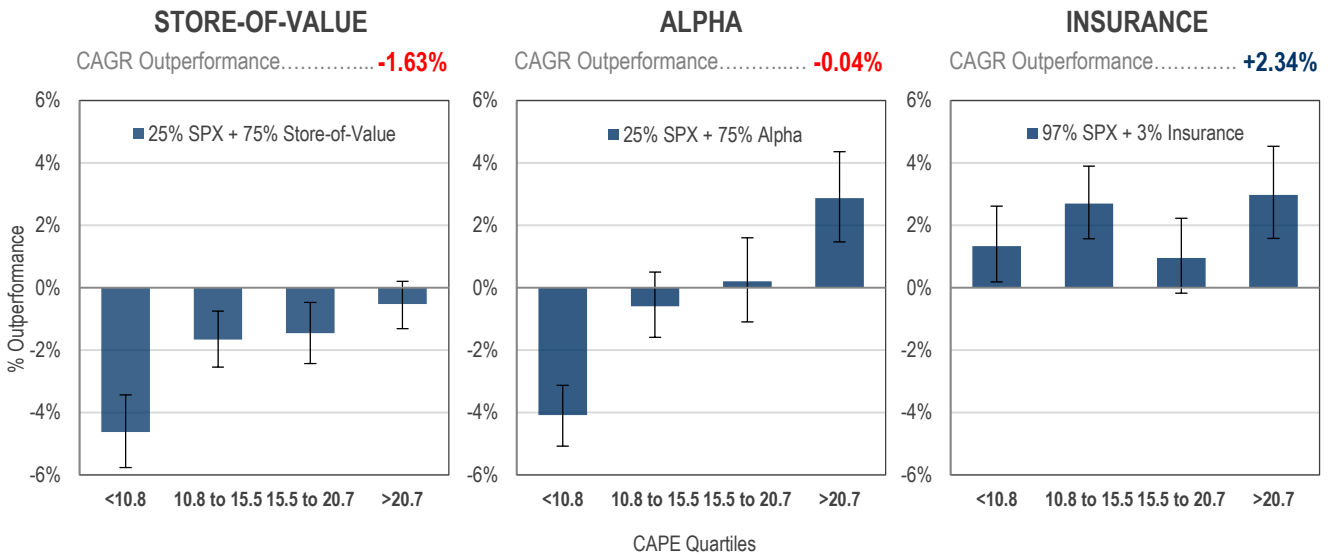
When equities were in their lowest three CAPE quartiles, all three hypothetical safe haven portfolios soundly outperformed the balanced portfolio. In fact, we cannot reject the hypothesis (with 95% confidence) that all three portfolios had the very same average CAGR outperformance over the balanced portfolio in each of these three respective lowest valuation quartiles. The big disparity came in the top CAPE quartile, where the insurance portfolio outperformed the balanced portfolio by about 3% on average, whereas the other two safe haven portfolios underperformed—a statistically significant difference with over 95% confidence.

The insurance portfolio historically outperformed the balanced portfolio quite independently of valuation levels; the outperformance of the other two safe haven portfolios depended heavily on valuation levels, with over 95% significance.

As always, we need to understand the sensitivity of these results to the chosen parameters, specifically the 10% allocation sizes, to avoid cherry-picking. By raising the allocation sizes of the store-of-value and alpha portfolios above 10%, it raised those portfolio's CAGR outperformances over the balanced portfolio in the high valuation quartile, but lowered their unconditional CAGR outperformance overall. And still, the store-of-value portfolio high-valuation CAGR outperformance never

FIGURE 5

5-Year CAGR Outperformance of Portfolio Versus 60% SPX + 40% Bonds (since 1917)



managed to turn positive. What's more, the alpha portfolio high-valuation CAGR outperformance only matched that of the 3% insurance allocation portfolio when the alpha allocation reached about 75%, in which case its overall unconditional CAGR outperformance dropped to -4 basis points (i.e., the cost of poorly predicting the crash was high). That's not effective risk mitigation. (Besides, such a high allocation to a "diversifying" risk mitigation strategy would be very impractical, even impossible, for many institutions.) See Figure 5 on the previous page.

"The right dose differentiates a poison from a remedy."
(Paracelsus)

Only the hypothetical insurance portfolio outperformed the balanced portfolio *both conditionally* (within each valuation quartile) and *unconditionally, and both on average and over the majority of the years* (going back 100, 20, 10, and even 5 years, despite only 9%, 10%, 10%, and 0% of the years, respectively, where the insurance paid off). Even more remarkably, lest a skeptic still worry that these results were skewed by 1929 or any other presumed un-repeatable historical crash, if we removed every year when the insurance ever paid off (meaning we kept only the years where the annual total return of the SPX was greater than -15%), the hypothetical insurance portfolio would still outperform the balanced portfolio *both on average and over the majority of the years* (and once again this would hold for the past 100, 20, 10, and 5 years). How was this possible?

The small 3% insurance allocation allowed the portfolio to outperform the balanced 60% SPX + 40% bonds portfolio in a crash. The remaining large 97% SPX allocation allowed it to outperform the balanced portfolio most of the rest of the time as well. The extremely asymmetric insurance payoff provided a conditional buffer to risk exposure when needed, and didn't when not needed. And this was done at a cost to the portfolio that was less than the opportunity cost to the balanced portfolio from a lower stock allocation (60% rather than 97%) and a higher bond allocation. Together, this created a surprisingly smooth path of consistent outperformance, even without a crash.

Here we have yet another highly counterintuitive, easy to miss result. When armed with information that stocks were expensive and crashes were far more likely (and that average returns would be much lower), it was optimal to be *more* invested in SPX (in a 97% SPX + 3% insurance prototype portfolio) rather than *less* (in a balanced 60% SPX + 40% bonds portfolio or even up to a 25% SPX + 75% store-of-value or alpha prototype portfolio—despite both prototypes always having positive carry). When armed with information that stocks were cheap and crashes were far less likely (and that average returns would be much higher), the *same* held true.

The hypothetical insurance portfolio was optimal no matter our information, and no matter our ability (or perceived ability) to predict a crash or not. It was thus the most universal and robust risk mitigation solution.

Incorporating the cartoon insurance payoff that required a crash in order to be profitable on its own need not have

WHICH SAFE HAVEN TYPE IS MOST EFFECTIVE FOR WHOM?

	Store-of-Value	Alpha	Insurance
Ability to Predict a Crash	✗	✓	✓
Inability to Predict a Crash	✗	✗	✓

been an expression of opinion about the likelihood of that crash. From the standpoint of model-error and non-ergodicity (i.e., “The future needn’t look like the past.”) as well as investor overconfidence, this is a very important result.

The insurance safe haven folds up neatly out of the way like an umbrella, ready to deploy at a moment’s notice when the weather turns; the other safe havens appear to stay holed up indoors in anticipation of financial storms.

Risk is time-varying, and there has been a very significant logical and empirical causal relationship between valuations and risk.

With today’s equity valuations solidly in their upper historical quartile, risk mitigation strategies are a particularly strong value proposition. But adding very high crash convexity, specifically, takes much of the

thinking and timing skill out of the decision-making process and allows for effective risk mitigation while remaining largely agnostic to potential (and even likely) systemic crises. Of course this is a tautology, as such agnosticism is the very point of effective risk mitigation in the first place.

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 2 through 5 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.
