An Introduction to Tail Risk Parity

Balancing Risk to Achieve Downside Protection

Ashwin Alankar, Senior Portfolio Manager, AllianceBernstein

Michael DePalma, CIO Quantitative Strategies, AllianceBernstein

Myron Scholes, Nobel Laureate, Frank E. Buck Professor of Finance, Emeritus, Stanford University

Much of the real world is controlled as much by the "tails" of distributions as by means or averages: by the exceptional, not the mean; by the catastrophe, not the steady drip; by the very rich, not the "middle class." We need to free ourselves from "average" thinking.

Philip W. Anderson, Nobel Prize Recipient, Physics

Executive Summary

- Tail Risk Parity (TRP) adapts the risk-balancing techniques of Risk Parity in an attempt to protect the portfolio at times of economic crisis and reduce the cost of the protection in the absence of a crisis.
- In measuring expected tail loss we use a proprietary "implied expected tail loss (ETL)" measure distilled from options-market information.
- Whereas Risk Parity focuses on volatility, Tail Risk Parity defines risk as expected tail loss—something that hurts investors more than volatility. Risk Parity is a subset of Tail Risk Parity when asset returns are normally distributed and/or volatility adequately captures tail-loss risk. Hence, when the risk of tail events is negligible, Tail Risk Parity allocations will resemble Risk Parity allocations.
- Tail Risk Parity seeks to reduce tail losses significantly while retaining more upside than Risk Parity or other mean-variance optimization techniques. It is very difficult to construct portfolios under symmetric risk measures (such as volatility as used in Risk Parity) that don't penalize both large losses and, unfortunately, large gains.
- Tail Risk Parity aims to protect investments against large losses when investors can least afford them, when systemic crises unfold and correlations spike unexpectedly. This is exactly when the marginal utility of an extra dollar is highest.
- Our research suggests that a Tail Risk Parity approach hedges the risk of large losses more cheaply than using the options market (historically we estimate savings of about 50%).
- We believe that Tail Risk Parity offers an attractive solution for investors seeking balanced investment portfolios that can cost-effectively reduce exposures to tail losses.

Introduction

It is often said that the most important decision an investor will ever make is choosing a portfolio's asset mix—more so than which individual stocks or bonds to hold. As a result, the marketplace is inundated with asset allocation products. Some are passive, maintaining a static asset mix (such as the traditional 60/40 equity/bond portfolio), while others are active. A shortcoming of passive strategies is that they fail to incorporate recent information such as changes in risks and correlations among asset classes. Hence, they pose the danger of assuming concentrated risks from time to time. Although active strategies aim to address these deficiencies, they historically take on parametric forms, such as mean-variance optimization, leading to a new set of shortcomings stemming from the difficulties in forecasting input parameters such as expected returns and correlations.

As a result of these shortcomings, investors are increasingly gravitating toward more robust risk-weighted allocation approaches that are less prone to calibration errors. In deciding what risks to weight in making portfolio allocation decisions, investors are becoming more interested in tail-loss risks. Extreme losses during crises, no matter how rare, often pose the greatest investment risk. Correlation structures change dramatically at times of shock: during the 2008 crisis, most investors experienced severe losses on multiple asset classes that were not offset with gains on other asset classes in their portfolios. That pattern has emerged during other crises too—in each case leading to potential extreme investment losses unless addressed directly.

Moreover, at times of stress, investors are often forced to incur dead-weight costs to adjust the risks of their portfolios. Often this adjustment takes the form of reducing risks first by liquidating the most liquid assets held to minimize transaction costs. This leaves the portfolio with illiquid assets and, if the crisis worsens, the liquidity of these assets all but dries up, leading investors to suffer very significant costs to reduce risks further. By exhausting the supply of liquid assets early on, the dead-weight adjustment costs are amplified if stress continues as the price of liquidity soars. Mitigation of risks in anticipation of crises would reduce these adjustment costs but might reduce returns in "normal times." We believe these alternatives—to react to crises or to plan for crises—must be part of the investment decision-making process. Currently, few investment strategies explicitly incorporate these adjustment costs into the allocation decision.

With this in mind, we developed and tested an asset allocation technology that reduces these dead-weight adjustment costs. It does this by combining the best elements of risk diversification offered by risk-weighted strategies with the loss insurance protection offered by tail-protection products. We call this innovative asset allocation technology "Tail Risk Parity" (TRP).

TRP should be thought of as an alternative to purchasing tail insurance. It provides a form of implicit insurance via its asset allocations, moving away from assets that pose greater tail risk in the future. This differs from tail insurance, which provides explicit protection, though at a potentially very high and unknown cost. We believe that Tail Risk Parity portfolios can offer a solution for price-sensitive investors seeking a balanced investment portfolio with materially smaller exposure to tail losses.

In this paper we will first briefly discuss traditional static asset allocation methods and some of their shortcomings. Second, we will introduce the concept of Risk Parity (RP), which addresses many of the key shortcomings of static asset allocation but, because it makes the crucial assumption of "normal markets," leaves portfolios vulnerable during stress environments. When the risk of experiencing a stress environment fades, TRP and RP are equivalent and hence RP can be thought of as a special case of TRP. Finally, we will describe TRP and its ability to address the shortcomings of Risk Parity.

Traditional Portfolio Construction: Static Allocation Leads to Concentrated Risk Exposures

For many years, investors were advised to choose an asset allocation consistent with their tolerance for risk and to keep this asset allocation fixed. This often boiled down to a decision about how much capital to allocate to an asset class. For example, investors might be advised to put 60% of their portfolio into stocks and 40% into bonds.

What this advice often missed, however, was that the risk concentrations of the portfolio were very different from that implied by a 60/40 capital allocation. Historically, the risk of stocks, measured by volatility, is far greater than that of bonds so, when measured as the contribution to risk, the portfolio allocation is in fact closer to 90% equity risk and 10% bond risk—not nearly as well diversified as investors might have believed (*Display 1*).

Display 1

Stocks Riskier than Bonds
1926–2012

	Equities	Bonds
Excess Return	6.0%	1.7%
Volatility	19.1%	5.2%
Frequency of Negative Returns	38.4%	33.2%
Maximum Drawdown	(84.4)%	(48.3)%

Past performance does not guarantee future results.
As of June 30, 2012
Equities are represented by the S&P 500. Bonds are represented by the Barclays Capital US Aggregate Index
Excess returns are above the 3-month Treasury bill rate.
Source: Barclays Capital, Bloomberg, S&P and AllianceBernstein

Risk Parity to the Rescue

To address a static allocation's limitations when it comes to balancing risks, risk-weighted allocation techniques, such as Risk Parity (RP), have been introduced. Proponents of Risk Parity argue that it is a better approach to asset allocation because it offers true diversification of portfolio risk. Rather than assets being diversified by dollar amount in a portfolio, under RP each asset class contributes equally to the overall volatility of the portfolio. Mathematically this can be expressed as, for all assets *i* and *j* of a portfolio:

$$\frac{d(\sigma_{portfolio})}{d(w_i)} * w_i = \frac{d(\sigma_{portfolio})}{d(w_j)} * w_j$$

Under this scheme, no asset class contributes more than any other to the total portfolio volatility.

In theory, assuming that volatility is the true measure of risk, a Risk Parity portfolio aims to adjust to capture the changing dynamics of risk. Unfortunately, volatility is a crude measure of risk with significant shortcomings. History has shown that the dynamics of portfolio volatility are not continuous. In other words, volatility can spike with no forewarning and this discontinuity makes it very difficult for RP to adjust in a timely fashion. Volatility spikes largely stem from jumps in correlations—the product of “cascades” of liquidation as investors cut risk simultaneously during a crisis. This leaves a Risk Parity portfolio unprotected at these times.

Nevertheless, historically RP has produced better risk allocations than the standard static 60/40 allocation. *Display 2* summarizes a performance study by Asness, Frazzini and Pedersen on portfolios allocated between US bonds and equities from 1926 to 2010. The table illustrates that RP outperformed a static 60/40 portfolio on a risk-adjusted basis. Additionally, the results in Display 2 show that it is very difficult to create a portfolio whose returns exhibit large positive skewness (in other words, having more upside than downside) using volatility-based allocation procedures (a theme we’ll return to later). Volatility management penalizes large gains and losses equally and, as a result, reins in both downside losses and upside gains and creates a portfolio payoff that has zero skew. Hence it is not surprising that the RP portfolio has a skew near zero and less than the 60/40 allocation (although neither portfolio exhibited much skewness).

Another argument in favor of Risk Parity is that it is actually a special case of mean-variance optimization. If all assets have the same pairwise correlations and Sharpe ratios (expected excess returns to standard deviations), then Risk Parity portfolios would be “Markowitz” efficient in that each would have the highest expected return for a given level of standard deviation of return.[1]

Display 2

Historical Annualized Performance of 60/40 and Risk Parity Portfolios
1926–2010

Statistic	60/40	Risk Parity
Excess Return	4.65%	2.20%
Volatility	11.68%	4.25%
Skewness	0.2	0.05
Sharpe Ratio	0.4	0.52

Past performance does not guarantee future results.
Excess returns are above the 1-month US Treasury bill rate.
Source: Asness, Frazzini and Pedersen; see References.

Moreover, although Risk Parity requires estimation of correlations among assets (a difficult task, as we know, because correlations can unexpectedly change very rapidly) it does not require specification of expected returns. This is because it assumes that returns on asset classes are proportional to their relative volatilities.

---

[1] *It can be argued that these assumptions imply all assets are redundant. For example, if two companies merge to form a new company, for this new company to have the same Sharpe ratio as all the other remaining assets, the correlation between the original two companies has to be one. Hence, extending the argument of equal pairwise correlations, the correlation among all companies has to be one.*

We believe that these assumptions make the approach more robust and less prone to calibration error than techniques that require an estimate of expected returns.[2]

However, a contentious assumption in constructing Risk Parity portfolios is "normal" or "well-behaved" markets. This is implicit in RP's use of volatilities and correlations as primary inputs. Volatility, by definition, measures normal dispersion. Correlation measures the degree of linear dependence between two variables. Theoretically there is no support that asset payoffs should be normally distributed, and economics would suggest otherwise. For example, risk premiums are dynamic, changing over time, with the result that returns on assets are shaped by changing means, variances and correlations. Over very short time horizons, perhaps days or weeks, investment decisions might reasonably be made assuming a normal distribution with known and constant expected returns, volatilities and correlations. But the longer the investment horizon, the more tenuous those assumptions become.

We know from experience that most reduced-form marginal distributions of returns on risky assets appear to have "fat tails." That is, there are too many returns in the extremes and too few in the middle of the empirical distribution when compared to a normal distribution (in other words, the probability of an extreme outcome is higher than that implied by a normal distribution). *Display 3* highlights a few tail events in the equity markets since 1976.

Display 3

Tail Risk Events in the US Equity Market

Realized Monthly Returns of S&P 500

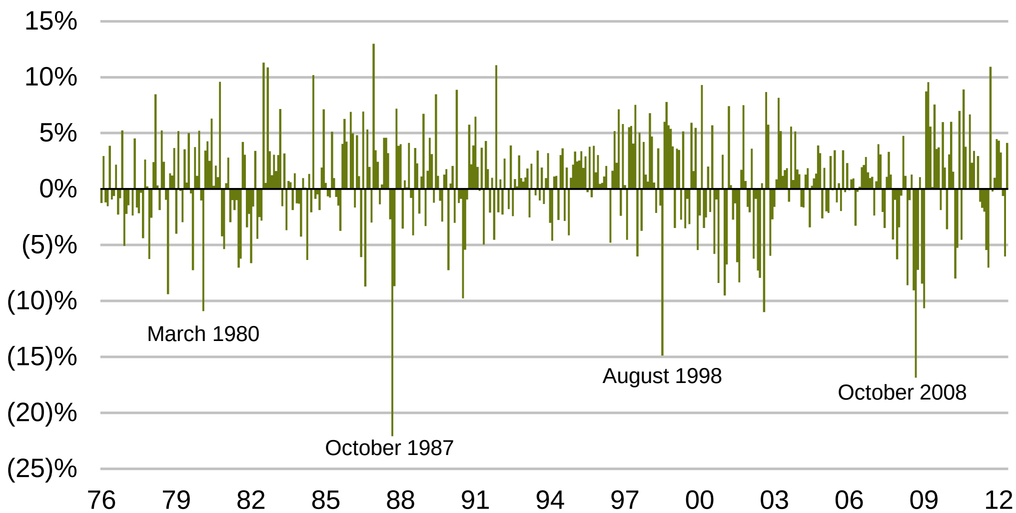

Past performance does not guarantee future results.
Through June 30, 2012
Excess returns are above the 3-month Treasury bill rate.
Source: Bloomberg, S&P and AllianceBernstein

---

[2] *In a two-asset case setting, an estimate of correlation is not required and Risk Parity weights are proportional to 1/volatility of each asset.*

In addition to stocks, other asset classes, such as corporate bonds, also exhibit asymmetric return behavior—in other words, a fat left tail and a truncated right tail. This is inherent in their economic structure: they have limited upside (given that bonds are redeemed at par, investment-grade bond holders don't expect prices to double or triple the way stock prices can) but they can suffer 100% loss of capital on the downside. A normal distribution cannot describe such asymmetric payoffs and neither can symmetric risk measures, like volatility, adequately reflect their risks. This necessitates the use of asymmetric distributions and risk measures, such as expected tail loss.

There are many reasons why tail events in the real world are more common than those implied by a normal distribution. One is that crises lead to cascades of liquidations and a "flight to safety"—particularly over short and medium time horizons. Cascades of selling cause the historical measured correlations and volatilities to understate the risk and correlations among asset classes during a crisis. As a result, the advantages of diversification in non-crisis periods are lost in crisis periods. In a crisis, asset classes that ordinarily tend to move somewhat independently, move together, leading to large drawdowns.

Display 4 illustrates this by plotting the monthly returns of the S&P 500 against various asset classes. While representative asset classes, such as high yield, commodities, and foreign-exchange carry, all draw down together, they show very little correlation otherwise.

Display 4

How the S&P 500 Has Behaved Versus Other Asset Classes
July 1983–June 2012

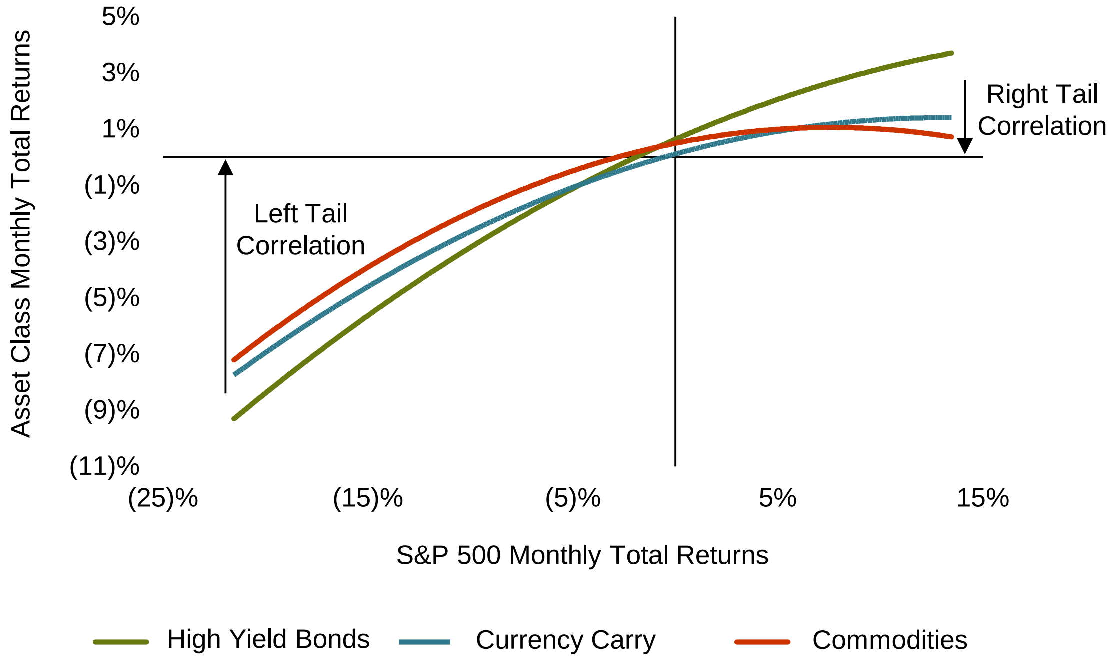

Past performance does not guarantee future results.
Through June 30, 2012
Source: Barclays Capital, Citibank, Dow Jones/UBS, S&P and AllianceBernstein

As explained above, Risk Parity deals well with some of the important flaws inherent in static asset allocation, or even mean-variance based allocation techniques. Importantly, it addresses the concentrated risk exposures that result from static approaches and eliminates the need to forecast expected returns. However, by focusing on volatility as a risk measure and assuming normal correlations, Risk Parity is ill equipped to deal with stress losses and liquidity-driven drawdown events—arguably the most important risks faced by investors. Tail Risk Parity is designed specifically to form portfolios that hold up well during these periods of high stress. The TRP approach subsumes the costly drawdowns or jumps to which the RP approach is susceptible by focusing on the jumps themselves.

Tail Risk Parity Navigates Turbulent Markets

Our Tail Risk Parity methodology generalizes the standard Risk Parity allocation strategies by forming portfolios that diversify tail risk exposure and reduce reliance on ordinary correlations. Our methodology seeks to provide diversification benefits by allocating capital across assets such that each asset contributes equally to the overall expected tail loss (ETL) of the portfolio.

Expected tail loss measures a portfolio's expected return over a pre-specified horizon in the "worst possible" scenarios, where "worst" is defined by a user-defined percentile. For example, a 5% monthly ETL corresponds to the expected average of the 5% worst monthly returns. ETL captures the risk of large losses or drawdowns that traditional measures that rely on a normal distribution underestimate. Both the return horizon and level of expected loss can be set based on investor preferences so that TRP is focusing on the level and duration of loss that investors are most concerned with. *Display 5* shows a fat-tailed distribution of returns. The position of the vertical dashed line is determined by the percentile below which returns are defined as tail losses. The ETL is the average value of the returns in the shaded region of the distribution to the left of the dashed line.

Display 5

Expected Tail Loss Captures the Risk of Large Losses

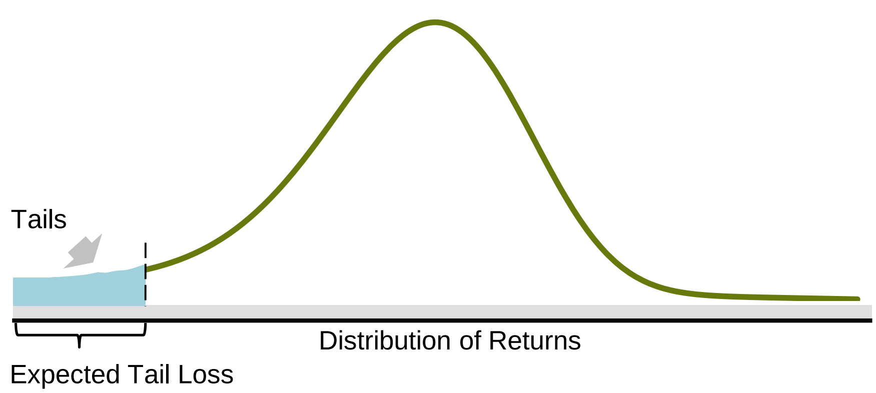

Source: AllianceBernstein

If the tails of each asset diversify, or assets do not show fat-tailed behavior, the distribution of portfolio returns is near normal and exactly normal if each asset is normally distributed. In these cases, volatility captures the risk of tail losses and Tail Risk Parity is equivalent to Risk Parity. Hence, Risk Parity can be thought of as a subset of Tail Risk Parity when the portfolio returns are normally distributed.

When investors experience drawdowns, they seek to reduce risk at a time when the costs of doing so are extremely high. The Tail Risk Parity approach aims to ensure that the portfolio has less tail loss risk exposure going into a crisis, in order to avoid deadweight high-liquidation costs resulting from forced portfolio risk reduction.

Tail Risk Parity gives up some returns (pays a small price) in order to reduce tail losses. From a theoretical point of view, this price is not a direct cost. By this we mean that this approach changes the shape of the distribution of outcomes by reducing some gains in good times for less loss in bad times. To measure these "costs" we would need to assess the utility differences between alternative distributional shapes and outcomes—a difficult task. Investors tend to value downside protection more than an equal gain on the upside. This is generally true except in cases where investors have very specific utility functions or distributions of returns are normal. For these reasons, we developed a methodology that focuses on how to mitigate tail losses at lowest cost. We think that a Tail Risk Parity approach can be useful to investors who use leverage in their portfolios or have capital calls, as their biggest worry is the risk of tail loss and the risk that creditors/beneficiaries force liquidations at the most punitive times when costs are high.

As discussed above, during a crisis, correlations among asset classes become very high as portfolio liquidations work themselves through the market—so-called liquidity-style events. While normal models such as Risk Parity ignore the possibility and effects of liquidity crises on returns and risks, the TRP portfolio is explicitly designed to take this diversification failure into account.

It does so by focusing on a portfolio's tail loss, which often is realized when diversification fails, such as during a crisis. In a crisis, a portfolio's tail loss essentially becomes the *aggregation* of the tail loss of each component asset given the high tail correlations across all assets in the portfolio. This contrasts with a portfolio's volatility during "normal" times, which is mostly driven by the *offsetting or sharing of capital* given lower ordinary correlations among the assets. Should correlations increase, volatility can provide a misleading measure of risk. There is little sharing of capital in a crisis; each asset class tends to stand on its own. Hence, portfolio tail loss incorporates diversification failure, while portfolio volatility does not.

Anecdotally, it is exactly during times of low volatility and low correlations that leverage and imbalances in the capital markets build up. This buildup increases the likelihood of correlation spikes, diversification failure and left tail events when liquidations are required to reduce the accumulated leverage. Observed low levels of volatility lead investors to believe that risk is low and will remain low—potentially followed by a rude awakening as cascades of selling unfold. This once again highlights the flaw in relying on volatility as a measure of portfolio risk. One example is the "Goldilocks" environment of 2006–2007[3], when observed low equity volatility in fact represented the "calm before the storm." In retrospect, at that time investors should have been reducing their equity risk levels rather than increasing them, as most did. In our simulations, TRP did in fact reduce equity exposures despite volatility being near record lows, because it was registering very

---

[3] *Another example would be the build-up of the "internet bubble" from 1996 to 1997. This was an environment where equity volatility was increasing due to "fat" expected right tails, not left tails. Downside risks were muted as signaled by decreasing option skews for example, despite increasing volatility. A RP manager would have reduced equity exposures in response to increasing equity volatility and missed much of the equity rally, while a TRP manager would have increased equity exposures in response to decreasing down-side risks. This example highlights another weakness of volatility as a risk measure, namely that volatility penalizes both down-side losses and upside gains equivalently and fails to distinguish between gains and losses. On the other hand, a downside risk measure differentiates between gains and losses.*

high levels of future equity tail-risk–based on signals such as the information contained in the equity option skew. Expected losses in the worst possible scenarios do not change during quiet or elevated periods of volatility; the worst possible scenario is the worst possible scenario. As a result, tail loss measures do not encourage increasing risk when measured volatility or correlation is low or decreasing it when measured volatility or correlation is high.

As we will discuss in more detail later, TRP distills information from the options market to forecast when the possibility of market stress has increased and adjusts the composition of the asset holdings proactively to protect against downside risk. Moreover, by using options-market forecasts in combination with historical inputs, the portfolio is adjusted to capture more of the upside when possibilities of stress loss are more muted.

By doing so, TRP seeks to deliver a portfolio that not only protects against large losses when systemic crises unfold, but also displays attractive returns on the upside. Our empirical work indicates that TRP portfolios experience drawdowns 20% smaller than alternative RP methodologies. And, as we'll discuss, if the investor wishes to buy put options to protect a portfolio, the TRP strategy, empirically, is more than 65% cheaper than an options-market strategy with similar drawdown protection.

In the sections that follow, we begin by demonstrating how the "tail-fatness" of asset classes affects the TRP allocation decision. After this, we will discuss how Tail Risk Parity and Risk Parity performed when hit by unexpected shocks in correlations. These unexpected shocks occur during periods of severe market stress and have profound impact not only on the current value of unprotected portfolios, but also on future values based on how a manager responds to the shock. Finally, and continuing on this theme, we will show that the implicit protection offered by Tail Risk Parity may come at a cost significantly below what comparable protection can be purchased for in the options market.

Tail Risk Parity Is Expected to Outperform When Assets Are Fat-Tailed

Tail Risk Parity will distinguish itself most from Risk Parity allocations when portfolios hold assets that exhibit seemingly low average measured volatility, but inherently pose large tail risks (for example high-yield credit and subprime mortgages that implicitly sell volatility). These asset classes offer high positive carry (investors earn a risk premium for owning them) as compensation for "crash" risk exposure. Volatility-based approaches, like RP, will tend to overweight these asset classes (due to their low volatility), while TRP will tend to underweight them (due to their high tail risks). This is illustrated in *Display 6* where we simulated the performance of TRP and RP across three assets, where the returns for each asset were drawn from a fat-tailed distribution, as outlined in the tables.[4]

The three assets were assigned return, volatility and expected tail loss properties. Each asset has a Sharpe ratio of 0.2, which is consistent with the long-term performance of the major asset classes, but assets 1 and 2 have fatter tails (as measured by ETL divided by volatility). The ETLs represent the average of the worst 10% of three-month returns. The pairwise correlations, both normal and tail correlations, among all three assets was 0.50.

[4] The "fat-tailed" distribution was modeled as a jump diffusion, whereby left tail returns were modeled as a negative jump process and, outside these tails, returns were modeled as a Gaussian process. Both the correlations of the jump process and Gaussian process across the three assets were set to 0.50. The portfolio weights and performance numbers represent average results over 10,000 return paths simulated of length 1000.

Based on those properties, the TRP portfolio allocates more money to asset 3 than the RP portfolio does, since asset 3 has a significantly lower expected tail loss than the other two assets.

The theoretical simulation results in the bottom table show Sharpe ratios (return per unit of volatility) and tail-based Sharpe ratios (return per unit of realized tail loss, or μ/TL). On a μ/TL basis, TRP outperformed RP by 19%, without giving up much of the RP return.

Display 6

Simulated Annualized Performance of Tail Risk Parity vs. Risk Parity Portfolios

Asset Return Properties

	Asset 1	Asset 2	Asset 3
Excess Return (μ)	4.7%	4.0%	2.7%
Volatility (σ)	27.3%	23.1%	14.7%
3-month 10% ETL	(39.7)%	(32.2)%	(12.3)%
Sharpe Ratio	0.2	0.2	0.2
μ/ETL	0.1	0.1	0.2
ETL/σ	3.2	3.1	1.8

Portfolio Weights

	Asset 1	Asset 2	Asset 3
Risk Parity	23.0%	27.0%	50.0%
Tail-Risk Parity	17.0%	21.0%	61.0%

Theoretical Simulation Results

	Risk Parity	Tail-Risk Parity	% Difference
Excess Return (μ)	3.5%	3.3%	(3.0)%
Volatility	13.6%	12.9%	(5.0)%
3-month 10% Tail Loss (TL)	(16.8)%	(13.6)%	(19.0)%
Sharpe Ratio	0.3	0.3	0.0%
μ/TL	0.2	0.3	19.0%

Excess returns can be thought of returns above the 3-month Treasury bill rate.
Source: AllianceBernstein; See Note on Simulation Returns at the end of the paper.

TRP Is Expected to Outperform When Tail Correlations Rise Unexpectedly

Traditional portfolio construction methods rely on low "normal state" correlations for diversification. By contrast, TRP focuses on losses during stress events, when "bad state" (tail) correlations drive performance. "Bad state" correlations are often much higher than "normal state" correlations. As a result, we think TRP portfolios are better equipped to handle shocks that increase correlations, which often happens when economic crisis hits.

In Display 6 above, the pairwise tail correlations between all three assets was 50%. Taking that as the base case, in *Display 7* we seek to answer the question of what would have happened if ex-post realized tail correlations had been higher or lower than the ex-ante tail correlation of 50%. The results are based on the same simulation framework governing the results in Display 6.

Display 7

Relative Simulated Annualized Performance Following Shocks in Tail Correlation

The Tail Risk Parity Portfolio Held Up Better

Tail Risk Portfolio vs. Risk Parity Portfolio

Tail Correlation	0%	25%	50% (base)	75%	100%
Relative TL Advantage	7%	14%	19%	24%	30%
Relative u/TL Advantage	4%	12%	19%	27%	37%

Source: AllianceBernstein; See Note on Simulation Returns at the end of the paper.

As first shown in Display 6, when the realized ex-post tail correlation was 50%, the TRP portfolio outperformed the RP portfolio by 19%. Display 7 shows that, when ex-post tail correlations increased to 100%, TRP extended its lead, outperforming RP by 37%. On the other hand, when ex-post tail correlations fell to zero, TRP shortened its lead over RP, outperforming by only 4%.[5]

These results suggest that Tail Risk Parity does a better job than Risk Parity in protecting against shocks in tail correlations. This makes intuitive sense: the higher the tail correlation, the larger the risk of tail losses to the portfolio and the more relevant TRP becomes given its focus on ETL as a risk measure, as opposed to standard volatility measures. While the merits of Tail Risk Parity are more pronounced when tail correlations are high, Risk Parity is expected to perform equally well when tail correlations are low and tails diversify. In such cases, the portfolio shows less fatness and volatility measures capture most of the portfolio risk.

[5] The return of each asset was modeled as a Gaussian random variable plus a left tail jump component. As the tail correlations become more negative, the portfolio distribution becomes more and more Gaussian because the jump component disappears. Hence, the underlying normality assumption of RP becomes correct, leading RP to "catch up" to TRP. In fact, when the tails completely diversify, the resulting portfolio distribution will be exactly normal.

Tail Risk Parity Is Expected to Provide Low-Cost Protection

While our results indicate that the Tail Risk Parity approach offers downside protection, a true test of its effectiveness is to compare the cost of this protection with the cost of a "plain vanilla" strategy of buying protection directly in the options market. We compared the return of a hypothetical TRP portfolio to the return of a 60/40 portfolio (60% S&P 500 and 40% Barclays Capital US Aggregate Index) that used at-the-money S&P 500 put options for protection (using data from 2003, when option price data became available) to assess the relative costs.

The chart on the left of *Display 8* shows the simulated annualized total returns of five portfolios: an unprotected 60/40 portfolio rebalanced monthly; 60/40 portfolios hedged with one- three- and six-month put options; and a TRP portfolio. As shown in the right-hand chart, the amount of put options purchased was set so that the portfolios using puts had the same level of protection (the same realized tail losses) as the TRP portfolio.

Display 8

Comparing the Costs of Insurance

Tail Risk Parity Portfolio vs. 60/40 Portfolio and 60/40 with Options (Simulated)

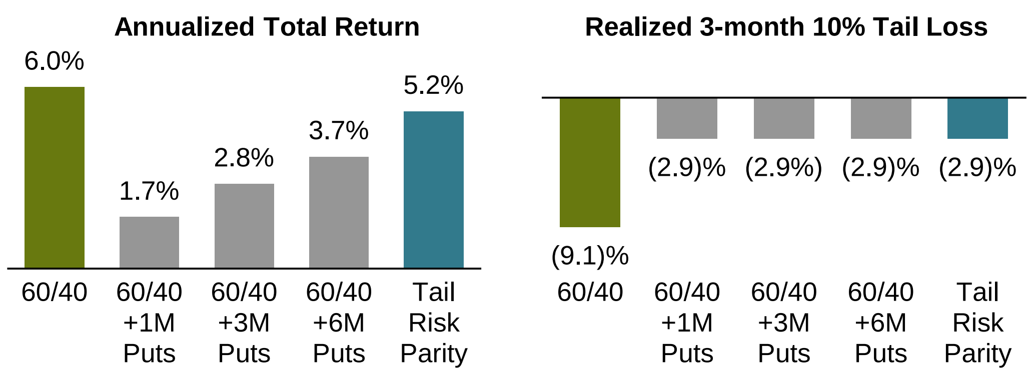

June 2003 to September 2012
Option market data for fixed income becomes available in 2003.
Source: Barclays Capital, Fama-French, OptionMetrics and AllianceBernstein; See Notes on Simulation Returns at the end of the paper.

The right-hand chart shows that the put-option and TRP approaches would all have reduced the 60/40 tail loss by two thirds. The left-hand chart shows that all of these approaches would have realized different returns. These return differences can be thought of as the relative cost of protecting the portfolio against similar losses following the different hedging strategies.

For example, the difference between the return on the 60/40 portfolio and the Tail Risk Parity portfolio is a proxy for the cost of insurance provided by TRP: about 0.8%. Portfolios that used put options for protection faced significantly higher costs of insuring against these losses, ranging from 2.3% to 4.3% a year.[6]

[6] *In fact, this is a lower bound cost since it was assumed that the buyer of the put options had the foresight to know exactly how much to purchase.*

There are three keys behind Tail Risk Parity's potential ability to offer cost-effective protection. The first is its proactive ability to reduce exposure to assets that pose greater drawdown risks by distilling information from the options market, as discussed in the next section.

The second is the likely imbalance in expected asset returns caused by managers seeking to hold high-carry (high-yielding) assets, which many do. This demand causes these assets to be over-bought, often over-leveraged, and susceptible to liquidation. While high-carry assets often have low volatility, economically the high carry represents a risk premium compensating investors for the possibility of large tail losses. High-carry assets—for example short put options—are synthetically short volatility. A short put option position earns a little amount each day (carry), but is exposed to large losses when a crisis hits and volatility spikes. In response to this large tail risk, unlike other approaches, TRP will generally underweight these carry assets. Underweighting these expensive assets helps offset the cost of protection TRP provides.

Third, by seeking to eliminate only downside losses (reducing the size of only the left tail) TRP is potentially able to retain proportionally more upside gains, which would allow it to deliver protection at significantly reduced costs. This contrasts with allocation methods that penalize both downside losses and upside gains, such as RP and other strategies that are based on volatility.

The above sections draw comparisons between Tail Risk Parity and Risk Parity through performance outcomes. In the sections that follow, we extend this comparison by contrasting the drivers behind the allocation choice of each approach. The drivers of TRP are the **tail loss** of each asset and the **tail correlations** across assets, whereas for RP the drivers are the **volatility** of each asset and the **ordinary correlations** across assets.

Tail Risk Parity Considers Left Fatness and Tail Correlations

In contrast to Risk Parity and other volatility-based allocation approaches, Tail Risk Parity portfolio weights are more dependent on the "tail fatness" of asset class returns. In *Display 9*, we followed the simulation methodology described in footnote 3 and simulated the return paths of three assets with constant volatility.

Display 9

As Drawdown Risk Increases, TRP Weightings Decline

Simulated Effect of Changing Tail Fatness in Two Portfolios (Volatility Remains Constant)

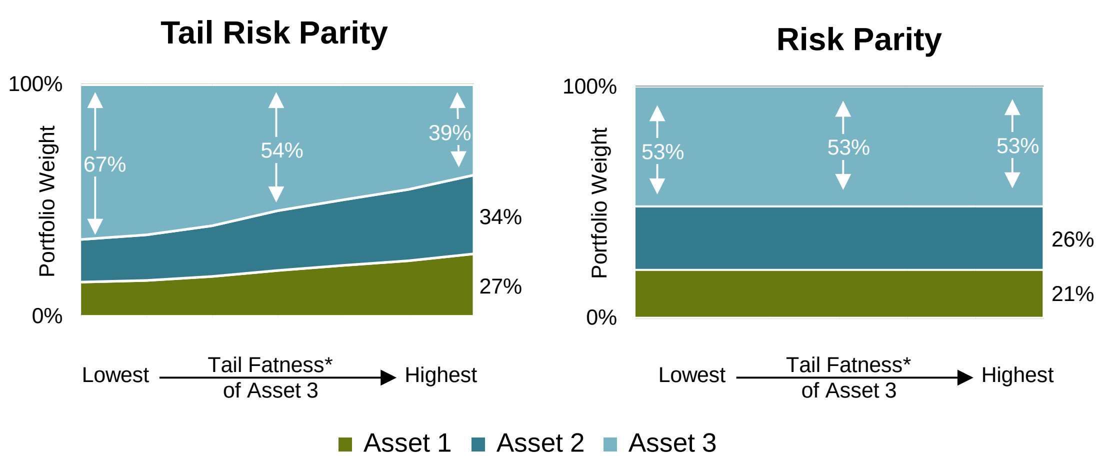

*Tail Fatness defined as (10% ETL–Mean Return)/Volatility.
Source: AllianceBernstein; See Notes of Simulation Returns at the end of the paper.

Assuming all the assets have thin left tails, asset 3 starts with a 67% weight in the portfolio. As we increase the left tail fatness of asset 3, while keeping volatility the same, TRP shifts capital away from that asset toward the others, until asset 3 has only a 39% weighting. By contrast, since RP uses only volatility measures, its exposures remain constant. Tail correlations among portfolio members are another key determinant of tail losses at the portfolio level. TRP portfolio allocations are more sensitive to tail correlations than RP allocations are.

In *Display 10*, we again simulated a portfolio with three assets, this time with changing tail correlations to illustrate this sensitivity.

Display 10

Asset Allocation Favors Assets with Lower Tail Correlations

Simulated Effect of Changing Tail Correlation (Ordinary Correlations Remain Constant)

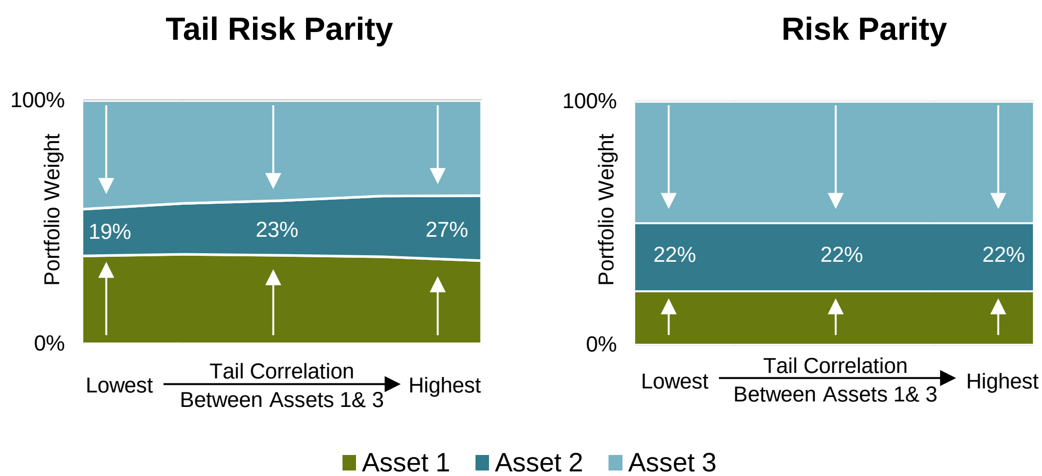

Source: AllianceBernstein; See Note on Simulation Returns at the end of the paper.

Three assets were simulated at different levels of tail correlation between assets 1 and 3, while keeping ordinary correlations constant. One would expect that, as the tail correlation of the two assets increased, TRP would shift capital to asset 2, which provides better tail diversification. The results confirm this intuition. As the tail correlation of assets 1 and 3 increases, the allocation to asset 2 increases from 19% to 27%. By contrast, again the Risk Parity allocation does not change, because it is concerned not with tail correlation but with overall correlation, which we deliberately kept constant.

Using Options-Market Data to Estimate Expected Tail Loss

To use Tail Risk Parity as a portfolio construction technology, we must estimate expected tail loss of assets. This brings us to a second innovation in the paper: extracting information from the options market to help arrive at a measure of future ETLs. We call these estimates "implied ETLs."

This is similar in spirit to using options information to gauge the market's assessment of hypothetical future volatility via implied volatilities.

One of the primary uses of options markets is to allow investors to protect their portfolios from large negative price moves. It is through the pricing and trading of options that the views of informed investors are manifested in the prices they are willing to pay for protection. The option skew reflects the severity of losses investors expect and contains information used to calculate implied ETLs.

When the marginal cost of protecting against large losses is higher than the cost of protecting against smaller losses, the skew is steep. So a steep skew implies that the market expects the underlying assets to suffer large losses with high probability. Our research has shown that these implied measures of possible drawdowns provide more accurate estimates of future realized drawdowns than methods that use historical returns only.

To assess the efficacy of implied ETLs vs. historical ETLs, we first constructed two sets of historical TRP portfolios that allocated between and the Barclays Capital US Aggregate Index. The first set used only historical ETL estimates, and the second set used only implied ETL estimates. We then compared the historical performance of these two sets of portfolios as well as how well each portfolio balanced each asset's contribution to the overall realized tail risk of the portfolios. As explained previously, balancing the tail loss risk of a portfolio is the goal of TRP, just like balancing the volatility of a portfolio is the goal of RP. On a risk-adjusted basis, the implied ETL portfolio outperformed the historical ETL portfolio by nearly 40% (Display 11). Furthermore, the contribution to the overall realized tail risk of the implied ETL portfolio was balanced, with Barclays Aggregate contributing 47% and equities contributing 53%. On the other hand, the contribution to the overall tail risk of the historical ETL portfolio showed no balance, with Barclays Aggregate contributing 1% and equities contributing 99% (*Display 11*).

Display 11

Comparing Ways of Estimating Tail Losses
Implied Measures Appear to Be More Accurate

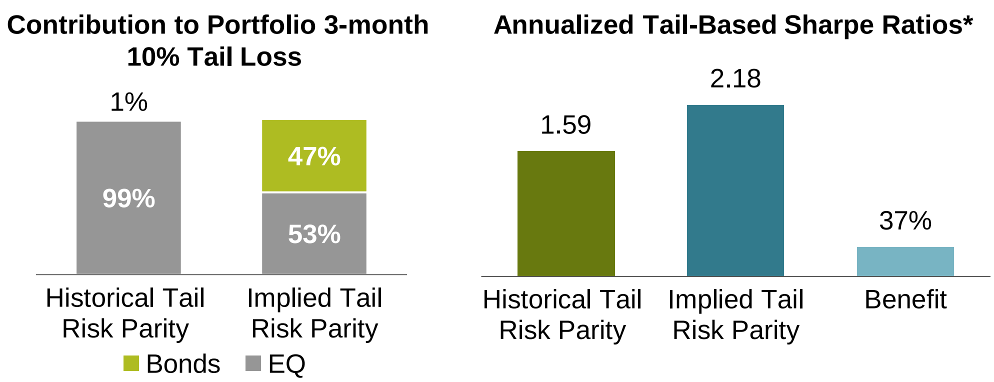

From June 2003 to September 2012
*Tail-based Sharpe Ratios are defined as return per unit of average tail loss.
Source: Bloomberg and AllianceBernstein; See Note on Simulation Returns at the end of the paper.

Most major asset classes such as stocks, government bonds, commodities, credit and currencies have active and deep options markets, due in part to the popularity of exchange-traded funds (ETFs) and the growth of an underlying options market on these ETFs. But many others do not, requiring care when using options markets that are smaller and less active. These markets may not provide the same robust estimates of implied ETLs as deeper markets. Many of these limitations, however, can be overcome with the guidance of economic theory, while retaining the robust and rich information of the options market.

Historically, Tail Risk Parity Has Outperformed Risk Parity on a Risk-Adjusted Basis

All the analysis we have cited focuses on simulation studies to compare the performances of Tail Risk Parity and Risk Parity across myriad scenarios. Because simulations can accommodate a variety of scenarios, they are one of the most robust tools for gaining insight into the future performance of an investment strategy. In our opinion, they offer much better insight than information contained in historical backtests. History is only one path among many possible scenarios, so the way a strategy performs in a historical backtest provides a limited view. Mining historical data can lead to dangerous conclusions.

Nevertheless, historical simulations do provide insight into how a strategy would have performed. So we staged a historical horse race between TRP and RP, allocating between four asset buckets (equities, Treasuries, credit and inflation)[7] starting in 2007, when our historical options data became available for US TIPS and commodities (*Display 12*).

---

[7] *The equity bucket consisted of the S&P 500; the Treasury bucket consisted of 7 to 10-year US Treasuries; the credit bucket consisted of 60% US investment-grade credit and 40% high-yield credit; and the inflation bucket consisted of 50% US TIPS and 50% Dow Jones UBS Commodity Index. TRP and RP constructions both used implied ETL and implied volatility information.*

Display 12

Historical Annualized Performance of Tail Risk Parity and Risk Parity Portfolios

March 2007–September 2012 (Simulated)

Statistic	Risk Parity	Tail Risk Parity	Advantage
Tail-Based Sharpe Ratios			
Excess Mean/10% 3-Month Tail Loss	0.8	1.0	17.3%
Excess Mean/5% 3-Month Tail Loss	0.5	0.6	21.0%
Sharpe Ratio	0.8	0.8	7.5%

Excess returns are above the 3-month Treasury bill rate.
Source: AllianceBernstein; See Note on Simulation Returns at the end of the paper.

Based on our simulation studies, we expected the TRP portfolio to show smaller tail losses than RP, but retain proportionally more of the upside, yielding higher risk-adjusted returns in tail loss space. Consistent with expectations, the TRP portfolio achieved significantly higher tail-based Sharpe Ratios of about 20% (return per unit of tail loss). The advantage that TRP realized over RP in this backtest is consistent with the relative performance between TRP and RP, simulated over tens of thousands of paths as described in the section above ("Tail Risk Parity Is Expected to Outperform When Assets Are Fat-Tailed.")

Conclusion

We believe there are several properties that make Tail Risk Parity a desirable portfolio construction methodology for balancing portfolios. These include:

- Focus on Losses

  Tail Risk Parity measures risk as loss, a risk that investors are really concerned about, rather than dispersion, a volatility risk measure. Instead of reducing dispersion by bringing in both the right and left tails equally, TRP aims to alter the return distribution by bringing in the left tail losses. This is done at the expense of some, but not all, of the upside. Economically, this delivers in a "bad state" of the world, when the marginal utility of an extra dollar is highest, at the expense of an extra dollar in a "good state" of the world, when the marginal utility is lower.

- Diversification of (Less Concentrated) Tail Risk

  A TRP framework will generally have a more diversified tail risk allocation than most portfolios, including static 60/40 portfolios and Risk Parity portfolios. Portfolios that have concentrated exposure to tail losses in their risk allocation will be left particularly vulnerable to stress events.

- Resilient Against Spikes in Correlation

  TRP proactively allocates between assets based on expected heightened correlations among assets that characterize stress conditions. As a result, a TRP portfolio is constructed to protect and preserve capital when correlations rise.

In summary, Tail Risk Parity is a robust portfolio construction methodology that addresses a key concern among investors: protection against material portfolio losses. By combining the best elements of parity-based diversification with the concept of tail insurance, our research suggests that investors can not only reduce the impact of severe market drawdowns more cheaply than they could by buying insurance in the options market, but also earn higher returns over full market cycles.

References

Anderson, Philip W. 1997. "Some Thoughts About Distribution in Economics." In W.B. Arthur, S.N. Durlaf and D.A. Lane, eds., *The Economy as an Evolving Complex System II.* Reading, MA: Addison-Wesley, 566.

Clifford Asness, Andrea Frazzini, and Lasse Heje Pedersen. 2012. "Leverage Aversion and Risk Parity." *Financial Analysts Journal*, vol. 68, no. 1.

Heston, Steven. 1993. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." The Review of Financial Studies, vol. 6, no. 2.

Appendix

Proof 1: Risk parity is mean-variance efficient when the pairwise correlation among all assets is equal to ρ and all assets have identical Sharpe ratios of K.

The variance of a portfolio can be written as:

$$\sigma_p^2 = \sum_n w_n^2 \sigma_n^2 + 2\rho \sum_n \sum_{m \neq n} w_n w_m \sigma_n \sigma_n$$

The first order conditions for risk parity are:

$$\frac{\partial \sigma_p^2}{\partial w_n} w_n = \frac{\partial \sigma_p^2}{\partial w_m} w_m$$

for all *m* and *n*. Plugging in σ_p^2 into the first order conditions gives:

$$w_n^2 \sigma_n^2 + w_n \sigma_n \rho \sum_{i \neq n} w_i \sigma_i = w_m^2 \sigma_m^2 + w_m \sigma_m \rho \sum_{i \neq m} w_i \sigma_i \qquad \text{EQ. 1A}$$

for all *m* and *n*. The solution to EQ. 1A is $w_n \propto \frac{1}{\sigma_n}$, where the proportionality constant ensures $\sum w_i = 1$.

The mean-variance objective is:

$$\min_w \sigma_p^2 \text{ such that } \sum w_i r_i = \mu$$

Upon substituting σ_p^2 into the above objective function, the first order conditions are:

$$w_n \sigma_n^2 + \sigma_n \rho \sum_{i \neq n} w_i \sigma_i - \lambda r_n = 0 \qquad \text{EQ. 2A}$$

Given $K = \frac{r_i}{\sigma_i}$ and multiplying by w_n, EQ. 2A can be re-written as:

$$w_n^2 \sigma_n^2 + w_n \sigma_n \rho \sum_{i \neq n} w_i \sigma_i - \lambda K \sigma_n w_n = 0 \qquad \text{EQ. 3A}$$

Since EQ. 3A has to hold for all *m* and *n*, the following condition also has to hold:

$$w_n^2 \sigma_n^2 + w_n \sigma_n \rho \sum_{i \neq n} w_i \sigma_i - \lambda K \sigma_n w_n = w_m^2 \sigma_m^2 + w_m \sigma_m \rho \sum_{i \neq m} w_i \sigma_i - \lambda K \sigma_m w_m \qquad \text{EQ. 4A}$$

One solution to EQ. 4A is $w_n \sim \frac{1}{\sigma_n}$, where the proportionality constant ensures $\sum w_i = 1$. Using EQ. 3A with this solution gives:

$\lambda \propto \frac{1+\rho N}{K}$, where N is the number of assets.

Hence the risk parity solution and the mean-variance solution are identical.

Proof 2: Conditions under which risk parity is mean-variance efficient, i.e. pairwise correlation among all assets is equal to ρ and all assets have identical Sharpe ratios of K, imply correlation of 1 between all assets or K = 0

Using the Capital Asset Pricing Model (CAPM) framework, the excess return on asset *i* can be written as:

$$E(r_i) = \beta_i * E(r_m) \quad \text{EQ. 5A}$$

where r_m is the excess return on the market and $\beta_i = \frac{cov(r_i, r_m)}{\sigma_m^2} = \frac{\rho_{i,m*}\sigma_i}{\sigma_m}$.

The Sharpe ratio, K_i, of asset i can then be written as: Re-writing EQ. 5A:

$$K_i = \frac{E(r_i)}{\sigma_i} = \frac{\beta_i * E(r_m)}{\sigma_i} = \frac{\rho_{i,m} * E(r_m)}{\sigma_m} = \rho * K_m \quad \text{EQ. 6A}$$

where the last equality results from the condition the "pairwise correlation amount all assets is equal to ρ". Hence for "all assets to have identical Sharpe ratio of K, either ρ = 1 in EQ. 6A, yielding $K_i = K_m$, or $K_i = K_m = 0$. The former implies all assets are redundant and the later implies all assets have zero excess returns.

Note on Simulation Returns

Hypothetical, Back-Tested or Simulated Performance. Hypothetical back-tested performance shown herein is for illustrative purposes only and does not represent actual performance of any client account. AllianceBernstein L.P. did not manage any accounts using the portfolio composition for the periods shown and does not represent that the hypothetical returns would be similar to actual performance had the firm actually managed accounts in this manner.

Hypothetical, back-tested or simulated performance has many inherent limitations only some of which are described herein. The hypothetical performance shown herein has been constructed with the benefit of hindsight and does not reflect the impact that certain economic and market factors might have had on the decision-making process. No hypothetical, back-tested or simulated performance can completely account for the impact of financial risk in actual performance. **Therefore, it will invariably show better rates of return.** The hypothetical performance results herein may not be realized in the actual management of accounts. No representation or warranty is made as to the reasonableness of the assumptions made or that all assumptions used in constructing the hypothetical returns have been stated or fully considered. Assumption changes may have a material impact on the returns presented. This material is not representative of any particular client's experience. **Investors should not assume that they will have an investment experience similar to the hypothetical, back-tested or simulated performance shown**. There are frequently material differences between hypothetical, back-tested or simulated performance results and actual results subsequently achieved by any investment strategy. Prospective investors are encouraged to contact the representatives of the investment manager to discuss the methodologies (and assumptions) used to calculate the hypothetical performance shown herein.

Confidential Material

The information in this document is confidential and intended for your exclusive use. By accepting this document, you agree that you will not (a) provide any information to any person other than your legal, tax, financial and accounting advisors for the purposes of making an assessment of the transaction or (b) use (or allow any of your advisors to use) the information for any purpose other than to make an assessment of the transaction. Notwithstanding anything expressed or implied to the contrary herein or in any document referred to herein, you and each of your employees, representatives and agents may disclose to any and all persons, without limitation of any kind, the tax treatment and tax structure of the transactions contemplated herein and all material of any kind (including opinions or other tax analyses) that are provided to any such person relating to such tax treatment and tax structure.

US Treasury Circular 230 Notice: Any US federal tax advice included in this communication was not intended or written to be used, and cannot be used, for the purpose of avoiding US federal tax penalties.

12-2896