

The Risk in Risk Parity: *A Factor-Based Analysis of Asset-Based Risk Parity*

VINEER BHANSALI, JOSH DAVIS, GRAHAM RENNISON,
JASON HSU, AND FEIFEI LI

VINEER BHANSALI
is a managing director
and portfolio manager
at PIMCO in Newport
Beach, CA.
vineer.bhansali@pimco.com

JOSH DAVIS
is a senior vice president
at PIMCO in Newport
Beach, CA.
josh.davis@pimco.com

GRAHAM RENNISON
is a senior vice president
at PIMCO in Newport
Beach, CA.
rennison@pimco.com

JASON HSU
is chief investment
officer at Research Affili-
ates, LLC in Newport
Beach, CA and adjunct
professor in finance at the
UCLA Anderson Business
School in Los Angeles, CA.
hsu@rallc.com

FEIFEI LI
is a director of research at
Research Affiliates, LLC
in Newport Beach, CA.
li@rallc.com

Traditional risk parity strategies are based on equal risk weighting of selected asset classes. The method seeks to equalize the risk contribution by the asset classes in an asset allocation portfolio. The intuitive appeal of this asset allocation scheme is in its improvement in asset class diversification versus the traditional strategic asset allocation approach of 60/40 (60% equity/40% bond), which is dominated by equity volatility.

In the most naïve approach, an asset-based risk parity portfolio would allocate roughly $1/\sigma_i$ to each of the N asset classes. A highly volatile asset class, like equities, would garner relatively low weight. Intuitively, the approach reduces the portfolio weight allocated to equities and therefore allocates more of its total variance to other sources of risk. However, this naïve approach ignores the covariance between the selected asset classes. As such, it is highly sensitive to the asset class inclusion heuristic. If the majority of the N selected asset classes are significantly correlated with equities, the resulting risk parity portfolio would still be dominated by equity risk. For example, if three of the six asset classes were U.S. equities, international equities, and emerging market (EM) equities, the naïve risk parity portfolio would have substantial risk allocation to equities. The portfolio would be significantly less equity-

dominated if only one of the six assets was equity like.¹

Two approaches have been taken to address the universe dependency of the naïve risk parity solution. The major product providers of risk parity strategies have taken on the approach of carefully selecting distinct asset classes into their risk parity universe. Generally, they blend all flavors of equities into one equity portfolio and a variety of investment-grade bonds into a bond portfolio, then augment the universe of assets with other seemingly unique asset classes like commodities, credit, currency, and real estate. For example, one well-known provider describes its risk parity philosophy as “targeting equal risk allocation from each of the four major risk sources: equity risk, fixed income risk, inflation risk, and credit/currency risk.” Another provider constructs its risk parity portfolio by “allocating 1/4th of the portfolio variance to global equity risk, interest rate risk, credit risk and commodity risk.”²

The alternative approach adopts a quantitative method for ensuring equal risk contribution by the selected asset classes by carefully accounting for the pairwise correlations between assets.³ Using this approach, highly correlated assets would be given lower weights than they otherwise would; this ensures a lower combined allocation to equity-like assets than would be achieved

through the naïve $1/\sigma_i$ approach. This line of research finds that while the exact risk parity portfolio does have better risk diversification than the naïve risk parity solution, the improvement appears quite modest in most relevant situations (Chaves et al. [2012a]).

DOES ASSET-BASED RISK PARITY REALLY PROVIDE “PARITY” IN PORTFOLIO RISK EXPOSURE?

In this article, we argue that the traditional asset-based risk parity, whether naïve or otherwise improved on, can still be highly concentrated in only one or two true risk exposures (especially in equity risk), and therefore be underdiversified in other risk exposures. This is true even if the direct capital allocation to equities appears small. We show evidence that an asset class-based approach to risk parity fails to achieve risk parity in the true underlying risk factor exposures. Having diversification in risk contribution from assets is generally not the same as having diversification in the primitive sources of risk underlying asset returns. An easy way to understand this argument is to think of assets as foods and risk factors as nutrients. While the body consumes foods, it actually needs the underlying nutrients to build bones and muscles. A healthy diet is not necessarily one that contains a diversified basket of foods but a diversified basket of nutrients.⁴ The advantage of using risk drivers

or risk factors to define and understand risk parity is that we properly account for the essential nutrients of our investment portfolios. This factor-based framework allows for a clearer insight and hence healthier portfolio construction (Bhansali [2011]).

Indeed, while there is a dizzy array of asset classes, only a handful of risk factors actually matter. We argue that two risk factors, driven by global growth and global inflation, largely dominate asset class risk and return. This is not surprising, since these two economic variables are indeed the dominant ones for economic and investment decisions. Using principal component analysis, we extract the orthogonalized representation of these two risk factors from a sample universe of nine conventional assets (U.S. equities, international equities, EM equities, real estate investment trusts (REITs), commodities, global bonds, U.S. long treasury, investment-grade corporate, and high-yield bonds).⁵ Exhibit 1 shows the total variance driven by the first five risk factors. The first two factors account for 68% of the variance in the co-movement of the nine assets. The other factors account for less than 10% each, confirming the assertion that two factors dominate the risk and returns of conventional assets. It is also straightforward to compute the statistical dependence or “loadings” of various assets on these core risk factors (see Bhansali [2011] and Chaves et al. [2011]).

EXHIBIT 1

Total Variance Explained by Each PCA Factor




EXHIBIT 2

Asset Class Factor Loadings for the First Two Factors




Exhibit 2 shows that the nine asset classes have intuitive loading on the two dominant risk factors. The pro-cyclical assets such as equities, commodities, REITS, and high-yield bonds exhibit significant loading on factor 1, which corresponds to the global growth risk factor. The counter-cyclical assets such as global bonds, U.S. treasury, and investment-grade corporates load heavily on the second factor, which largely maps to the global inflation risk factor. Note that high-yield and investment-grade corporate bonds have positive and economically sizeable loadings on both the global growth and the inflation factors. This makes sense because defaultable bonds are sensitive to both changes in interest rate levels and credit conditions, which are clearly impacted by economic growth and inflation. Because growth and inflation shocks are dominant risks in our capital market, it is unsurprising that most asset classes exhibit significant covariance with one or both factors. It is also unsurprising that other factors play relatively lower importance.

A RISK FACTOR FRAMEWORK FOR UNDERSTANDING ASSET-BASED RISK PARITY

We propose, in this section, a risk factor-based framework for examining the traditional asset-based risk parity approaches. For tractability, we consider only the two dominant risk factors: growth shocks and inflation

shocks.⁶ This is not an unreasonable abstraction from reality given the PCA presented in the previous section. The full intuition of our proposed analytical framework is concisely illustrated in this setup, which can be easily extended to include more than two factors.

While we can represent the growth risk factor using a number of different tradable portfolios, including one extracted from the PCA exercise, operationally, it is more intuitive and transparent if we proxy growth shocks with equity market shocks. Generally, the equity market is one of the most responsive to shocks to the underlying growth environment; disruptions to global growth instantaneously result in stock market declines and vice versa. Similarly, we proxy inflation shocks with bond market shocks. In our narrow, two-factor analytical framework, risk can be boiled down into equity-like risk, bond-like risk, and residual risk not spanned by the equity and bond factors. Other risk factors like inflation, liquidity, volatility, and momentum exist, but not in a form independent from equity and bond risk. Even if the connection of the economic variables (growth and inflation) to the investment variables (equities and bonds) may be imperfect, the resulting analyses remain extremely instructive and relevant qualitatively. Because equity and bond returns are easily accessible for investors, this analytical framework is relatively simple to execute and can be performed in spreadsheets that support multi-factor regressions.

The return of an asset class can be decomposed into the following two-factor relationship—return associated with exposure to growth (equity) risk, inflation (bond) risk, and a term (catch-all residual risk) not spanned by equity and bond factors:

$$\gamma_p = \beta_E \gamma_E + \beta_B \gamma_B + e \quad (1)$$

The total volatility of the portfolio can be decomposed in the following manner:

$$\sigma_p^2 = \beta_E^2 \sigma_E^2 + \beta_B^2 \sigma_B^2 + 2\rho_{E,B} \beta_E \sigma_E \beta_B \sigma_B + \sigma_{resid}^2 \quad (2)$$

If we allocate the covariance term equally to the bond and the equity component using the simplest approach, then variance due to the growth (equity) risk is given by

$$var_E = \beta_E^2 \sigma_E^2 + \rho_{E,B} \beta_E \sigma_E \beta_B \sigma_B$$

The variance due to the inflation (bond) risk is then given by:

$$var_B = \beta_B^2 \sigma_B^2 + \rho_{E,B} \beta_E \sigma_E \beta_B \sigma_B \quad (3)$$

and the variance due to “other” risk factors is given by the term σ_{resid}^2

Although there are several variations of the risk parity strategies offered commercially and discussed in research papers, the common stated goal involves achieving “parity” in risk exposure. That is, the portfolio variance driven by each of the key risks is comparable. In our framework, this implies $var_E = var_B$ for any well-constructed risk parity portfolio. For this analysis, we remain mute on σ_{resid}^2 . There are sensible reasons to require $\sigma_{resid}^2 >, <, \text{ or } =$ to $var_E = var_B$. When the world is truly two factor only, we would expect σ_{resid}^2 to tend to zero. If the world is actually three factor, risk parity would mean $\sigma_{resid}^2 = var_E = var_B$. More than three factors, true risk parity would imply $\sigma_{resid}^2 > var_E = var_B$.

The above two-factor variance decomposition allows us to quantitatively examine the returns of any risk parity strategy to determine 1) whether their growth (equity) and inflation (bond) risk are indeed in parity and 2) whether other orthogonal risk exposures actual exist in

the portfolio to any meaningful degree. To illustrate our decomposition framework as simply as possible, we analyze four commercially available risk parity portfolios using the S&P 500 as the equity factor and 10-year Treasury as the bond factor. We do not disclose fund names or fund company names to avoid issues with SEC regulations regarding the marketing of financial products. All four products claim to target equal allocation of portfolio variance to their selected asset categories, which represent distinct and meaningful risks in the economy. All products include equities and bonds as two of the key risk categories.

Our objective is to illustrate our simple analytical framework for researching risk parity strategies; we do not seek to provide analyses on commercial products for investors. One might think of our framework as a first step in constructing a risk factor-based risk parity index. This first step defines a simple “touchstone” by which to measure the true risk exposure in risk parity portfolios. In Exhibit 3, we display the basic fund performance characteristics. In Exhibit 4, we compute for each risk parity strategy the growth (equity) versus inflation (bond) risk exposure decomposition. We note that the sample period is extremely short for most of the risk parity strategies examined in this article. This is unsurprising as risk parity has only seen an increase in investor interest and, therefore, product offerings in the past three years. We note that RP #1 is based predominantly on a simulated history rather than actual net-of-fees fund returns. The extent to which costs and fees are taken into account is unclear.

The variance decomposition exercise shows that different risk parity strategies can have very different exposures to the two key risk factors and other risks. The first and third risk parity portfolios are dominated by equity risk exposure in the common sample, while the

EXHIBIT 3

Commercial Risk Parity Product Performance

Sep 30, 2010–May 31, 2012	RP #1	RP #2	RP #3	RP #4
Arithmetic Mean (Annualized)	13.5%	14.4%	5.9%	3.7%
Portfolio Volatility (Annualized)	11.7%	8.8%	9.2%	8.4%
Pairwise Correlation				
RP #1	1	0.75	0.93	0.61
RP #2		1	0.78	0.70
RP #3			1	0.61
RP #4				1

EXHIBIT 4

Risk Parity Portfolio Risk Decomposition

Sep 30, 2010–May 31, 2012

	RP #1	RP #2	RP #3	RP #4
% Variance in Growth (Equity) Risk	73	30	80	30
% Variance in Inflation (Bond) Risk	–6	41	–8	26
% Variance in other Risks	32	29	28	43

	RP #1	RP #2	RP #3	RP #4
Start Date	1/31/1990	6/28/1996	10/29/2010	7/31/2009
End Date	5/31/2012	5/31/2012	5/31/2012	5/31/2012

% Variance in Growth (Equity) Risk	30	27	80	34
% Variance in Inflation (Bond) Risk	10	24	–8	30
% Variance in other Risks	60	49	28	36

others are much more balanced across risks for the same time horizon. Consequently, the first and third portfolios display a very high correlation of 93%; this suggests that they have similar non-equity/non-bond risk allocations as well. However, they differ substantially in realized Sharpe ratios. This suggests that the two portfolios differ significantly in 1) their alphas from managing the asset class exposure or 2) their costs.⁷ The equity risk factor exposures exhibited by the first and third risk parity portfolios are indeed characteristically similar to a 60/40 portfolio in that equity risk dominates the portfolio variance. This is likely a result of selecting an initial universe of risk categories that are procyclical, which would load heavily on the “growth” (equity) risk. The second and fourth products display more parity in risk exposure, with nearly one-third of total portfolio variance allocated to equity risk, bond risk, and other risks. If there is only one other important risk in the economy, which is captured by “other,” then RP #2 and RP #4 may indeed have achieved “parity” in risk exposure. We note that all four RP strategies allocate about one-third of total portfolio variance to the “other” risk bucket.

We can gain additional insight by investigating the time-varying nature of the risk factor exposures by looking at the variance decomposition on a rolling five-year basis. Particularly, we are able

to examine the manager’s deviation away from his or her “neutral” risk allocation and assess whether those active deviations are successful. For brevity, we focus only on RP #1 and RP #2. Exhibit 5 displays the risk decomposition for the first risk parity portfolio (RP #1). Notice that over the life of this product, the “other” risks are dominant, with bond (inflation) risk explaining on average less than 10% of the variance over time. We also note that the equity risk and the “other” risk appear to exhibit cyclical behavior, suggesting that the portfolio manager may be timing his or her underlying factor exposure at

the expense of maintain constant risk parity. Specifically, his or her reduction in “growth” risk allocation before the global financial crisis and the subsequent rebalancing into “growth” risk appear to be successful.

Exhibit 6 shows that the second risk parity product exhibits different risk factor decomposition through time, with the risk factor exposures trending away significantly from the parity default for years until the financial crisis of 2008. The portfolio manager appears to have actively reallocated toward “inflation” risk while

EXHIBIT 5

Rolling Window Risk Decomposition for RP #1





EXHIBIT 6

Rolling Window Risk Decomposition for RP #2



reducing risk allocation to “growth” and “other.” Following the crisis, the strategy appears to revert back to a more passive and more “risk-parity-like” allocation. The active deviations appear to be enormously prescient given how poorly “growth” assets performed and how well “inflation” assets performed during the global financial crisis.

WHAT ARE THE IMPLIED RETURN ASSUMPTIONS FOR TRADITIONAL RISK PARITY PORTFOLIOS?

The risk parity approach generally remains mute on the expected returns for the asset classes included and hence on the expected return for the resulting portfolio. This is characteristic of portfolio approaches that ignore the first moments of the asset-return joint distribution and focus primarily on the second moment, such as minimum variance, low volatility, or other risk-managed portfolio strategies. This feature ought to seem somewhat unsettling for investors. Indeed, one often-cited concern associated with the risk parity approach is the potential concentration in extremely low-volatility

assets, which might carry unattractive ex ante Sharpe ratios (e.g., Treasury notes or bills today).

In this section, we propose another analytical framework to tease out the return assumptions embedded in the risk parity approach. This approach allows us to 1) assess the reasonableness of a particular risk parity portfolio from its implied expected returns for the asset classes included and 2) compute the expected returns for the risk parity portfolio. Our approach is similar to the approach used in Black and Litterman [1992] to compute inferred return assumptions on assets based on their portfolio weights. To do so, we need to fix the return of one asset so that the other asset-implied returns can be determined.⁸ Here we choose to set the Barclays U.S. Agg to a 2.8% term premium, which is derived from the long-term historical average.

We illustrate this approach using the following simulated naïve risk parity portfolios constructed from subsets of the following universe of assets: S&P 500 Index, MSCI EAFE, MSCI EM, BarCap Aggregate Bond Index, DJ-UBS Commodity Index, and NAREIT Index. We are unable to use portfolio weights from commercial products analyzed in the above section as we do not have access to the information. The simulated

EXHIBIT 7

Risk Parity Portfolios: Implied Expected Returns on Assets and Portfolio

Implied Excess Returns over Risk Free	S&P 500	MSCI EAFE	MSCI EM	Bonds	Commodities	REITs	Portfolio
60% Stocks 40% Bonds	114.7%			2.8%			70.0%
Risk Parity							
Bonds + S&P 500	12.3%			2.8%			4.6%
Bonds, Commodities + S&P 500	16.0%			2.8%	19.9%		7.3%
Bonds, Commodities, REITs + S&P 500	21.0%			2.8%	34.3%	20.6%	11.1%
Bonds, Commodities, REITs + S&P 500 + MSCI EAFE	28.6%	35.8%		2.8%	42.9%	25.4%	16.2%
Bonds, Commodities, REITs + S&P 500 + MSCI EAFE + MSCI EM	35.0%	44.1%	56.1%	2.8%	49.7%	30.5%	22.0%
Risk Parity Portfolio Weights							
Bonds + S&P 500	18.6%			81.4%			
Bonds, Commodities + S&P 500	16.0%			69.9%	14.1%		
Bonds, Commodities, REITs + S&P 500	14.5%			63.3%	12.8%	9.4%	
Bonds, Commodities, REITs + S&P 500 + MSCI EAFE	12.9%	11.0%		56.4%	11.4%	8.4%	
Bonds, Commodities, REITs + S&P 500 + MSCI EAFE + MSCI EM	11.9%	10.1%	7.7%	52.0%	10.5%	7.7%	

risk parity portfolios are nonetheless illustrative of our point regarding the potential issues arising from traditional asset-based risk parity approaches. Exhibit 3 displays various risk parity portfolios, the implied expected returns on the included assets and the resulting implied portfolio expected return.

While risk parity portfolios do not claim to be mean-variance optimal portfolios, it remains an illustrative exercise to ask about the expected excess return assumptions that are implied if investors were to hold such portfolios as “optimal.” We also include the 60/40 portfolio for comparison. Again, we note that the 60/40 portfolio was never meant to be mean-variance optimal;⁹ however, the illustration highlights immediately the appeal of risk parity over traditional asset allocation from a mean-variance efficiency improvement perspective.

Contrasting the implied equity risk premiums derived from the 60/40 and the naïve risk parity portfolio constructed from just equities and bonds, we see that the implied equity returns are wildly different: 114.7% vs. 12.3%, respectively. This stark contrast underlies the criticism levied against the 60/40 portfolio—that the equity risk so dominates the portfolio that there is little to no risk-diversification benefit from the asset allocation. An investor would have to assume extremely high equity risk premium to concentrate his or her portfolio in only equity risk.¹⁰

Exhibit 7 also suggests that the traditional asset-based risk parity approach can often imply a set of required return assumptions that are both unreasonable and inconsistent. We observe that the expected returns required for the pro-cyclical assets (equities, commodi-

ties, and REITS) increase as more pro-cyclical assets are added to the risk parity universe. As more correlated cyclical assets are added to the portfolio, the portfolio becomes more concentrated in “growth” (equity) risk, even though the asset weights appear more diversified. It would only make sense for investors to concentrate on the growth risk and forgo portfolio risk diversification if pro-cyclical assets are priced to pay enormously high returns. This implied asset return calculation exercise helps illustrate, from the “return” perspective, the potential problem associated with the asset-based risk parity approach, while our risk factor decomposition analyses demonstrate the same problem from a “risk” perspective. The common message is that diversified asset weights and diversified risk contribution from assets can often overstate the true underlying risk factor diversification. This dichotomy highlights the danger of examining concentration risk in asset allocation from an asset-based diversification perspective. It also demonstrates the danger of constructing risk parity portfolios based on assets without understanding the underlying risk factors driving the asset returns.¹¹

CONCLUSION

To conclude, our risk-factor approach to characterize portfolio risk diversification is fundamentally different from asset-based risk parity approaches. We begin by identifying the risk-factors that are stable and essential using a principal component analysis. Economic intuition, as well as statistical evidence, tells us that such risk factors are few in number. Our analytical framework is valuable in two ways: 1) it allows investors to assess a portfolio for its true risk diversification from a factor perspective and to assess whether parity in risk allocation has indeed been achieved and 2) it provides investors with a tool to examine the risk parity manager’s active decisions to deviate from their long-term risk allocation. Using our framework, we find that many traditional asset-based risk parity portfolios (actual and simulated) can often concentrate too much in just one of two risk exposures, particularly equity risk. We also show that the traditional asset-based risk parity approach can often require unreasonable implied-asset return assumptions due to its unintended risk concentration associated with its universe dependency. We highlight the danger of asset-based risk diversification measures versus a more robust factor-based risk diversification measure.

For further research, our framework clearly sets us up to construct a passive factor-based risk parity portfolio, where allocations to the major risk factors are maintained mechanically in parity. We argue that this approach would serve as the more sensible passive risk parity portfolio from which actively managed risk parity portfolios can be benchmarked.

ENDNOTES

¹The universe dependency problem for naïve risk parity portfolios is carefully documented and discussed in Chaves et al. [2011].

²We reference these two providers because they were the only firms that readily provide verbiage on their risk parity philosophy on their website.

³See Maillard, Thierry, and Teiletche [2010] and Chaves et al. [2012b] for detailed discussions on the mathematical definition of equal risk contribution risk parity portfolios as well as the computation algorithms.

⁴We borrowed this example of mapping assets to foods and factors to nutrients from Andrew Ang’s 2012 Research Affiliates Advisory Panel presentation.

⁵PCA, originally developed for reducing the dimensionality of the data by Pearson [1901], is now a common technique used in finance for extracting orthogonalized risk factors. For a given universe of assets, PCA identifies the independent factors (the so-called principal components), which explain the majority of the covariance in the time series of returns. In the APT framework, these PCA factors would correspond to the systematic risks factors that drive the co-movements in returns. See Connor and Korajczyk [1993] for more details.

⁶To stretch the nutrient analogy used previously, note that while hundreds of macro- and micronutrients make up foods, for most people it is generally sufficient to focus on protein, fat, and carbohydrates.

⁷It does also bring into question the potential look-forward bias of RP #1’s simulated history.

⁸Given expected excess return and covariance matrix assumptions, a mean-variance optimal portfolio can be computed as: $w^* = m \Sigma^{-1} \mu$, where μ is a column of expected excess returns and Σ is the covariance matrix. Thus, assuming a given portfolio is optimal, MVO-implied expected returns are given by $\mu^* = \frac{1}{m} \Sigma w$, where m is the normalization factor to get portfolio weights to sum to 1.

⁹The 60/40 is generally meant to target an above 8% return or 10% portfolio volatility without using leverage. It is now understood that the leverage constraint comes at the expense of extreme portfolio risk concentration.

¹⁰The large implied equity premium comes from the fact that variance of equities is about 19 times that of bonds. The 60/40 allocation means that the equity expected return can be estimated by $1.5 \times 19 \times 2.8 \approx 80\%$, where $1.5 = 60/40$; $19 = \text{equity variance/bond variance}$ [note that the variance ratio translates into a volatility ratio of about 4]; 2.8 is the assumed rate of return for bonds. The negative correlation between equities and bonds then amplifies the effect, as the willingness to forgo the hedge benefit of bonds must further imply added equity attractiveness.

¹¹Chaves et al. [2012] show that the optimal asset-based risk parity portfolio, which properly accounts for the correlation information between asset classes, does not produce significantly different portfolios from the naïve risk parity approach. This suggests that asset-based risk parity approaches, generally, are exposed to unintended concentration risk driven by the universe dependency.

REFERENCES

Bhansali, V. "Beyond Risk Parity." *The Journal of Investing*, Vol. 20, No. 1 (2011), pp. 137-147.

Black, F., and R. Litterman. "Global Portfolio Optimization." *Financial Analysts Journal* (September 1992), pp. 28-43.

Chaves, D., et al. "Risk Parity Portfolio vs. Other Asset Allocation Heuristic Portfolios." *The Journal of Investing*, Vol. 20, No. 1 (2011), pp. 108-118.

———. "A PCA Approach to Risk Factor Parity." Research Affiliates White paper, 2012a.

———. "Efficient Algorithms for Computing Risk Parity Portfolio Weights." *The Journal of Investing*, Vol. 21, No. 2 (2012b).

Connor, G., and R. Korajczyk. "A Test for the Number of Factors in an Approximate Factor Model." *Journal of Finance*, Vol. 48, No. 4 (1993), pp. 1263-1291.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.