MTree Capital PARTNERS

INSIGHT SERIES

in cooperation with AMAC RESEARCH

BEHIND THE VEIL: RISK PARITY ASSET ALLOCATION

Abstract

Risk considerations are key to any type of portfolio management processes. Whether it is in a multi-asset portfolio or a portfolio linked to a specific asset class, dollar weighted allocations or allocations purely based on expected returns are not considered optimal. This article carries out a review of risk parity asset allocation, including the role of risk and the logical inconsistencies it entails. In order to place risk parity in the context of modern portfolio theory, the author reverse engineers the risk parity portfolio back through the mean-variance optimiser to identify the implied assumptions an investor makes when deciding to adopt risk parity allocations. Once put in the context of modern portfolio theory, the paper investigates the trade-offs risk parity investors implicitly make.

1. Introduction

Risk considerations are key to any type of portfolio management processes. Whether it is in a multi-asset portfolio or a portfolio linked to a specific asset class, dollar weighted allocations or allocations purely based on expected returns are not considered optimal. It is well understood that expected risk needs to play a role in the ultimate decision. Risk-based thinking in asset allocation is everywhere in modern portfolio theory, including the Black-Litterman approach, alpha-beta separation, alternative beta-management and even liability driven investments (LDI). All of these investment management techniques are valuable components of our industry and can play a role in many portfolios. There is also no doubt that all of these techniques are part of a positive evolution of the investment industry. However, it is much less clear whether it is a good idea to base an asset allocation methodology purely on risk expectations.

Risk Parity as an asset allocation approach and as a group of multi-asset fund products have had a meteoric rise over the past ten years. Not only has it significantly outperformed the traditional 60/40 portfolio, it also provides a welcomed relief for investment committees during an unprecedented period in the global economy with unconventional central bank interventions abound. Its attractions are clear: it is easy to explain; does not require any conviction on individual assets; quantitative in construction and even has the all important R-word in the name.

These risk parity fund products appear in several guises, ranging from basic constructions widely offered by asset managers to more sophisticated methodologies which for example, aim to allocate equally to different economic scenarios in risk terms.

With this fashionable accessory in their briefcases, investment managers are heard proudly pronouncing their fondness for risk parity in boardrooms and conferences around the world:

"We do not try to forecast returns, we manage risk!";

"It's ALL about risk management";

"We allocate risk, not alpha".

Risk parity's relative outperformance in recent years and the intuitive nature of its philosophy has attracted many fans. However, it is unclear whether the implications and possible side effects of risk parity asset allocation are fully appreciated. This article aims to first explore the strengths and weaknesses of the traditional risk parity approach, with particular attention given to the context and measurement of risk. This is followed by an exercise to reverse engineer the risk parity approach back through the modern portfolio machine. This reverse engineering exercise effectively asks the following question: "What inputs and parameters would lead the mean-variance optimiser to produce a risk parity portfolio as the efficient portfolio?". The answers to this question provide the specific assumptions an investment manager implicitly makes regarding information ratios and correlations when deciding to adopt a risk parity approach. Once these assumptions are brought into light, we can further appraise the risk parity approach within the context of modern portfolio theory and investigate its vulnerabilities.

2. The good, the bad and the statistically insignificant

For many, intuitiveness is perhaps the number one reason for adopting the risk parity approach. Its mechanics can be sketched out on the back of an envelope and the quantitative construction and passive rebalancing give a sense of consistency and robustness.

Risk parity can also be thought of as a natural progression from notional (dollar-based) equal weighting. In this context, it is a successful evolution towards maximum diversification and therefore rightly deserves the attention from the investment community. Furthermore, this wider diversification, relative to 60/40, can provide smoother long run returns over multiple economic cycles.

On the other hand, risk parity presents some major problems and challenges. In most of its manifestations, risk parity funds employ significant leverage, particularly in the fixed income portion of the portfolio. This extensive and concentrated use of leveraged beta, in order to close funding gap or to match future uncertain spending, may prove to be one step too far. This problem was elegantly stated by Rodney Sullivan, editor of the Financial Analyst Journal, in a 2010 article, "Beating the market requires a strategy informed by a special insight that might add an alpha to returns. Leverage is not such an insight". In addition, despite reduction in transaction costs and more advanced derivative management in recent years, there is still significant friction to leverage, therefore further eroding effectiveness of risk parity portfolios (Anderson, Bianchi and Goldberg, The Decision to Lever, 2013).

Beyond the practical and frictional issues of leverage, the rationale for using any leverage at all presents a ***logical inconsistency***. This decision to lever is based on the expectation that an unlevered version of risk parity would unlikely be able to provide the desired level of returns due to the large fixed income component. How ironic is it then, to see return expectations being used to justify and determine leverage in a strategy that is supposed to be indifferent or agnostic to return expectations in the first place. As it turns out, this is not the only logical inconsistency with risk parity application. The devotees of risk parity are not in the business of forecasting returns… or are they?

Furthermore, central to risk parity is the indifference to return expectations and valuation. This might be fine in theoretical isolation but how would this fit in with the wider fiduciary responsibility performed by trustees and investment committees? Could these trustees justify such a completely reactive and mechanical process that ignores real world fundamental or valuation changes?

Risk parity, in its many guises, has outperformed the traditional 60/40 portfolio mix in recent years. A leveraged exposure in fixed income during the period of the "great moderation" turbo boosted by unprecedented central bank bond buying made sure the wind is fully behind the risk parity sail. The key question for investors, however, is whether this period of outperformance is a validation of this return-indifferent asset allocation methodology, or simply a retrospectively opportune period to be leveraged long fixed income assets.

This question certainly cannot be answered definitively but R. Anderson et al identified a number of aspects of risk parity research that might be less robust than suggested by the recent outperformance (Anderson, Bianchi and Goldberg, Will My Risk Parity Strategy Outperform?, 2012). Over the longer term research, risk parity looks a lot less attractive, with its outperformance easily eroded by, for example, start/end date choices, transaction costs and bond index choices. Recent outperformance is indeed very different to long term robustness.

Beyond the strengths and weaknesses mentioned above, central to the risk parity framework is the use of "risk" as the unit of measurement. Despite its allure as a successful marketing slogan, ("Risk is central to our asset allocation process!"), this isolated use of risk can be a practical and philosophical minefield.

3. The meaning and measurement of risk

Given risk parity portfolios are indifferent to expected returns, the spotlight must surely be front and centre on the definition and measurement of "risk". The exercise of identifying and evaluating every measure of risk used by risk parity practitioners is far beyond the scope of this article. However, in practice, it can be generalised that some measurement of historic volatility is by far the most common definition of risk. These measurements can be calculated in many ways and even enhanced or conditioned by some other measures such as options-implied volatilities.

Ultimately though, these mostly backward-looking risk measures suffer from issues in terms of accuracy and practical application as covered by many academics and practitioners. A more dynamic measurement of volatility (e.g. shorter half-life or measurement period) may give more accurate short-term forecasts of volatility but would also be likely to cause higher turnover and hence transaction costs during rebalancing. Another issue with volatility is that it is not equally "transparent" in every asset class. In such cases, the volatility of some real assets or private equity investments cannot be compared "apple-to-apple" with traditional liquid assets that have more transparent mark-to-market information.

Another potential problem (or at least a philosophical challenge) for value investors is that the risk parity methodology tends to overweight an asset with rich valuation after it had a low volatility up-run in the recent past (e.g. US government bonds in 2013).

Volatility as a measure of risk has its limitations and many detractors but, interestingly, very few competitors. As elegantly put by Cliff Asness (and fully agreed by the author) in a 2014 FAJ article "My Top 10 Peeves", when properly used in conjunction with expected returns, volatility is an important tool in the modern quantitative finance framework. In this context, volatility (as a measurement of risk) is the uncertainty around return forecasts. This is central to modern portfolio theory but unfortunately, this is also the place where we discover risk parity ***logical inconsistency no. 2***. If "risk" is indeed the uncertainty around return forecasts, why is it front and centre in a strategy that is

indifferent to expected returns in the first place? In other words, why are risk estimates being generated as uncertainty levels around return forecasts that were intentionally NOT made?

4. Back through the grinder

Beyond looking at its practical pros and cons and logical (in)consistencies, it is important to examine this asset allocation methodology within the context of modern portfolio theory. This article will take a two-step approach, which starts by asking "under the mean variance optimisation framework, what specific inputs of expected return, expected variance and correlation would lead the optimiser to produce the risk parity portfolio as the optimal portfolio?". This reverse engineering is crucial in the understanding of risk parity and effectively provides the technical subtitle to the risk parity movie.

Under the mean-variance framework, three sets of inputs are required: expected returns, expected variances and expected correlations. The reverse engineering exercise states that for risk party to be optimal, both of the following two specific conditions would need to be met:

1) All assets have the same information ratios (expected returns divided by expected standard deviation);

2) Correlations between all assets are the same;

This can be demonstrated with the following set of hypothetical assets:

Asset A - ExpRet = 5%, ExpSD = 5%

Asset B - ExpRet = 10%, ExpSD = 10%

Asset C - ExpRet = 15%, ExpSD = 15%

Asset D - ExpRet = 20%, ExpSD = 20%

Correlations between all assets are 0.5.

The mean-variance optimiser (with long-only constraints) produces the following set of results in a 100-portfolio optimisation:

	Portfolio Weights						Information
	Asset A	Asset B	Asset C	Asset D	Risk	Return	Ratio
Portfolio 23	60%	20%	11%	8%	6.6%	8.3%	**1.2556**
Portfolio 24	59%	21%	12%	8%	6.7%	8.5%	**1.2580**
Portfolio 25	57%	21%	13%	9%	6.9%	8.6%	**1.2599**
Portfolio 26	56%	22%	13%	9%	7.0%	8.8%	**1.2615**
Portfolio 27	54%	22%	14%	10%	7.1%	8.9%	**1.2627**
Portfolio 28	53%	23%	14%	10%	7.2%	9.1%	**1.2636**
Portfolio 29	51%	23%	15%	11%	7.3%	9.2%	**1.2643**
Portfolio 30	50%	23%	15%	11%	7.4%	9.4%	**1.2647**
Portfolio 31	49%	24%	16%	12%	7.5%	9.5%	**1.2649**
Portfolio 32	47%	24%	16%	12%	7.7%	9.7%	**1.2649**
Portfolio 33	46%	25%	17%	13%	7.8%	9.8%	**1.2647**
Portfolio 34	44%	25%	17%	13%	7.9%	10.0%	**1.2643**
Portfolio 35	43%	26%	18%	14%	8.0%	10.2%	**1.2637**
Portfolio 36	41%	26%	19%	14%	8.2%	10.3%	**1.2630**
Portfolio 37	40%	26%	19%	15%	8.3%	10.5%	**1.2622**
Portfolio 38	38%	27%	20%	15%	8.4%	10.6%	**1.2613**
Portfolio 39	37%	27%	20%	16%	8.5%	10.8%	**1.2602**

As demonstrated above, the optimal portfolio is the one where the portfolio weights are in the same inverse proportions to expected standard deviation or volatility, i.e. the unlevered risk parity portfolio. We can further demonstrate that as long as the correlations are the same, it does not matter what the exact correlation value is.

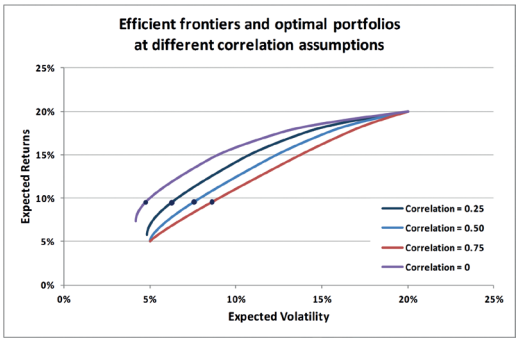

Similarly, as long as the information ratios are the same, it does not matter what the exact return and volatility numbers are or how widely diverged they are.

So now we know, when an investor chooses to adopt a risk parity asset allocation approach, he or she is implicitly describing this specific set of expectations around returns and correlations, whether he/she is aware of them or not.

5. Behind the frontier

Now that we have identified the specific conditions where the use of risk parity would be "optimal", we can move to the next part of understanding the role of risk parity in real world applications.

One by-product of the earlier reverse-engineering exercise is the inference that the risk parity portfolio cannot be optimal as those specific assumptions cannot be met in reality. This allows us to put the sub-optimal risk parity portfolio behind the efficient frontier in the conceptual framework and more clearly understand where the illusion of effectiveness might come from.

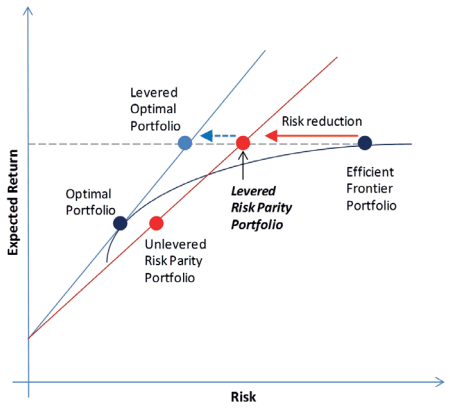

As the diagram above demonstrates, a levered risk parity portfolio can indeed achieve a certain level of expected return with a lower level of expected volatility than an unlevered efficient frontier portfolio. However, this is simply an illusion, as a levered optimal portfolio can achieve an even better return-to-risk ratio.

Some risk parity practitioners argue that the Capital Asset Pricing Model (CAPM) does not reflect reality due to leverage aversion (Asness, 2012; Black, 1972; Frazzini and Pedersen, 2010; similar to arguments used to promote minimum volatility strategies) and effectively the real world risk parity portfolio can sit beyond the efficient frontier rather than behind it.

Whether CAPM reflects reality or not is a Nobel Prize-worthy debate and is out of the scope of this article. However, by proposing such an argument, the risk parity practitioners have exposed **logical inconsistency no. 3**. This leverage aversion argument centres on the assertion that low volatility assets have higher return-to-risk ratios and therefore a methodology that overweights bonds should outperform the market portfolio. So in the end, the most sophisticated argument to defend risk parity is one that admits expected returns are produced somehow. And surely, if this was the case, one can simply enter a higher expected return assumption into the mean-variance optimiser for the fixed income assets. Choosing to allocate through risk parity is simply an arbitrary (though admittedly marketing-friendly) way to justify such a leveraged fixed income allocation. In other words, if one thinks bonds have a higher information ratio than other asset classes, that in itself is the reason to overweight (or even to leverage), but not through a methodology that appears to be indifferent to expected returns.

6. Challenges

As mentioned in section 3, risk is often identified as some measure of historic volatility in risk parity methodologies. However, not all assets have the same transparency of volatility. Private equity and some real assets can have very low volatility in the context of mark-to-market relative to more liquid assets. This could lead to underestimation of the volatility of such asset classes and in turn leading to very high weightings in a traditional risk parity approach that are unintuitive and lead to liquidity compromises.

In addition, as mentioned in earlier parts of the article, risk parity is largely a volatility parity methodology in practice. It is important to recognise that risk is a much more complex idea than just volatility. It should include considerations relating to the shape of the whole expected return distribution, which often can have non-normal skew and kurtosis as well as other asymmetrical characteristics.

7. Conclusion

No one can argue with the success of risk parity portfolios in recent years when compared to the, surely by now, mythical and defenceless 60/40 portfolio mix and it certainly has many practical advantages and elements of this approach rightly deserve a place in the evolution of our industry. It could well continue to perform and reward its practitioners for years to come but investors should not be complacent about its true identity and vulnerabilities.

Perhaps the most intriguing aspect of risk parity are its logical inconsistencies - it proudly puts indifference to expected returns in the centre of its construction, yet by using “risk” as the key parameter, the users have somehow managed to put very specific levels of uncertainty around return forecasts that they intentionally did not make. Or maybe they did make them after all, otherwise there would be no justification to lever in the first place. Lastly, the evidence of higher return-to-risk offered by low volatility assets cannot in itself be an argument for risk parity. It could certainly argue for a higher fixed income allocation and leverage, but the link to risk parity is arbitrary and ignores much of modern portfolio theories.

Taking the logical inconsistencies, theoretical vulnerabilities and both practical advantages and disadvantages together, we can start to see behind the veil of risk parity. Stripped of its marketing glitter and put backward through the modern portfolio machine, risk parity is starting to look a lot like something that gives diversification in one hand but takes away efficiency in the other, therefore cancelling out the diversification benefits with the need to lever and an arbitrary allocation methodology.

Generating outperformance require insights or good timing. It seems risk parity had good timing in abundance over recent years but not necessarily the insight.

8. Reference

Rodney N. Sullivan, 2010. “Speculative Leverage: A False Cure for Pension Woes”. Financial Analysts Journal, vol. 66.

Robert M. Anderson, Stephen W. Bianchi, Lisa R. Goldberg, 2013. “The Decision to Lever”. University of California, Berkeley, Working Paper Series.

Robert M. Anderson, Stephen W. Bianchi, Lisa R. Goldberg, 2012. “Will My Risk Parity Strategy Outperform?”. University of California, Berkeley, Working Paper Series.

Clifford S. Asness, 2014. “My Top 10 Peeves”. Financial Analysts Journal, vol. 70.

Clifford S. Asness, Andrea Frazzini, Lasse H. Pedersen, 2012. “Leverage Aversion and Risk Parity”. Financial Analysts Journal, vol. 68.

Fischer Black, 1972. “Capital Market Equilibrium with Restricted Borrowing.” Journal of Business, vol. 45, no. 3 (July):444–455.

Andrea Frazzini, Lasse H. Pedersen, 2010. “Betting Against Beta.” NBER Working Paper 16601 (December)

Owen Thiers *is the founder of MTree Capital Partners, a vision based service provider to the investment management industry. It offers bespoke institutional distribution services as well as management consulting services for the (alternative) investment management industry. Owen has been 17 years in the industry, has worked for blue chip investment firms like Deutsche Asset Management and Citigroup Alternative Investments, and held a variety of senior roles related towards the institutional investment community.*

Aston S. K. Chan *is the founder of AMAC Research, a research consultancy specialising in providing tactical asset allocation and macro research to asset managers and family offices. Aston holds a Masters in Finance degree from the London Business School and previously worked in the field of global macro and tactical asset allocation in Deustche Asset Management, Auriel Capital Management LLP and was co-founder and portfolio manager of the GLC Global Macro Fund.*