ALTERNATIVE INVESTING

—

Risk Parity Is Even Better Than We Thought

June 1, 2015

No, I don't mean tactically. I like to think I'm bold, sometimes a bit too bold, but not crazy! We think both stocks and bonds are both quite expensive versus history and that this typically, though not always, leads to lower than normal long-term returns.[1] But valuation is a poor timing method. We are not writing articles, belying my title, about how risk parity[2] or for that matter traditional 60/40 investing is tactically attractive now.[3] The risk parity versus 60/40 argument has always been about strategic long-term — not tactical short-term — asset allocation (tactical arguments are fine, they just are not the point here; if you can time the stock and bond markets, you could do so whether using a traditional or risk parity type strategic allocation as a baseline). Here I argue that, when viewed strategically, the empirical work on risk parity, including some of our own, understates its potential advantages. Moreover, all you need is basic finance theory to see it; and, as we'll see, you have to have rather pessimistic assumptions about risk parity to not add any of it to an existing 60/40 portfolio.

There has been a lot of dispute about the historically achievable Sharpe ratio of risk parity versus 60/40. Issues abound. What period is a fair one to judge the key assets (stocks, bonds, commodities, inflation protected bonds, etc.)? How much should we subtract from a backtest to account for costs, both trading and leverage? We have taken, to nobody's shock, a pro-risk parity position. To see some of the back and forth, check out articles here, here, here, and here. In fact, we have written a new rejoinder, which will come out soon in the *Financial Analysts Journal* (here's a preview of our comments: the paper we're responding to flat-out mistakes a difference in volatility due to risk targeting for transactions costs and the cost of leverage).

One thing these challenges and defenses generally have in common is that they focus on the direct comparison of a portfolio like traditional 60/40 to risk parity.[4] They ask, which has the better Sharpe ratio after considering all real life implementation issues? This is the right way to compare if one is considering between two alternatives: 1) staying completely traditional, in this case 100% invested in regular old 60/40, or 2) moving 100% to risk parity. But this is trying to answer a question few are asking; very few will actually consider a 100% move to risk parity. Rather, investors interested in risk parity are mostly interested in making only a partial allocation to it. And, if one is considering a partial allocation, comparing Sharpe ratios, as both we and the risk parity critics have done too often, is not the right approach at all.

Think about how you decide to allocate a small part of your portfolio to something new. Do you say, "Does this have a larger net Sharpe ratio than my entire portfolio?" No, you don't. Whether examined formally through assumptions and quantitative analysis, or informally, the proper question is: "Does this allocation make my total portfolio better?" In that case, the assumed Sharpe ratio of the new investment is, of course, important, but so is its correlation to your existing portfolio. If the correlation is low enough, an allocation to an investment with a modest Sharpe ratio offers improvement over your starting portfolio.

To formalize this a bit, let 60/40 represent the traditional portfolio and RP represent the risk parity portfolio. If you had to choose between 100%/0% and 0%/100% it would come down to, after all costs, whether[5]:

$$SR_{RP} > SR_{60/40} \qquad (1)$$

If that were satisfied, a 100% allocation to risk parity would be superior to a 100% allocation to 60/40.

But, if you were willing to consider less extreme allocations, anywhere between 0% and 100% to risk parity, the question is different, it's whether or not $\alpha_{RP} > 0$ in the simple regression:

$$r_{RP} = \alpha_{RP} + \beta_{RP} * r_{60/40} + noise \qquad (2)$$

That is, if 60/40 is your starting portfolio, you'll want to allocate some amount away from it to risk parity if your best guess of α_{RP} is positive. That's basic portfolio math.[6] In English, it says that if risk parity is not perfectly correlated to 60/40, and 60/40 is your entire portfolio, the hurdle to allocate some money to risk parity is whether it adds value net of any common variation with your starting 60/40 portfolio, not whether risk parity completely dominates 60/40 head-to-head (in a comparison of 100% in risk parity to 100% in 60/40).[7]

I think we sometimes forget this Finance 101 idea when thinking about risk parity, often implicitly arguing about whether an investor should be *all* 60/40 or *all* risk parity. That's just wrong! We don't seem to make this mistake for most other investments. We don't require them to be better than our entire current portfolio — and then, if so, assume we move into them 100% leaving our current portfolio entirely. But for some reason, perhaps because the topic itself is strategic allocation, we seem to evaluate risk parity this way. But there's really no reason

why it should be thought of differently. If a partial allocation makes your portfolio better, a partial allocation should be made.

Let's consider a simple example. We often use a "simple" risk parity strategy in backtests.[8] It does not include everything we use in real implementations today — neither all the assets nor all the portfolio management and risk control. For one thing, we desire a long backtest, and many of the assets that we think worthwhile to include today don't have a long-enough history (e.g., inflation-protected bonds in the U.S.). For another, we want to keep it as free from data mining as possible.

Over 1947-2015 we find the gross Sharpe ratio of 60/40 to be 0.52 and of this simple risk parity backtest to be 0.75. We also find that both — 60/40 by outcome and risk parity by design — deliver long-term volatility of just about 10%. That means risk parity outperforms 60/40 by almost 2.5% a year (almost but not quite the Sharpe difference times a 10% volatility — in backtests, the volatilities come out close to but not exactly 10%). It also means that you could take off up to this amount per annum in costs and haircuts of all kinds (trading, leverage, a belief that future risk-adjusted returns will strongly favor equities) and still prefer risk parity to 60/40 if those two extremes were the only choices. But is 2.5% actually about the right hurdle?

Let's return to our regression based approach and estimate equation (2) empirically using the 1947-2015 monthly returns (all numbers annualized):

$$r_{RP} = 4.0\% + 0.70 * r_{60/40} \qquad (3)$$

This says first that for every 1% move in 60/40, risk parity tends to move 70 bps. This makes sense, as 60/40 is dominated by equities and risk parity has fewer equities and more of everything else.

Of course the 4.0% figure is our focus. The roughly 2.5% return difference we found earlier was the amount by which we could reduce risk parity returns and still prefer it to 60/40 in a binary "choose only one" scenario. In contrast, the 4.0% in this regression represents (and this is just portfolio math) how much we could knock off risk parity and still want to move *some of our portfolio* to risk parity. That is, unless you believe risk parity's performance will be more than 4.0% a year worse than the backtest you'd allocate some of your starting portfolio to it, in preference to 100% in 60/40. Please note, knocking 2.5% off was to take risk parity's Sharpe ratio down to equal to 60/40. The fact that we'd allocate some to it even after knocking off anywhere less than 4.0% per annum means we'd allocate some to it if it's Sharpe ratio were decently worse than 60/40. Such is the power of diversification.

Essentially, the amount of pessimism you have to have not to do *any* risk parity is nearly double what you need not to do *all* risk parity. You could sum up this essay with the obvious, yet often forgotten principle that doing *any* requires far less belief than doing *all*. This can be thought of in a few equivalent ways:

- You always get the biggest bang out of adding a diversifying asset from the first part of your allocation to it; it's not a linear process. Using our numbers, moving 10% out of 60/40 (with its Sharpe of 0.52) to risk parity (with its Sharpe of 0.75) creates a portfolio with a Sharpe of 0.56 (an 8% increase on the unrounded numbers). But say you've already moved 90% of your funds to risk parity, leaving 10% in 60/40. Well, your Sharpe is now higher (as the allocation to risk parity is now starting out much higher) coming in at 0.75. Moving the last 10% to go to all risk parity raises it all the way to, wait for it, 0.75. That is, with rounding to the second decimal it doesn't move (it actually gets about 0.2% better). This observation is not unique to this situation. It just hasn't been applied to this situation enough. The biggest bang for the buck comes from the allocation of the first buck.

- Those who argue that risk parity's Sharpe, after all their assumed deductions, isn't better than 60/40, have often found it to be the same (no advantage, though of course some believe it's actually worse). That tells you that even after all of their assumptions knocking risk parity down, it turns out neither risk parity nor 60/40 are optimal, but rather some combination of them probably represents the better portfolio. In fact, even under some severe assumptions that drive the risk parity Sharpe down all the way to a zero advantage over 60/40, you'd still want a combination of risk parity and 60/40. In other words, for those not doing any yet, you'd want to go out and get yourself some risk parity!

- For an economic interpretation, if 60/40 is closer than risk parity to the true portfolio of investor wealth, then risk parity is producing its returns while having a lower beta to the risk we actually care about (losing when the overall wealth portfolio is losing). In an economic sense it's "lower beta," and thus the hurdle to invest in it is also lower. If it's "lower beta" but produces the same return and Sharpe (under pessimistic assumptions), that's still a good deal, just not one warranting 100% of your money.

Backing this up with some analytics, we can run optimizations on our simple risk parity backtest and 60/40 that shed some more light. Using our full gross backtest and optimizing Sharpe while targeting 10% volatility, we effectively want 100% in risk parity (with some rounding). If you take 50 bps off of risk parity's backtest to account for costs and any other cynicism, you still want mostly risk parity (about 90% risk parity and 10% 60/40). Take 1% off — maybe for high cost estimates and fees — and you want 80% risk parity and 20% 60/40. Now let's say you go mad (or, the equivalent of mad, you listen to some of risk parity's ardent critics!) and assume the full historical advantage of risk parity is gone. That is, assume all else is the same but risk parity returns almost 2.5% less per annum versus our simple backtests. This, again, gets to the same Sharpe ratio as 60/40. In this case, an optimizer still wants about 50% in each alternative (not a shock given they are two equal Sharpe assets that aren't perfectly correlated).[9] It is not until, and this is again just portfolio math, you assume that real life risk parity is actually 4.0% per annum worse than the simple gross backtest, and, miraculously, 60/40 remains unchanged and costless, that you don't want any of your portfolio in risk parity. This is a gigantic amount of cynicism involving, among other things, very large and unrealistic cost assumptions.

Let's look at this other ways. First, let's look at the full gross risk parity backtest and compare its rolling 10-year Sharpe ratio to that of 60/40 (that is, compute the risk parity and 60/40 Sharpe ratios over every 10 year period and plot the difference).[10]
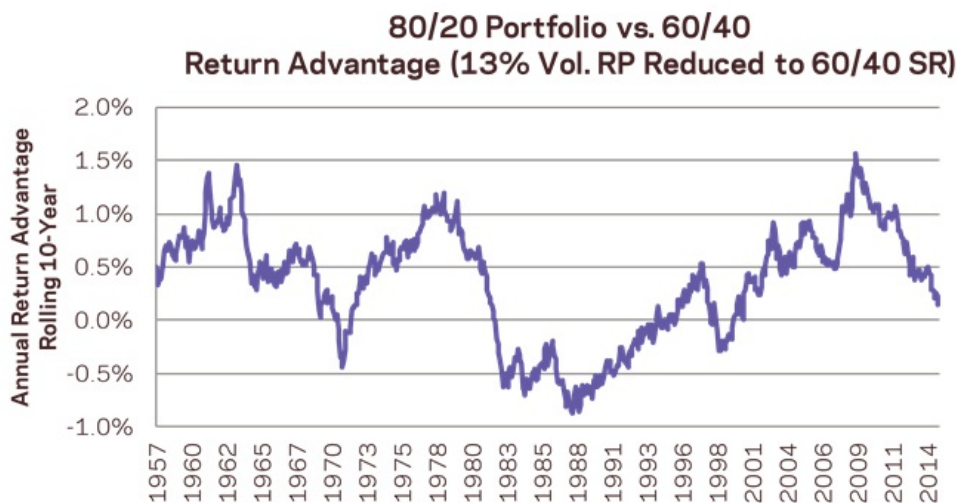
80/20 Portfolio vs. 60/40
Return Advantage (13% Vol. RP Reduced to 60/40 SR)

Risk parity wins 81% of the time.[11] But, again, that's a gross backtest.

Let's be draconian now and subtract from the risk parity returns the full difference in means between it and the 60/40 portfolio; by assertion declaring risk parity no better than 60/40. Specifically, we take out a constant amount each month from the risk parity backtest so that its mean return is the same as that of 60/40, but of course does not have the same time pattern of returns. If we look at a risk parity backtest reduced by this massive 2.5% per annum, and plot the rolling 10-year Sharpe of risk parity minus that of 60/40 you get this picture of randomness:
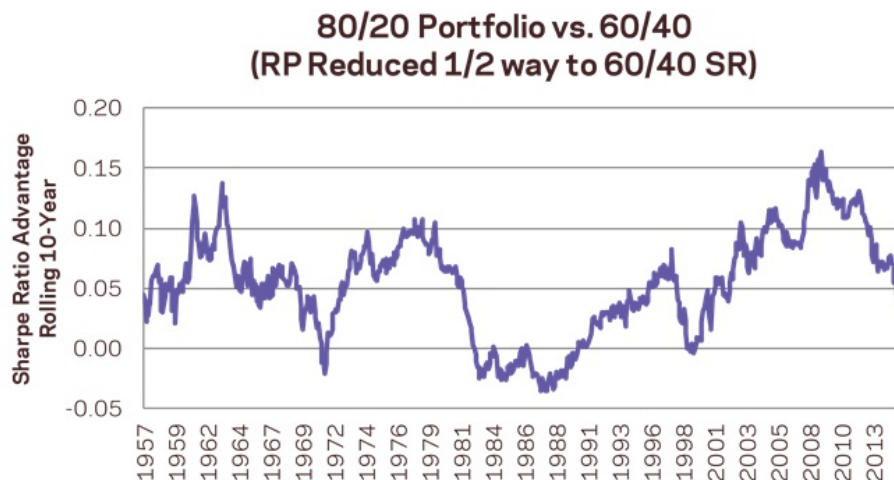
80/20 Portfolio vs. 60/40
(RP Reduced 1/2 way to 60/40 SR)

Of course, this is to be expected. If by construction you assert that two portfolios are the same realized Sharpe ratios, and neither is horribly skewed in either direction, they will each win about half the time (in fact 60/40 wins about 60% of the time above, but risk parity wins by more when it wins, making the means equal).

But, alas, I've temporarily forgotten my own point! I'm foolishly comparing 100% 60/40 to 100% risk parity in the above two graphs. What if, instead, we plotted the rolling 10-year Sharpe ratio of a portfolio 80% invested in 60/40 and 20% in risk parity minus the rolling 10-year Sharpe of one 100% in 60/40? That is, what is the gain from a move just one-fifth the way toward risk parity? Well, perhaps not surprisingly (given that improving the Sharpe of 60/40 with a modest allocation is an easier task than improving by fully replacing 60/40) an 80/20 portfolio of 60/40 and risk parity, without any haircut to the backtests, delivers a superior Sharpe ratio in 96% of the 10-year periods:

80/20 Portfolio vs. 60/40
(RP Reduced to 60/40 SR)

If we, again, cut risk parity's advantage in Sharpe to *zero*, removing about 250 bps a year, an 80/20 allocation still delivers a superior Sharpe to 60/40 in 75% of the rolling 10-year periods, just due to the power of diversification (a power that does not kick in when you fully replace one portfolio with another, but kicks in big when you allocate modest amounts like 20%):

80/20 Portfolio vs. 60/40
(Unreduced RP)

Now let's add a (perhaps most realistic?) middle-ground case to the two just considered, namely splitting the difference. Doing the same exercise with risk parity reduced half-way to 60/40 (taking off about 125 bps per annum of its approximately 250 bps per annum advantage) yields an improvement in Sharpe in 86% of the rolling 10-year periods:

Risk Parity vs. 60/40
(RP Reduced to 60/40 SR)

Source: AQR; the simulated Risk Parity strategy is based on a hypothetical portfolio described in greater detail at the conclusion of this document. The U.S. 60/40 portfolio consists of a 60% allocation to the S&P 500 index and a 40% allocation to U.S. 10 year Treasuries, rebalanced monthly. The 80/20 portfolio consists of an 80% allocation to the 60/40 portfolio and a 20% allocation to the 10% Risk Parity strategy. Please read important disclosures at the conclusion of this document.

Finally, I've framed this discussion in terms of Sharpe ratio. A superior Sharpe ratio can generally be used to increase return or reduce risk — though the more leverage-averse one is, and the more pessimistic one is about risk parity's Sharpe advantage, the more its use turns to risk reduction (or, with an extreme amount of pessimism, to no use at all). If one reduces risk parity's Sharpe in backtests to that of 60/40 (again an assumption we think draconian) and targets the same volatility in risk parity as 60/40, then allocations between them can reduce risk but not affect average returns.[12] However, if one is willing to invest 20% in a slightly more aggressive version of risk parity, targeting almost 13% instead of 10% volatility (10% is, again, the approximately volatility of 60/40), you gain a return advantage even while targeting a portfolio that in total still realizes only 10% volatility. I chose 13% volatility so that the volatility of a portfolio of 80% in 60/40 and 20% in this more aggressive 13% volatility risk parity has the same approximate 10% volatility as 60/40 alone (backtested worst drawdowns are also actually a bit less painful for this new portfolio than for 100% 60/40, so this is not just a consequence of using volatility to measure risk). Under these assumptions (again we think very pessimistic for long-term risk parity!) the 10-year average annualized value-add to the overall portfolio return is plotted below:

Risk Parity vs. 60/40
(Unreduced RP)

Source: AQR; the simulated Risk Parity strategy is based on a hypothetical portfolio described in greater detail at the conclusion of this document. The U.S. 60/40 portfolio consists of a 60% allocation to the S&P 500 index and a 40% allocation to U.S. 10 year Treasuries, rebalanced monthly. The 80/20 portfolio consists of an 80% allocation to the 60/40 portfolio and a 20% allocation to the 13% Risk Parity strategy. Please read important disclosures at the conclusion of this document.

This comes out to an average return advantage of 33 bps per annum on the whole portfolio with no increase in volatility or worst drawdown from only a 20% allocation to a risk parity strategy that we assumed has no edge over a 60/40 portfolio.[13] And, because we assumed no Sharpe advantage, all of this return is coming from monetizing the power of diversification.

Of course, if we assume a less severe, and we'd argue more reasonable, reduction of ½ way between the backtested risk parity historical results and 60/40 we find this 20% allocation to 13% volatility risk parity adds 64 basis points per annum to the whole portfolio without adding to portfolio volatility or making the worst drawdown more severe over 1947-2015.

Basically, again what we've done here is portfolio math 101, not rocket science. The hurdle in belief, assumptions, etc., for adding something as a modest allocation in a portfolio is, and should be, far lower than moving to it lock, stock and barrel (or, in this case, lock, bond and barrel — OK, I feel a bit of shame now for that one…). However, because risk parity is often framed as head-to-head versus 60/40 or the like, we believe this important perspective has been missing from the debate.

I examined this question using backtests. That's one way to do it. But it's far from the only way. If you never looked at a backtest, but formed your views on risk parity from any other economic or portfolio construction principles, you would always find it far easier to justify moving a modest amount of your portfolio into it versus moving a very large amount into it.

More generally, while some might still want to make the case to move entirely into risk parity, a case we too could support with the pure numbers, we think a partial move is far more realistic.[14] Investing is not about creating the absolute best portfolio regardless of realism. It's about creating the best portfolio you can actually stick with, and subjecting your own beliefs and evidence, especially when they differ from market-wide long-term consensus, to a reasonable (not draconian) standard of conservativism. No matter how much we might wish it were not the case, risk parity still uses things (i.e. leverage and simple derivatives) that, perhaps rationally, scare some investors, and risk parity still is a little unconventional (and is also still a strategy that by definition cannot be done by everyone as it's a tilt, not capitalization weights). Essentially it is still "tracking error" for most investors, and there's a real-world limit for almost all of us in what we can endure. Our advice to investors when asked how much to allocate is, for risk parity and in general, that it's rare you'll be able to approve or tolerate allocating what optimizers (even given realistic assumptions) say you should. We think you normally should allocate the amount you are very confident you can stick with through the inevitable bad times (we admit that coming up with this figure involves some art and some science, and we encourage conservatism).

It turns out that this approach— partial allocations, and often modest ones out of the whole portfolio — are also far easier to justify, and require far less optimism about risk parity (or permit far more cynicism about it). We think that when they view it this way, most investors will find it more difficult to justify not moving some assets to risk parity (strategically). Of course, that is opinion. You can still be cynical enough to move nothing to risk parity, but you have to be far more cynical to reject a modest partial move than to reject the full 100% switch to risk parity, a standard we argue many implicitly, and wrongly, use. Furthermore, again despite my provocative title, we make no tactical statement. If you think you can forecast crashes, more power to you. Our argument is about long-term strategic allocation. In our example, you would want some long-term allocation to risk parity even if it were anywhere better than 4% per annum worse than historical backtest relative to 60/40. In contrast, if your only choice was to move 100% from 60/40 to risk parity, implicitly asking the question "Which has the better net Sharpe ratio?" you'd reject risk parity if it were more than about 2.5% worse than the same history. A 1.5% difference in return hurdle is fairly large! We think this should have been obvious (including to us) but, for various reasons, we've all been wrongly comparing 100% 60/40 to 100% risk parity and not considering anything in between.

By the way, risk parity is a specific area where this logic is needed, but it applies everywhere. Whenever you find yourself asking "should I do A or B?" at least pause to ask "what if I do some of A and some of B?" Sometimes they will be too similar for this to matter, and sometimes splitting things up will bump into cost and logistical issues (like having too many managers to monitor effectively or giving up a size-based fee discount if you split your investment). But, sometimes, as I believe for risk parity and 60/40, the right answer really might be to do both.

One of the most basic lessons of investing is to think about how each investment impacts your overall portfolio not just its characteristics standalone. You don't evaluate your other investment options based on, "Would I trade my whole portfolio for this?" but rather on, "Does it make my portfolio better?" Risk parity should be evaluated likewise. We think it's a little better strategically, and it's diversifying, a powerful argument for adding some to a traditional portfolio.

—

The simulated Simple Risk Parity Strategy is based on a hypothetical portfolio of global equities, global bonds and commodities. Equities are based on a GDP-weighted global equity market portfolio until January 1970 using those global markets we have data for at each point in time; after that we use the MSCI World Equity index for the equity portion. Fixed income is an allocation to GDP-weighted global 10-year bonds, currency hedged. Commodities are an equal-weighted basket of whatever global markets we have data for at each point in time.

---

[1] Where we, perhaps, differ from consensus is stocks vs. bonds. We believe many look at stocks vs. bonds over too short a time frame. Some are way too short-term, looking at it over say the last 20 years, a period dominated by the very high stock valuations of the 1999-2000 technology bubble. Stocks will likely always look cheap vs. that! Looking at it over the last century both stock and bond valuations are high, but the spread between them is surprisingly close to average.

[2] A very quick introduction if you're not familiar with all of the jargon. A common model of a traditional institutional portfolio is allocating 60% of the assets to a diversified U.S. stock portfolio (S&P 500), and 40% to a portfolio of U.S. government bonds (10 year Treasuries). We call this 60/40. Our (and other's) objection to this is that by our estimates 85% or more of the risk comes from the stocks. The risk parity idea is that investors should diversify by risk instead of dollars. If stocks happen to be about three times as volatile as bonds, at the moment, a portfolio with 25% of the dollars in stocks and 75% of the dollars in bonds will have equal volatility in each bucket. Actual risk parity portfolios will usually expand to additional assets classes (such as commodities) and use leverage to get that low risk portfolio up to a risk level about the same as that of a 60/40 portfolio. Moreover, there are many very important implementation details. For one thing, few institutional investors use exactly 60/40. This essay is not concerned with those issues. It is comparing pure domestic 60/40 to a simplified version of risk parity that uses only global stocks, bonds and commodities. We hope, and believe, the implications are relevant to real traditional portfolios and real risk parity strategies. Some of our thoughts on risk parity itself are here and here.

[3] We stress again that our, and everyone's, ability to time markets is quite modest and it's very difficult to add value from timing even if you have some ability; we also point out that risk parity survives and can even thrive in rising-rate environments depending on the speed of the increase and the shape of the yield curve, and what happens to the equally important other asset markets. We certainly acknowledge that if you think you can forecast a short sharp severe bear market in bonds (or for that matter probably commodities – though we just kind of had one of these and risk parity came through OK) you also likely forecast short-term better returns for 60/40 than risk parity. Likewise in reverse if you forecast an equity crash.

[4] As an aside, we compare, as many others, to 60/40 here, but 60/40 is not a market capitalization-weighted portfolio, but rather an actively rebalanced one. Empirically, 60/40 seems to do somewhat better than the market capitalization-weighted portfolio, so effectively we're assuming a tougher competitor by choosing 60/40 from the start.

[5] Yes this assumes a relatively well behaved world where Sharpe ratio is the right measure. Mostly this is an issue for another day, but if investors are rationally leverage averse then risk parity's Sharpe ratio advantage may be fair compensation, not a free lunch. Also, again if they are leverage averse, depending on assumptions, risk parity may be more useful for reducing risk than for increasing expected return (if you are not very leverage averse it can easily be used for risk reduction, expected return enhancement or both — more on that later).

[6] For those who want to get more into the details, this would be a good place to dig in. The same author also offers some very good discussion of why "parity" is an oversold concept (indeed, precise equality has never been our point; we believe traditional portfolios have too much equity and too little of "everything else," but precise "parity" is only implied with very specific Sharpe and correlation assumptions).

[7] Of course, most starting portfolios are more complex than simple 60/40. However, usually a lot of that complexity diversifies away and I'd wager an analysis like this still gets you most of the way there.

[8] Some places we have used similar methods include here, here, here, and here.

[9] The backtested Sharpe ratio of this 50/50 combo is now 0.57 versus 0.52 for pure 60/40. A more modest improvement than for our less discounted numbers of course but still a 9% increase in Sharpe.

[10] It's important to note that this is the difference in Sharpe ratios between two portfolios, not the Sharpe ratio of the difference (long/short) portfolio (e.g., if risk parity's 10-year Sharpe ratio was 0.50 and 60/40's was 0.30, the number plotted in the graph for that month is 0.20).

[11] We admit that your eye is probably suggesting to you that since the line is heading down, in 10 years it is likely to be below zero. We don't put faith in this kind of eyeball analysis for financial data, and if you do careful statistical analysis, the level or slope of this graph is not a good way to "time" risk parity versus 60/40. However, one implication of this essay is that even if you are convinced risk parity will have a modestly lower Sharpe than 60/40 over the next 10 years, you likely still want some of it in your portfolio.

[12] Allocations can't affect arithmetic returns. They can mildly improve geometric returns through risk reduction.

[13] Please don't take numbers like "33 bps" too seriously. Running any numbers often leads to a false precision. I think the arguments here are directionally right but we don't know anything with precision like "33 bps"!

[14] We've actually been yelled at by some risk parity fellow travelers for our lack of adherence to dogma on this issue. We too would argue that risk parity is fairly close to the long-term optimal portfolio (among these asset classes) and the math warrants a near full allocation (admittedly "optimal" is always a flat surface guess). But we also know that the first small allocation helps way more than the last one. We believe that investors making very good allocations that they can hopefully stick with through thick and thin is far better than them making "perfect" allocations they are perhaps forced out of (for all the many reasons we all know can happen in bad periods for unconventional choices) at precisely the wrong time. If this be risk parity treason, make the most of it!