ABSTRACT

Several large plans, including the San Diego County Employees Retirement Association and the State of Wisconsin, have recently considered or decided to pursue risk parity strategies.[1] *Risk parity portfolios are examples of portfolios with higher Sharpe ratios than typical institutional portfolios. Like most other higher Sharpe ratio portfolios, risk parity portfolios are generally expected to return less than typical institutional portfolios. Therefore, they are typically leveraged to meet return requirements.*

We have three concerns with leveraged risk parity portfolios. First, the risk-adjusted advantage of these portfolios narrows and may disappear depending on the Fund's borrowing behavior (both costs and term structure). Second, using the Sharpe ratio as a comparative metric ignores left tail risk, which a leveraged portfolio could introduce and amplify. Third—and related to the second—global interest rates are relatively low and increasing a Fund's exposure to rising interest rates may not be advisable in such an environment.

OVERVIEW

Risk parity asset allocation strategies have recently been gaining popularity, and several large plans have already adopted them. As broadly defined, risk parity asset allocation strategies aim to spread portfolio risk equally across asset classes—or, in some cases, risk factors—resulting in portfolios that generally have greater expected risk-adjusted returns than a typical institutional portfolio.[2] However, though risk parity portfolios typically have greater risk-adjusted returns, they often do not have greater expected *absolute* returns.[3] This can be seen in Figure 1—the risk parity portfolio has a Sharpe ratio—a common measure of risk-adjusted returns—of 0.73, while the typical institutional portfolio has a Sharpe ratio of 0.58. At the same time, the risk parity portfolio has an expected return of 6.6%, versus the 7.5% expected return of the typical institutional portfolio.

[1] See "Plans Increasingly Consider Risk-Parity," Asset International; "Public Pension Funds Are Adding Risk to Raise Return," The New York Times.

[2] Whether the risk parity focuses on risk factors or asset classes, the general points of this paper remain valid.

[3] Without leverage; leverage in the context of risk parity strategies will be discussed later.

Figure 1. Capital and risk allocations of a typical institutional portfolio and a corresponding risk parity portfolio.

	Typical Institutional Portfolio	Risk Parity Portfolio
Capital Allocation	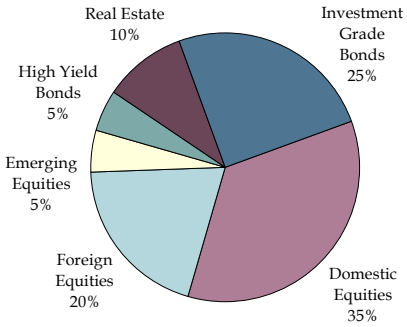	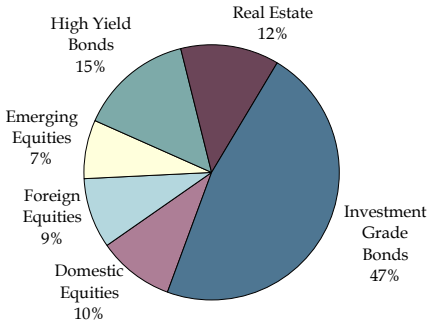
Risk Allocation	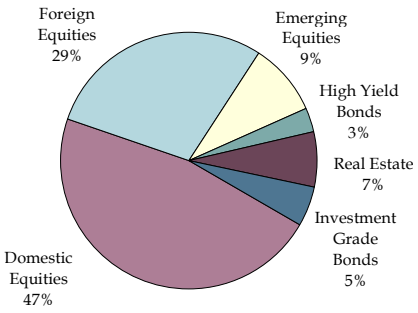	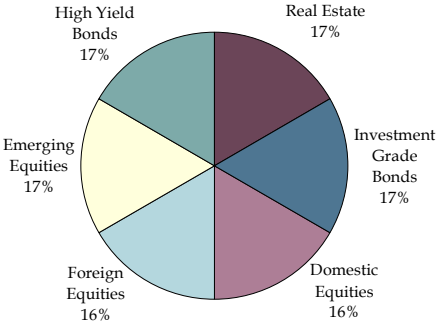
Exp. Return	7.5%	6.6%
Exp. Std. Dev.	12.0%	8.3%
Sharpe Ratio[4]	0.58	0.73

A higher Sharpe ratio implies that an investor in the risk parity portfolio will achieve his return with proportionally less risk *as defined by the standard deviation*. In this sense, higher Sharpe ratio (hereafter, "Sharper") portfolios are desirable[5]; unfortunately, Sharper portfolios are likely to have lower expected returns than a typical institutional portfolio. When viewed through this lens, there is nothing special about the risk parity portfolio except that it is an example of a portfolio that has a higher Sharpe ratio than the typical institutional portfolio. In fact, *any* portfolio that travels down the "efficient frontier" (or, at the very least, is above the line drawn from the risk-free rate to the typical institutional portfolio) will have a higher Sharpe ratio (see Figure 2).[6]

[4] The (expected) Sharpe ratio is equal to the (expected) return minus the risk-free rate, divided by the (expected) standard deviation. In this example, all assumptions are based on Meketa Investment Group's 2010 Annual Asset Study and 0.5% is used as the risk-free rate.

[5] A higher Sharpe ratio does not directly address the overall Fund's objectives; it is but one metric used to evaluate the desirability of an asset allocation.

[6] Some risk parity portfolios will lie below the efficient frontier but still have a higher Sharpe ratio.

Figure 2. Efficient frontier and three sample efficient portfolios: the typical institutional portfolio, the risk parity portfolio and the highest attainable Sharpe ratio portfolio.

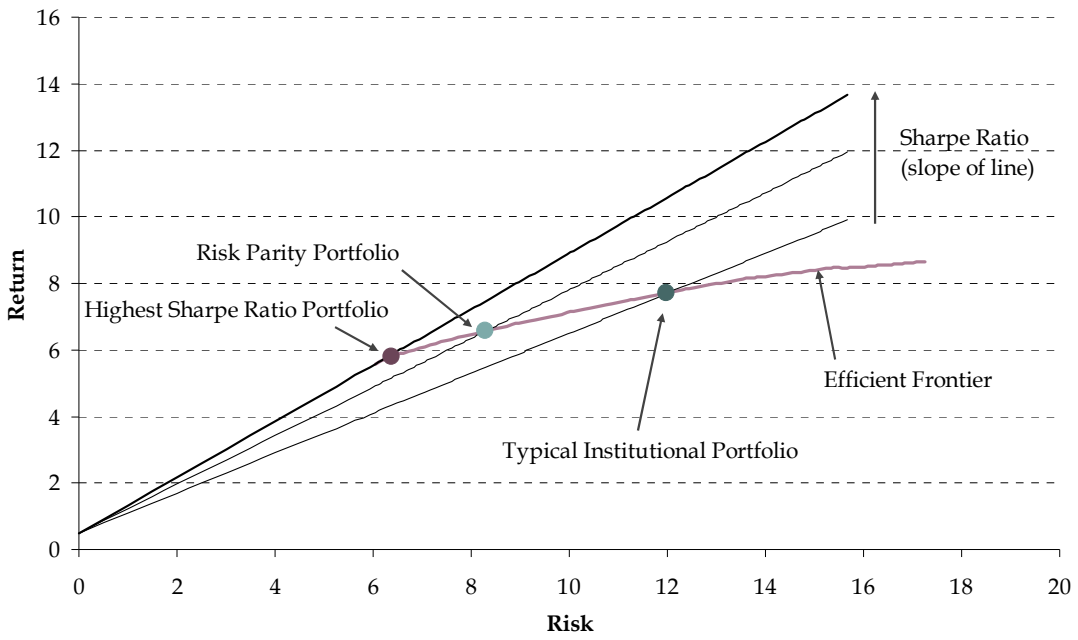

Sharper portfolios—including the risk parity portfolio—typically have lower expected returns compared to the typical institutional portfolio, as can be seen in Figure 2. In response, risk parity and other higher Sharpe ratio advocates suggest leveraging the Sharper portfolio such that the expected returns of the leveraged Sharper portfolio equal those of the typical institutional portfolio. Specifically, they recommend that an institutional investor either (a) borrow money through a plan-wide borrowing facility and invest the proceeds in the risk parity portfolio or (b) use derivatives to attain leveraged exposure to the higher Sharpe ratio portfolio. In the case of the risk parity portfolio shown above and assuming the borrowing or financing cost is roughly equal to the risk-free rate of 0.5%, in order to achieve the typical institutional portfolio return of 7.5% the risk parity investor would need to borrow or finance approximately 15% of their fund's value.[7] Under a strong assumption[8], the leveraged risk parity portfolio will achieve this 7.5% return with a standard deviation of (1.15 x 8.3% =) 9.6%, which is below the 12.0% standard deviation of the typical institutional portfolio. The leveraged risk parity portfolio's Sharpe ratio remains 0.73, higher than the institutional portfolio's 0.58. Graphically, this is represented in Figure 3.

[7] 1.15 x 6.6% - 0.15 x 0.5% = 7.5%

[8] This strong assumption is that borrowing costs can be treated as fixed. The validity of this assumption will be discussed in the next section.

Figure 3. Leveraged risk parity portfolio with the same expected return as the typical institutional portfolio. Borrowing at the risk-free rate pushes the unleveraged portfolio along the line with a slope equal to its Sharpe ratio, 0.73.

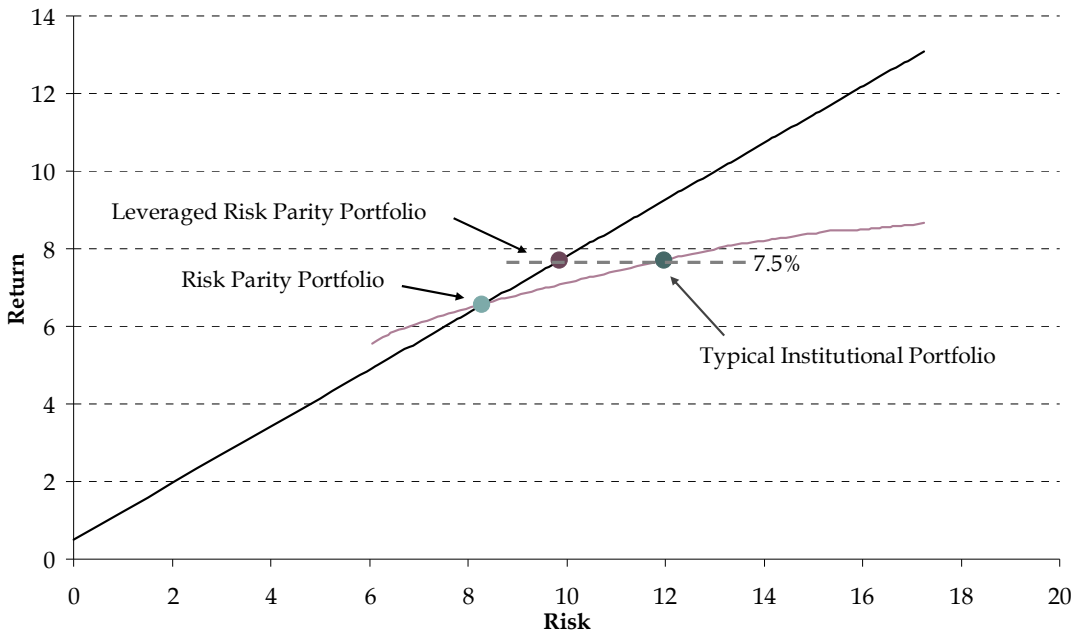

HISTORICAL PERFORMANCE

Figure 4 shows the relative historical performance of a leveraged risk parity portfolio compared with the typical institutional portfolio. Each portfolio is composed of domestic large cap stocks and investment grade bonds only; the risk parity portfolio is 20% stocks and 80% bonds, while the typical institutional portfolio is 60% stocks and 40% bonds. The risk parity portfolio's total leverage is determined on an annual basis as a function of both the difference in the expected returns of the two portfolios and the borrowing costs.[9]

[9] The expected return to stocks is assumed to be 5% plus the yield on intermediate term government bonds, while the expected return to investment grade bonds is assumed to be the current yield on a common bond index. The borrowing costs are assumed to be the intermediate term government bond yield divided by three. The precise choice of borrowing costs negligibly affects the results.

Figure 4. Historical performance of the risk parity portfolio.

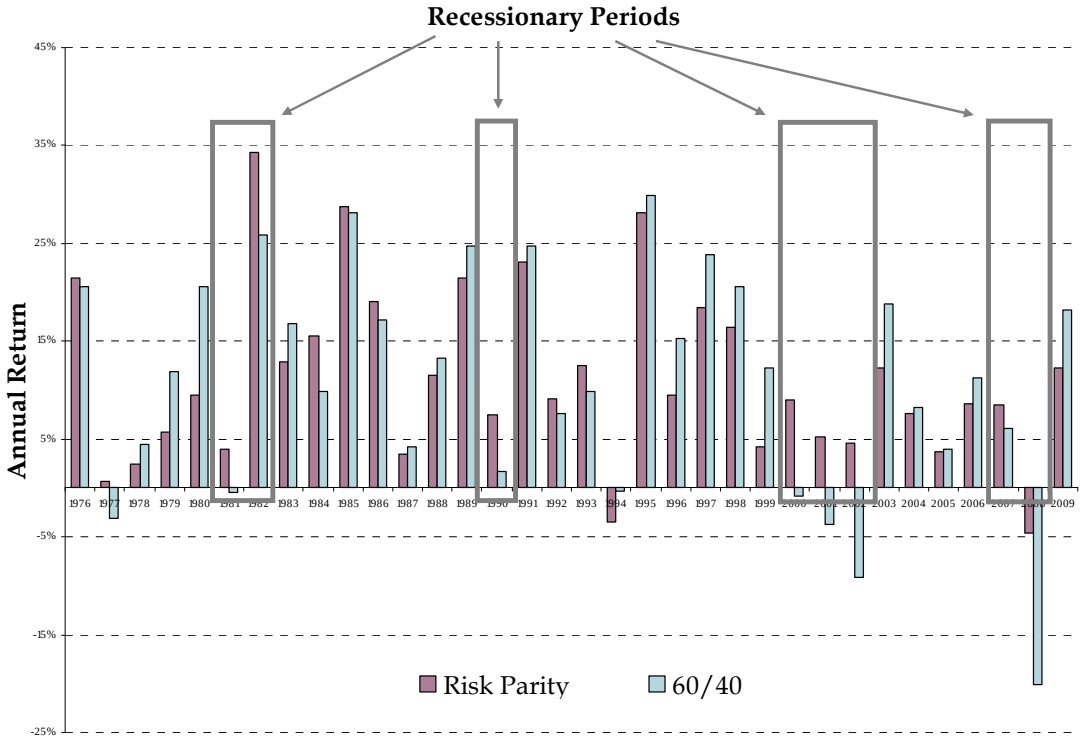

This is a highly simplified example (i.e., volatility assumptions are not updated on a yearly basis and there is no cap on leverage), but it proves a general point: risk parity and other higher Sharpe ratio portfolios generally perform better than typical institutional portfolios during times of economic weakness, while the opposite is generally true during periods of economic growth. Notable deviations of this pattern are the result of interest rate movements (e.g., the mid-1980s and the early 1990s featured high and declining interest rates, which translated into high bond returns that more than compensated for the underweight equities position in the risk parity portfolio).

In this simplified example, the leveraged risk parity portfolio outperformed the typical institutional portfolio by about 40 basis points per annum with significantly less volatility. Of course, it did so while underperforming for extended periods of time during boom times in the economy; *peer comparisons during these times may shake the faith of risk parity investors*. Most importantly, however, the example does not capture any tail risk due to the mechanics of leverage and, furthermore, featured an extended period of high and falling interest rates that may not be likely in the future.

CONCERNS

Three concerns plague leveraged risk parity and other leveraged Sharper portfolios: one is important but relatively trivial, one shakes the very foundations of the concept, and one is tactical.

Borrowing costs are higher than the risk-free rate and often variable.

First, borrowing money (either through a facility or derivatives) costs more than the risk-free rate. As borrowing costs increase, the Sharpe ratio advantage disappears. For example, assume that in the initial concept discussion, the borrowing cost of the Fund is 1.5%. Now, instead of 15%, the Fund must borrow 18% in order to achieve a 7.5% return. This leads to a Sharpe ratio that is now 0.71, down from 0.73 (but still higher than the typical institutional portfolio's 0.58).

Borrowing costs need to be moderately high before the Sharpe ratio advantage is eliminated. In the context of the earlier example, borrowing costs must rise to 4.6% before the Sharpe ratio advantage disappears. Historically, short-term borrowing *can and have* reached this point, most notably from 1996 through 2000. However, since short-term borrowing costs are currently well below this level, only longer-term borrowing facilities are cause for concern. Longer-term borrowing costs may very well approach 4.6%, eliminating the Sharpe ratio advantage of the risk parity portfolio. Indeed, yields on long-term (20+ year) investment grade municipal bonds are approximately 5.0%, leading to a structural preference for short-term borrowing to preserve the Sharpe ratio advantage of leveraged risk parity portfolios.

This leads directly to the next point. In the simple introductory example, borrowing costs are assumed to be fixed. This may be the case if an investor is able to set up a long-term facility and borrow fixed-rate funds. However, given the current cost of borrowing long-term funds, many longer-term leveraged Sharper portfolio investors will be borrowing on shorter time horizons than the return assumptions guiding their investments—effectively resulting in a duration mismatch and exposing the Fund to rollover risk, a risk that many adjustable-rate mortgage holders have experienced firsthand. In these cases, borrowing costs need to be treated as a *random variable*.

Treating the cost of borrowing as a random variable means that it will carry a standard deviation and will likely be correlated with any portfolio of risky assets. In fact, the correlation may be assumed to be negative: when times are good and markets are up, borrowing costs go down (for non-government borrowers); when times are bad and markets are down, borrowing costs go up (again, for non-government borrowers). A negative correlation between borrowing costs and a portfolio's return implies that borrowing *adds* risk to the portfolio above and beyond the contribution from its own standard deviation.[10]

The mathematical impact of variable borrowing costs is about as trivial as assuming higher borrowing costs: in the risk parity example, an assumption of 1.5% as the standard deviation for short-term borrowing costs and a -0.8 correlation with the risk parity portfolio yields a Sharpe ratio of 0.71, down from 0.73 (but still higher than the typical portfolio's 0.58). Even all together the impact of these two effects is relatively trivial: higher and variable borrowing costs as assumed above lower the Sharpe ratio to 0.70. Nevertheless, investors must consider

[10] As opposed to a negative correlation between two long positions (different assets within a portfolio), a negative correlation between a short position (borrowed funds) and a long position (the portfolio) adds risk to a total position.

these factors before deciding on the overall benefits of investing in a leveraged risk parity or other Sharper portfolio.

Sharpe ratios do not account for tail risk.

A significant drawback when relying on the Sharpe ratio as a comparative metric is its reliance on standard deviation as a measure of risk. Specifically, while a Sharper portfolio may have a smaller standard deviation *it may simultaneously carry more tail risk*. That is, the set of Sharper portfolio outcomes could be narrowly centered around the expected return (yielding a proportionally smaller standard deviation), but at the same time have a significant left tail (see Figure 5). Consequently, a high Sharpe ratio portfolio may be less desirable than a lower Sharpe ratio portfolio precisely because the lower Sharpe ratio portfolio has less left tail risk.

Figure 5. A distribution with a left tail.

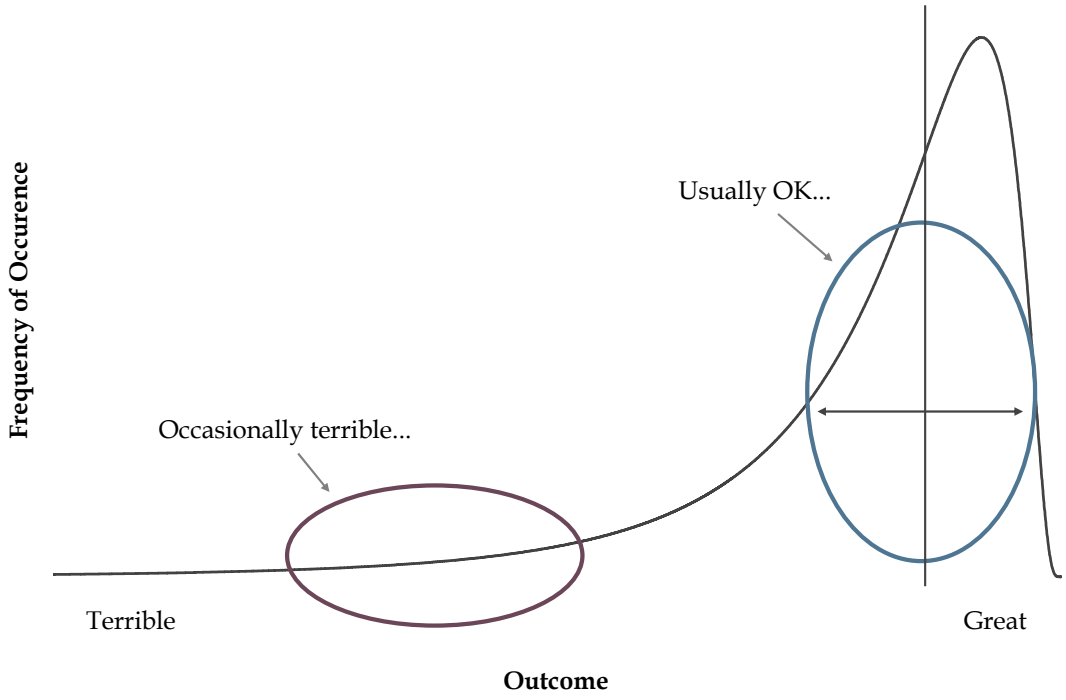

As with any portfolio, an unleveraged Sharper portfolio may or may not have more exposure to sources of a left tail risk than the typical institutional portfolio. The risk parity portfolio, however, *does*: risk parity portfolios are generally more heavily weighted to fixed income, which is an asset class whose upside is relatively limited but whose downside is not.[11] In addition, a *leveraged* Sharper portfolio has another identifiable source of increased left tail risk: the leverage itself. Leverage may seem innocuous most of the time. But as markets become exuberant and expected returns to risk assets go down, more and more market participants respond by increasing leverage. At some point, as in 2008, the market recognizes that the excessive borrowing cannot be paid back, and a financial crisis ensues.

[11] Of course, negative bond returns cannot exceed -100%.

Leveraged risk parity portfolios are particularly sensitive to a liquidity crisis, as they essentially behave like banks, borrowing short and lending long. During a liquidity crisis, investors who use leverage to fund their operations see a huge rise in borrowing costs—if borrowing is available at all. If there is a duration mismatch, the Fund may be forced to (a) sell assets in a panicked market in order to repay loans, or (b) rebalance to meet margin or collateral requirements. If assets are sold, the fund will not be invested in the leveraged Sharper portfolio as it was envisioned, perhaps consigning itself to a lower return than expected. In the case of a forced rebalancing, sponsors of the Fund need to be psychologically comfortable with committing large amounts of money in what will have been a seriously losing investment. While either of these events need not happen, they could—that is the nature of left tail risk: a low probability of a truly terrible outcome.

This is just one example of a possible tail event resulting from leverage. Another could be as simple as what happened in the historical performance analysis: borrowing costs rise to such an extent that the need for borrowing explodes, amplifying any relatively modest downturn in the risk parity portfolio.

Increasing longer-term interest rate risk in a low interest rate environment.

Almost all Sharper portfolios—including risk parity portfolios—increase the allocation to investment grade fixed income instruments (see Figure 1). Currently, longer-term interest rates are near record lows. While a moderate allocation to investment grade bonds can be justified as a hedge against a deflationary or slow growth environment, there is the medium-term risk that interest rates will rise because of (1) inadequate savings flows, (2) increasing inflation expectations, or (3) a strong economic rebound. In general, a Sharper portfolio with a high allocation to investment grade bonds will underperform a well structured institutional portfolio in any of those three scenarios: using the example in Figure 1, a 2% rise in longer-term interest rates would translate into approximately a 5% loss for the *unleveraged* risk parity portfolio versus only a 2.5% loss for the typical institutional portfolio. If this rise occurs in the context of a strong economic rebound, the risk parity portfolio's relative performance will be even worse because it is underweight risk assets relative to the typical institutional portfolio. From a tactical perspective, it does not appear advisable to overweight investment grade bonds unless the Fund adopts a firm belief in a medium-to-longer term deflationary or low-growth environment.

IMPLEMENTATION

Constructing a leveraged risk parity portfolio need not be difficult. Leveraged risk parity portfolios are typically set up through a plan-wide borrowing facility (for big, well known plans) or through derivatives. The former is a complicated endeavor and involves hiring investment banks to issue the debt, but can result in better duration management compared with the latter. Using derivatives is far less complicated, and is usually effected by hiring a leveraged bond manager. This bond manager likely purchases Treasury bond futures on margin and uses the remainder of his capital to ensure that the allocation tracks the specific index that the investor wants, such as the Barclays U.S. Aggregate. (Otherwise, the manager might engage in swaps, which introduces counterparty and tracking error risk.) Focusing

the leverage on one manager allows for a simpler process of portfolio monitoring, rebalancing, and policy changes, but also adds a possible additional layer of management fees.[12]

Furthermore, the Fund would need to carefully monitor the liquidity of their overall portfolio. As noted above, tail events could motivate a forced repayment or forced rebalancing operation. If during one of these forced events the Fund is caught with a significant portion of its assets in illiquid investments, then these repayment or rebalancing demands may go unfulfilled, resulting in severe disruption.

For plan sponsors who would prefer to avoid the details of implementation and monitoring, there are at least five active managers who offer "turnkey" risk parity solutions. These managers offer a range of risk parity solutions, including simple asset-based risk parity portfolios to more complicated risk factor parity (so-called "all weather") portfolios. Management fees start at 0.5% and can go much higher. Note that these turnkey portfolios do not invest in private equity partnerships, which complicates the asset allocation process for Funds with private investments (that is, there will be a slice of the strategic portfolio that is not accounted for in the risk parity portfolio).

CONCLUSION

A risk parity portfolio is a type of portfolio with a higher Sharpe ratio than typical institutional portfolios. Because they typically carry a higher Sharpe ratio, risk parity portfolios are expected to return more per unit of risk than a typical institutional portfolio. However, like most other higher Sharpe ratio portfolios, risk parity portfolios are generally expected to return less than typical institutional portfolios. As a result, risk parity portfolios are usually leveraged.

We have three concerns with leveraged risk parity portfolios. First, the risk-adjusted advantage of these portfolios narrows and may disappear depending on the Fund's borrowing behavior (both costs and term structure). Second, using the Sharpe ratio as a comparative metric ignores left tails risk, which a leveraged portfolio could introduce. Third global interest rates are relatively low and increasing a Fund's risk exposure to rising interest rates may not be advisable. For most plan sponsors who wish to adopt a leveraged higher Sharpe ratio portfolio, we recommend hiring a leveraged bond manager.

[12] In the vast majority of cases, the portfolio effect of leveraging the risk parity portfolio through either a plan-wide borrowing facility or derivatives (in one or more asset classes) is the same. However, there are limits to effecting portfolio-wide leverage through one allocation within that portfolio. For example, in the rare cases in which a risk parity strategy calls for leveraging the unleveraged portfolio two times, if the bond allocation represents only half of that portfolio, there is not enough capital to both maintain collateral for the bond futures program and to invest the required amount in the non-bond assets at the same time. Most Sharper portfolios call for around 1.3 to 1.6 times leverage.