Journal of Investment Strategies **6**(3), 13–30
DOI: 10.21314/JOIS.2017.084

Research Paper

Black–Litterman, exotic beta and varying efficient portfolios: an integrated approach

Ricky Alyn Cooper[1] and Marat Molyboga[2]

[1] Illinois Institute of Technology, 565 West Adams Avenue, Chicago, IL 60661, USA; email: rcooper3@iit.edu

[2] Efficient Capital Management, 4355 Weaver Parkway, Warrenville, IL 60555, USA; email: molyboga@efficient.com

(Received October 24, 2016; accepted January 16, 2017)

ABSTRACT

This paper brings Black–Litterman optimization, exotic betas and varying starting portfolios together into one complete, symbiotic framework. The approach is unique because these techniques are often viewed as alternatives rather than as complements to each other. We first demonstrate the approach using exotic beta as the "views" in the Black–Litterman optimization. This framework benefits investors who already utilize the classic Black–Litterman approach and appreciate advances in the exotic beta research, and also those who focus on practical implementation of exotic betas. We then explore the framework using the risk-parity portfolio as an efficient starting portfolio for Black–Litterman optimization on both theoretical and practical grounds. We demonstrate that risk parity is a highly effective starting point in many situations. Finally, as part of our discussion, we derive conditions under which almost any completely diversified portfolio may be used as a starting portfolio in the Black–Litterman process. The integrated methodology developed is robust, flexible and easily implemented, which means that a wide range of investors can benefit from this framework.

Keywords: Black–Litterman; exotic betas; risk parity; reverse optimization; alpha model; risk-adjusted returns.

Corresponding author: M. Molyboga

Print ISSN 2047-1238 | Online ISSN 2047-1246

1 INTRODUCTION

The Black–Litterman model has contributed to the field of quantitative portfolio management by elegantly applying Bayesian statistics to marry two seemingly contradictory ideas: the efficiency of the market portfolio and the efficacy of alpha modeling. However, the practical implementation of the model is often difficult because it relies on expert opinions, which often are hard to obtain and of uncertain quality, and the capitalization weights of the market portfolio, which are not always available. Further, the market portfolio is assumed to be nearly efficient. The latter assumption has often been questioned over the last two decades.

Exotic beta is a well-established concept emanating from a large body of research that presents a ready-made set of prior beliefs for people who trust the established literature more than a privately built model. This paper explores two aspects of Black–Litterman and exotic beta. First, we consider the likelihood that exotic beta will improve the performance of an index portfolio inside a Black–Litterman framework. Second, we consider using the Black–Litterman framework to derive an implementation portfolio for exotic beta. Our conclusion is that this combination is effective on both counts.

From a practical point of view we wish to address situations that may come up in the real world.

(1) A manager likes the Black–Litterman framework but does not have access to a good proprietary view. In this case, exotic beta may perform as the options.

(2) A manager likes the idea of the Black–Litterman framework but does not necessarily believe that a capitalization-weighted portfolio is efficient, or cannot get accurate capitalization weights. In this case, the risk-parity portfolio may perform as a starting portfolio.

(3) A manager likes exotic beta tilts to a portfolio and seeks a good implementation methodology. In this case, the Bayesian construction of Black–Litterman scales the exotic beta and provides returns that can be optimized.

(4) A portfolio manager wishes to use the Black–Litterman model, but does not wish, or is unable, to start with the market portfolio. In this case, we derive a more general reverse optimization that can be used with any well-diversified portfolio.

The net result of this research is that a wide range of investment professionals – including the portfolio manager interested in applying exotic betas, the fund-of-funds manager or commodity trading advisor interested in applying Black–Litterman and anyone interested in extending their tool set of allocation techniques – should find these results appealing.

The paper unfolds as follows. First, we demonstrate the manner in which exotic betas may be integrated with the Black–Litterman framework using a simple ten-stock example with an exotic beta of low volatility. Second, we illustrate the difference between starting Black–Litterman with a risk-parity portfolio and starting it with a capitalization-based portfolio using the same ten-stock example. Finally, we demonstrate how Black–Litterman, risk parity and exotic beta can be integrated within the Bayesian risk-parity framework using a three-asset-class example of stocks, bonds and commodities, which is particularly interesting because capitalization weights are not available. For this example, cross-sectional momentum is chosen as the exotic beta.

Along the way, we will also demonstrate an algorithm by which almost any very well-diversified portfolio may be used as the starting portfolio in the Black–Litterman framework, without sacrificing theoretical integrity.

2 STANDARD BLACK–LITTERMAN OPTIMIZATION

The original Black–Litterman model changed the landscape of quantitative portfolio management by combining into a single framework the two seemingly contradictory ideas of the efficiency of the market portfolio and alpha models. Black–Litterman optimization takes the implied returns from a cap-weighted index, which represents the market portfolio, combines them with a personal view on expected returns (the prior) and reinverts the linear combination of expected returns (the posterior) into a final portfolio.

The capital asset pricing model (CAPM; see Sharpe 1964; Lintner 1965) suggests that the market portfolio is an excellent starting point because, under fairly restrictive assumptions, in equilibrium the expected return from diversifiable risk is equal to zero and the market portfolio should have the highest Sharpe ratio.[1] However, Black–Litterman also allows personal views on expected returns for assets that are deemed to be away from their equilibrium value. As we will show, the Black–Litterman solution represents returns that are a weighted average of the market portfolio (or, more generally, the data model portfolio) and a portfolio that

[1] Roll (1978) dimmed some hopes by pointing out that the ultimate market portfolio is unobservable in the real world, but he also pointed out that any very well-diversified cap-weighted portfolio would have little unsystematic risk, and that portfolios of these indexes would be nearly efficient and still have returns directly related to their betas. These insights led to the philosophy that, with the preponderance of people all keeping the market efficient through analysis, trying to beat the market was a fool's errand. As Ellis (1975) stated, indexing and long-term goal planning were the way to avoid playing a "loser's game".

fully incorporates the private expected return model (the prior). This suggests that the many desirable properties of the Black–Litterman methodology, documented in Black and Litterman (1992) and Bevan and Winkelmann (1998), are driven by diversification.

As Markowitz (1952) suggests, the efficient frontier with a risk-free asset available is defined by a linear function of portfolio weights, $\boldsymbol{w}^*$, that solves the problem

$$\min_{\boldsymbol{w}} \tfrac{1}{2}\boldsymbol{w}'V\boldsymbol{w} - \lambda(T - r_f - (\boldsymbol{r} - r_f)'\boldsymbol{w}), \tag{2.1}$$

where $\boldsymbol{r}$ is a vector of expected returns, $\boldsymbol{w}$ is the vector of portfolio weights, V is the covariance matrix of excess returns, λ is a Lagrange multiplier closely related to the inverse of risk aversion, T is a target return and r_f is the risk-free rate. It is easily shown that the capitalization-weighted portfolio is a solution to (2.1), assuming the assumptions of the CAPM are not violated. This vector of returns (the starting solution) may be expressed as

$$\boldsymbol{r}^* = r_f\boldsymbol{e} + \left(\frac{\overline{r_m - r_f}}{\sigma_{r_m}}\right)^{-1}\sigma_{r_m}V\boldsymbol{w}^*, \tag{2.2}$$

where r_m is the capitalization-weighted (market) return, $\boldsymbol{e}$ is a conformable vector of ones and $\boldsymbol{w}^*$ is the vector of capitalization weights. Note that the term in parentheses represents the Sharpe ratio. It may be estimated from the historical data, or from other outside considerations. Its appearance comes from the fact that, for a target rate in (2.1) equal to the market return,

$$\lambda = \left(\frac{\overline{r_m - r_f}}{\sigma_{r_m}}\right)^{-1}$$

is a necessary part of the solution.

Equation (2.2) gives returns consistent with the capitalization-weighted portfolio being on the minimum variance set, along the tangency line from the risk-free rate, and is used in deriving implied returns of the Black–Litterman model. Following Black and Litterman (1992) and Meucci (2010), we allow the possibility that, besides the "public" opinion of returns (public in that it is based on the market being in equilibrium), a portfolio manager gives credence to a private model of returns: a prior belief that some additional factors are needed to determine the equilibrium return. The Black–Litterman model imposes such beliefs using a matrix of constraints on the distribution of returns. For example, if we wanted to model the belief that all returns would be equal next period (we are not advocating this), it could be written within the standard Black–Litterman framework as

$$P\boldsymbol{r}^{**} = \begin{bmatrix} 1 & -1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & -1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -1 & 0 \\ 0 & 0 & 0 & \cdots & 1 & -1 \end{bmatrix} \times \begin{bmatrix} r_1^{**} \\ r_2^{**} \\ \vdots \\ r_{n-1}^{**} \\ r_n^{**} \end{bmatrix} \sim N(\boldsymbol{\nu}, \Omega). \tag{2.3}$$

Equation (2.3) implies that the difference in the forecasted returns of any two assets is distributed normally around some values, which in this case would be a vector of zeros. This particular set of beliefs is that the difference in return of any two assets will be expected to be zero with some variability.

To account for the variability, a good choice is to follow Meucci (2010) and set

$$\Omega = \frac{1}{c} PVP^{\mathrm{T}},$$

where c represents the strength of conviction about the prior.[2]

The brilliancy of Black–Litterman optimization is that the final (posterior) return vector may be written as

$$\boldsymbol{r} = (1-x)\boldsymbol{r}^* + x\boldsymbol{r}^{**}, \tag{2.4}$$

where

$$x = \frac{c}{1+c} \quad \text{and} \quad \boldsymbol{r}^{**} = \boldsymbol{r}^* + VP^{\mathrm{T}}(PVP^{\mathrm{T}})^{-1}(\boldsymbol{\nu} - P\boldsymbol{r}^*).$$

Equation (2.4) implies that the posterior return is a portfolio of the data-model-based return and the prior-based return. This also means that the portfolio formed from the posterior returns will be a linear combination of the portfolios formed from the data-based returns and prior returns. Moreover, the portfolio resulting from the posterior returns (by (2.2)) will also be a linear combination of the portfolios obtained from the starting returns and the views-based returns.[3] Since these portfolios will not generally have perfect correlation, this yields diversification benefits that can potentially

[2] A c near zero would mean no belief in the prior (and thus it should not be considered), whereas c tending toward infinity would indicate certainty that the prior is the correct forecast for the next period (this return is a conditional return, based on the state of the efficient frontier, even if we think it is, at least in part, completely incorrect).

[3] One minor problem is that the Bayesian return vector might not reinvert into a portfolio with normalized weights using (2.2) if assets are highly correlated. In this paper, we simply rescale the weights to equal 1 in total. An alternative would be to take the returns and apply them to Markowitz optimization or a more general quadratic program directly.

improve the risk-adjusted return of the data portfolio, as long as the prior returns are sufficiently strong.[4]

3 THE BLACK–LITTERMAN FRAMEWORK WITH EXOTIC BETAS

What became known as "exotic beta" appeared in the literature in the very early days of equilibrium asset pricing theory. Deriving from the literature originally referred to as the anomaly literature, meaning things that are outside the standard canon of CAPM orthodoxy, such non-CAPM risk factors began to really take shape in work by Fama and French (1992) and, later, Carhart *et al* (1997). Carhart *et al* (2014) explored the notion of exotic beta rather fully and concluded that exotic beta is a powerful portfolio management tool. They defined exotic betas as exposures to risk factors that are uncorrelated with global equity markets and have positive expected returns. This definition suggests that exotic betas are not stock-specific "alphas" in the traditional sense of being an unsystematic return specific to a stock, but rather are sensitivities to risk premiums that are outside of the CAPM, and presumably more stable than the market risk premium.

Exotic betas (or any other "alpha" model) might seem to invalidate the Black–Litterman framework, which assumes that the market portfolio is nearly efficient. If the CAPM were strictly true, the market portfolio would account for all systematic risk, and no prior return vector (due to exotic beta, or otherwise) would be needed. So, when we are using a Black–Litterman framework, the starting index portfolio may be thought of as approximately on the efficient frontier, but there is some inefficiency (or weighted group of inefficiencies) that can be added to the market returns to make more accurate asset return forecasts.

Another interpretation of the Black–Litterman approach is to treat it as a mixture model, where we believe there is a $1/(1+c)$ chance that the market is efficient by itself and the implied expectations are correct, and a $c/(1+c)$ chance that a different set of return expectations is the efficient portfolio. The combination of the two probabilities results in hybrid forecasts of asset returns that are part market portfolio return and part exotic beta return. Of course, many methods of applying exotic betas are possible, but the additional benefit of this method is that it provides a distinctly different approach to extracting returns from exotic betas from those documented in the literature.

The nice thing about (2.3) is that it allows the reverse-optimization returns to scale the exotic beta returns. Looking closely at (2.3), we note that there are $n-1$ equations and n assets. This allows infinitely many solutions. However, (2.4) specifies that the

[4] The required strength of the prior returns can be expressed using a simple inequality. Prior returns y will improve performance of the data-based returns x if and only if $\mathrm{Sharpe}(y) > \mathrm{corr}(x, y) \times \mathrm{Sharpe}(x)$. Though the CAPM suggests that there should be no priors strong enough to improve performance, a rich literature on market anomalies presents evidence that such priors exist.

solution chosen will be the one that puts the exotic beta returns on the same scale as the reverse-optimization returns.

To make this discussion more concrete, Table 1 gives descriptive statistics for ten Dow Jones industrial stocks, whose returns are observed monthly, for the period from January 1995 to May 2015. This small number of stocks was chosen to make covariance estimation a trivial issue that would not interfere with the main points of the paper.[5]

We illustrate how to incorporate exotic betas in the Black–Litterman framework by starting with a capitalization-weighted portfolio of the ten stocks and using the low-volatility anomaly, described in Jagannathan and Ma (2003), as an example of an exotic beta. In this case, the future expected Sharpe ratio is considered to be an inverse function of previous volatility. An easy way to represent this prior within the current framework is to state, for any pair of assets, the expected difference in their Sharpe ratios going forward as a percentage of the trailing difference in their inverse volatility.

Specifically, this prior return, consistent with the low-volatility anomaly, may be written for all assets as

$$\boldsymbol{P}\boldsymbol{r}^{**} = \begin{bmatrix} \dfrac{1}{\sigma_1} & -\dfrac{1}{\sigma_2} & 0 & \cdots & 0 & 0 \\ 0 & \dfrac{1}{\sigma_2} & -\dfrac{1}{\sigma_3} & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -\dfrac{1}{\sigma_{n-1}} & 0 \\ 0 & 0 & 0 & \cdots & \dfrac{1}{\sigma_{n-1}} & -\dfrac{1}{\sigma_n} \end{bmatrix} \times \begin{bmatrix} r_1^{**} \\ r_2^{**} \\ \vdots \\ r_{n-1}^{**} \\ r_n^{**} \end{bmatrix} \sim N(\boldsymbol{v}, \Omega), \quad (3.1)$$

with

$$v_i = \alpha\left(\frac{1}{\sigma_i} - \frac{1}{\sigma_{i+1}}\right).$$

For this example we use alpha equal to 0.001, and the prior return equation indicates there is an expectation that the difference in Sharpe ratio between two assets will be directly related to the difference in their inverse volatilities.

---

[5] There is no consensus as to the optimal estimation method for larger covariance matrixes, but a number of approaches have been introduced. Ledoit and Wolf (2004) suggest shrinking a sample covariance matrix; Jagannathan and Ma (2003) consider using daily returns and factor models in addition to shrinkage; Pafka *et al* (2004) argue for applying filtering based on the random matrix theory. The issue of estimating the covariance matrix is beyond the scope of this paper. For this particular small problem, covariance estimates rely on a simple sixty-month sample.

TABLE 1 Descriptive statistics of the ten Dow Jones 30 stocks.

Stocks	Annual excess return (%)	Annual SD (%)	Sharpe ratio	Pairwise correlations									
				HD	DIS	NKE	MCD	PG	WMT	KO	CVX	XOM	TRV
HD	10.63	26.94	0.39	1.00									
DIS	7.83	25.72	0.30	0.33	1.00								
NKE	13.74	30.31	0.45	0.32	0.28	1.00							
MCD	8.15	21.34	0.38	0.31	0.37	0.34	1.00						
PG	7.48	20.29	0.37	0.09	0.10	0.16	0.25	1.00					
WMT	8.06	22.13	0.36	0.51	0.13	0.35	0.30	0.16	1.00				
KO	5.19	21.40	0.24	0.14	0.30	0.26	0.40	0.46	0.26	1.00			
CVX	8.32	19.87	0.42	0.25	0.32	0.25	0.37	0.07	0.15	0.25	1.00		
XOM	8.34	16.74	0.50	0.14	0.29	0.16	0.34	0.22	0.08	0.33	0.72	1.00	
TRV	7.20	25.58	0.28	0.31	0.36	0.37	0.32	0.13	0.26	0.34	0.37	0.30	1.00

Sample period: January 1995 to May 2015. SD denotes standard deviation.

TABLE 2 Black–Litterman applied to the ten-stock example using the low-volatility anomaly.

	$c = 0$	$c = 1/3$	$c = 1$	$c = 3$	$c = \infty$
Annualized excess return (%)	3.64	3.92	4.21	4.50	4.79
Annualized standard deviation (%)	12.17	12.15	12.45	13.10	14.08
Sharpe ratio	0.299	0.323	0.338	0.344	0.340

Monthly rebalancing. Out-of-sample period: February 2000 to May 2015.

TABLE 3 Comparison of two implementations of the low-volatility anomaly.

	Standard implementation	Bayesian prior approach
Annualized excess return (%)	0.68	1.15
Annualized standard deviation (%)	11.78	8.39
Sharpe ratio	0.06	0.14

Standard refers to a typical long–short implementation. The Bayesian prior approach is a long–short portfolio that goes long Black–Litterman with infinite c (100% weight to the low-volatility prior) and short market portfolio. Out-of-sample period: February 2000 to May 2015.

Table 2 outlines the results of applying the Black–Litterman model to a market-weighted reverse-optimization model and a low-volatility prior with several levels of c that represent 25%, 50%, 75% and 100% weights to the low-volatility anomaly prior.[6] In this case, the maximal Sharpe ratio of 0.344 is accomplished at $c = 3$, representing a 75% weight to the prior, which is superior to 0.299, the Sharpe ratio of the market portfolio; this suggests that the low-volatility anomaly view used within the Black–Litterman framework can substantially improve performance.

Another interesting experimental result is shown in Table 3, which compares the performance of the standard implementation of the low-volatility anomaly and our implementation, based on the Black–Litterman framework. The standard implementation ranks stocks based on their in-sample volatility and then buys the bottom quintile of stocks (low-volatility stocks) and sells short the top quintile of stocks (high-volatility stocks). Portfolio weights are inversely proportional to historical inverse volatilities. The Black–Litterman approach purchases the highest Sharpe ratio portfolio implied by Black–Litterman returns with $c = \infty$ (100% weight to the low-volatility anomaly prior) and sells short the market portfolio.[7]

[6] For example, if $c = 1/3$, this represents a $1/(1+c) = 25\%$ weight assigned to return expectations of the market portfolio and a 75% weight assigned to those of the prior.

[7] Once we have the Black–Litterman returns we can use standard Markowitz optimization or constrained quadratic optimization. The results in Table 3 use standard Markowitz optimization to find the highest forecast Sharpe ratio portfolio.

The Sharpe ratio of the Bayesian prior approach is equal to 0.14, which is more than twice 0.06, the Sharpe ratio of the standard implementation of the low-volatility anomaly. Though the relative performance of the two implementations of the low-volatility anomaly (or any exotic beta in general) can be sensitive to the time period, portfolio constituents and choice of parameters, the Bayesian prior approach substantially expands the toolbox of exotic beta implementation with potentially significant performance implications for investors.

Carhart *et al* (2014) suggest that utilizing a limited version of Black–Litterman with exotic betas as portfolio constituents is unlikely to diminish its power. We have extended this result by showing that the standard Black–Litterman implementation can be combined with exotic betas, by using them as views, to achieve multiple useful results. In the next sections, we extend the notion of reverse optimization to utilize portfolios other than the capitalization-weighted portfolio. This paper is, to the best of our knowledge, the first to investigate the implications of alternative efficient portfolios for Black–Litterman optimization and suggest a version of Black–Litterman optimization that extends its application to many new investment situations.

4 THE MARKET PORTFOLIO, THE RISK-PARITY PORTFOLIO AND EFFICIENCY

While the CAPM suggests that a capitalization-weighted market portfolio should have the highest Sharpe ratio, there are situations in which the market portfolio is either suboptimal or even inappropriate. For example, fund-of-hedge-funds allocation decisions, and decisions involving futures contracts generally, do not readily admit capitalization calculations, which makes capitalization-weighted market portfolios unattainable. Moreover, Asness *et al* (2012) suggest that the market portfolio is not efficient, and instead argue that the risk-parity portfolio approach, which equalizes the contribution to portfolio risk from each constituent, is more efficient due to leverage aversion.[8] Qian (2006) provides a comprehensive analysis of risk-parity portfolios. We outline the technical details of the risk-parity approach in the online appendix.

Since capitalization weights are available in the ten-stock Dow Jones example, we can easily compare the performance of the capitalization-weighted market portfolio and risk parity. Table 4 reports the out-of-sample performance of the two portfolios.

The risk-parity portfolio delivers a Sharpe ratio of 0.45, which is higher than 0.30, the Sharpe ratio of the market portfolio. However, we need to be careful about drawing

[8] Asness *et al* (2012) argue that leverage aversion changes the predictions of modern portfolio theory, because investors without access to leverage are unable to benefit from higher risk-adjusted returns of safer (low-beta or low-volatility) assets. Risk parity portfolios overweight safer assets relative to the market portfolio and benefit from their higher risk-adjusted returns after applying leverage.

TABLE 4 Comparison of capitalization-weighted market portfolio and risk parity.

	Market	Risk parity
Annualized excess return (%)	3.64	5.74
Annualized standard deviation (%)	12.17	12.89
Sharpe ratio	0.30	0.45

Out-of-sample period: February 2000 to May 2015.

conclusions about the efficiency of risk parity from this simple example. Anderson *et al* (2012) argue empirical studies that make claims about the efficiency of risk parity might be very sensitive to the time period studied and the transaction costs assumed. Our simple example is also not immune to their criticism.

To summarize, there is reason to believe that, at least sometimes, an equal-risk-weighted index portfolio may be more efficient than a capitalization-weighted portfolio. Also, there are times when capitalization weights are unavailable. In either of these situations, the risk-parity portfolio is a good candidate for the data-based starting point in the Black–Litterman framework.[9]

5 AN ALTERNATIVE REVERSE OPTIMIZATION

The importance of the Black–Litterman framework is that it provides a way of mixing market-data-based returns with prior views about assets that are not priced properly by the market. A key thing to realize is that reversing the first-order condition in the Markowitz model is theoretically permissible for any portfolio on the minimum variance set. The mathematics are only slightly more complicated, and the basics are outlined below.

The Markowitz model may be written in an alternative form as

$$\min \tfrac{1}{2}\boldsymbol{w}'V\boldsymbol{w} - \lambda(T - \boldsymbol{r}^{\mathrm{T}}\boldsymbol{w}) + \gamma(1 - \boldsymbol{w}'\boldsymbol{e}).$$

In this formulation there are two constraints. The risky assets' weights must sum to 1, and the target expected return must be achieved. By varying the target, we map the entire minimum variance set. This formulation leads to the following reverse-optimization returns:

$$\boldsymbol{r} = \frac{V\boldsymbol{w}_{\mathrm{p}} - \gamma_{\mathrm{p}}\boldsymbol{e}}{\lambda_{\mathrm{p}}}, \tag{5.1}$$

where $\boldsymbol{w}_{\mathrm{p}}$ are the weights of the presumed efficient portfolio.

[9] Although some people may find the notion that the equal risk portfolio is possibly efficient strange, even less likely is the fact that diversified portfolios have been shown to be efficient in some cases. DeMiguel *et al* (2009) maintain that, for their universe, the equal-weighted $1/N$ portfolio was more efficient than more conventional alternatives.

FIGURE 1 Reverse-optimization scenarios.

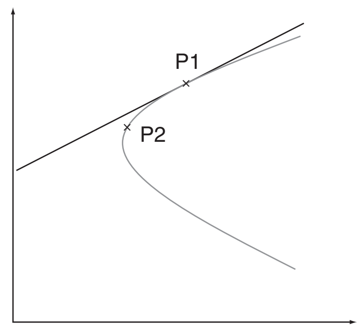

Now, to complete the solution of (5.1), we need to define r_{mv} as the return of the minimum variance portfolio, whose weights are given by

$$\boldsymbol{w}_{\mathrm{mv}} = \frac{V^{-1}\boldsymbol{e}}{\boldsymbol{e}V^{-1}\boldsymbol{e}},$$

which is a function only of the elements of the covariance matrix, not expected returns. Now,

$$\left.\begin{aligned} \lambda_{\mathrm{p}} &= \frac{\sigma_{\mathrm{p}}^2 - \sigma_{\mathrm{mv,p}}}{\bar{r}_{\mathrm{p}} - \bar{r}_{\mathrm{mv}}}, \\ \gamma_{\mathrm{p}} &= \lambda_{\mathrm{p}}\bar{r}_{\mathrm{p}} - \sigma_{\mathrm{p}}^2. \end{aligned}\right\} \tag{5.2}$$

The individual reverse-optimization returns may be obtained using (5.1) and (5.2). In this formulation, we have replaced the use of information about the risk-free rate and the capitalization-weighted portfolio with the use of information about the minimum variance portfolio and our assumed minimum variance portfolio.

Now that we have two reverse-optimization methodologies, which should we use? Figure 1 helps us answer this question. If we believe our portfolio is a minimum variance portfolio and is the one with the highest Sharpe ratio, then it is at point P1 and we should use (2.2). This is the market portfolio, if the CAPM is strictly true, but may be another portfolio if it is not. Only if the CAPM is true will P1 be the capitalization-weighted portfolio. If we believe that we have a minimum variance portfolio (mainly a very well-diversified portfolio) but not that it is at P1 (perhaps it is at P2), then we should use (5.1) and (5.2). Equations (5.1) and (5.2) are useful when the only thing we know is that our starting portfolio is very well diversified with respect to the investment universe. Equation (2.2) is useful when we believe our starting portfolio has not only the minimum variance but also the maximal Sharpe

TABLE 5 Descriptive statistics of bonds, stocks and commodities.

	Bonds	Stocks	Commodities
Annualized excess return (%)	2.79	3.36	–0.71
Annualized standard deviation (%)	5.45	14.63	19.26
Sharpe ratio	0.51	0.23	–0.04

This table shows the annualized excess return, standard deviation and Sharpe ratio of bonds, stocks and commodities for March 1976 through May 2015. The Barclays US Aggregate Government Index is used as a proxy for the bond market, The MSCI World Index is used as a proxy for the stock market, the S&P Goldman Sachs Commodities Index is used as a proxy for the commodities market and the three-month treasury bill secondary market rate is used as a proxy of the risk-free rate.

ratio with respect to the universe. That this portfolio is the capitalization-weighted portfolio is a function of also believing the CAPM is strictly true. This discussion opens up reverse optimization as the starting point for a very large number of possible portfolios.

Having established in the previous section that the risk parity can be considered a reasonable maximal Sharpe ratio portfolio, we will take it as a starting point for Black–Litterman using (2.2). One interesting thing to note (discussed in the online appendix) is that the risk-parity portfolio simply says that each asset's contribution to total risk should be made equal. This does not mean that each asset should have the same Sharpe ratio or the same expected return.

6 TWO EXAMPLES OF A RISK-PARITY STARTING PORTFOLIO

We illustrate the risk-parity and exotic beta approach by considering two simple examples with investments in three major asset classes: stocks, bonds and commodities. The data for these experiments covers the period from March 1976 through May 2015, and consists of the MSCI World Index, as a proxy for the stock market, the Barclays US Aggregate Government Index, to represent the bond market, and the Standard & Poor's (S&P) Goldman Sachs Commodity Index, to represent commodities. We use the three-month treasury bill secondary market rate as a proxy for the risk-free rate. Inclusion of commodities is particularly interesting because they do not have capitalization weights, and therefore the capitalization-weighted market portfolio is unattainable.[10]

Table 5 presents some relevant descriptive statistics of the data. One interesting statistic is that the three asset classes perform very differently over this time period, with commodities performing worst, with a Sharpe ratio of –0.04, and bonds performing best, with a Sharpe ratio of 0.51.

[10] Portfolio allocations that involve hedge funds are another example of limitation in the capitalization-weighted approach.

However, the relative performance of the three asset classes across time is very inconsistent. Figure 2 displays the rolling twenty-four-month Sharpe ratio of the three assets.[11]

Absolute and relative performance is inconsistent across time, with the Sharpe ratios ranging between around -2.5 and around $+2.5$.

The first step in this extended Black–Litterman approach is calculating risk-parity weights and the corresponding implied expected excess returns using (2.2). The second step involves imposing exotic-beta-based views. We consider two examples that can have broad applications to many investors. The first example considers momentum as an exotic beta that is robust across most asset classes, as documented in Asness *et al* (2013). The second example considers an equal Sharpe ratio prior.

6.1 Risk parity with momentum

Momentum is a pervasive anomaly that has been extensively documented in the literature. We express belief in momentum in Sharpe ratios using (3.1) with the same matrix P and vector

$$v_i = \alpha\left(\frac{\bar{R}_i}{\sigma_i} - \frac{\bar{R}_{i+1}}{\sigma_{i+1}}\right).$$

We set $\alpha = 0.05$, which represents the belief that the difference in the future Sharpe ratios of any two assets is expected to equal 5% of the most recent Sharpe ratios. In this simple example we use a time window of twenty-four months to estimate recent Sharpe ratios. As before, we use the same levels of c, representing 25%, 50%, 75% and 100% weights to the exotic beta belief. The results of this experiment are presented in Table 6.

In this case, the maximal Sharpe ratio of 0.35 occurs at $c = 3$, representing a 75% weight to the prior, which is superior to 0.253, the Sharpe ratio of the risk-parity portfolio; this suggests that the momentum in the risk-adjusted return prior used within the Black–Litterman framework can improve performance.

6.2 Risk parity with equal Sharpe ratios

Although belief in equal Sharpe ratios is unrelated to exotic beta, it is an interesting case to study. There is a sizable group of investors who hold this belief about diversified asset classes. In addition, there are fund-of-(hedge) funds managers who have very high requirements for their hedge funds; this is reflected in their rigorous due diligence steps, which result in approximately the same Sharpe expectations across hedge funds

[11] In this section, we choose to use rolling twenty-four-month sample estimates for all covariance and variance estimates. With only three assets, a longer period is not required.

FIGURE 2 Rolling twenty-four-month Sharpe ratios of stocks, bonds and commodities.

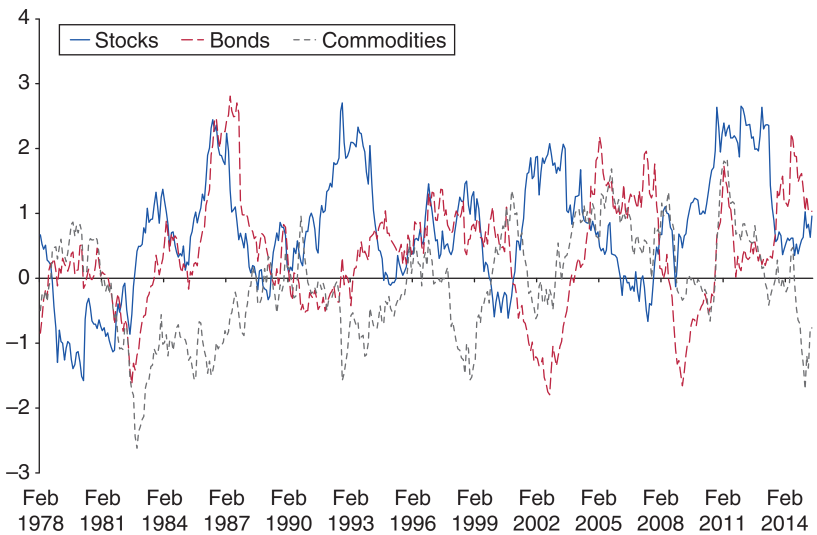

TABLE 6 The three-asset example with risk parity and momentum.

	c				
	0	**1/3**	**1**	**3**	∞
Annualized excess return (%)	2.08	2.07	2.02	1.96	1.90
Annualized standard deviation (%)	8.21	6.52	5.89	5.59	5.44
Sharpe ratio	0.253	0.318	0.343	0.350	0.349

in the portfolio. Finally, it is illuminating to show how even a fairly weak prior performs in the Black–Litterman framework.

Table 5 shows that commodities substantially underperformed stocks and bonds over the time period in this study. However, other studies, such as Gorton and Rouwenhorst (2006), argue that an equally weighted index of commodity futures' monthly returns should deliver a Sharpe ratio comparable to that of equities. The equal Sharpe belief can be expressed using (3.1) with the mean vector set to zero.

Table 7 reports results of the out-of-sample performance. The maximal Sharpe ratio of 0.316 corresponds to $c = 1$, representing a 50% weight to the prior, which is superior to 0.253, the Sharpe ratio of the risk-parity portfolio. This suggests that the prior of equal Sharpe ratios used within the Black–Litterman framework can add value.

TABLE 7 The three-asset example with risk parity and equal Sharpe priors.

	c				
	0	**1/3**	**1**	**3**	∞
Annualized excess return (%)	2.08	2.03	1.93	1.81	1.69
Annualized standard deviation (%)	8.21	6.65	6.10	5.87	5.79
Sharpe ratio	0.253	0.305	0.316	0.308	0.292

Monthly rebalancing. Out-of-sample period: March 1978 to May 2015.

FIGURE 3 Risk-parity performance

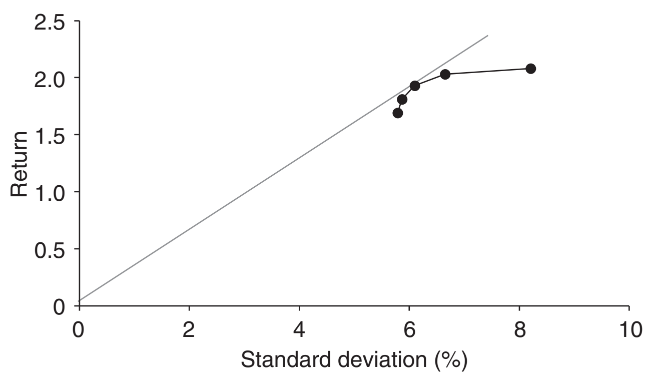

The figure displays results from Table 7. $c = \infty$ is the point furthest south and $c = 0$ is the point furthest northeast, representing the classical risk-parity approach. The line from the origin is tangential to the frontier at the point with the highest Sharpe ratio, which is approximately $c = 1$. The axes show excess returns and excess standard deviation.

The above approach gives combinations of starting returns and view returns that allow us to reach Sharpe ratios unobtainable through either approach alone. The exact nature of this improvement is apparent in Figure 3, which presents the data in Table 7 in graphical form, and once again demonstrates the diversifying power of Black–Litterman even in the face of a rather weak prior.

The curve, reminiscent of a Markowitz frontier, reiterates the benefits of diversification produced by the Bayesian risk-parity method. The straight line from the origin to the combination line of strategic outcomes shows that the maximum obtainable Sharpe ratio is at approximately $c = 1$ for this particular combination of strategies and assets over this particular period. While it is true that this diversification benefit would occur with almost any prior, the better the prior actually is, of course, the better the opportunities are.

7 CONCLUSION

In this paper, we demonstrate that using the exotic beta as a prior alpha model in Black–Litterman optimization is attractive to investors that already utilize the classic Black–Litterman approach and seek to incorporate advances in the exotic beta research, and to those who focus on practical implementation of exotic betas. The reason for this behavior is that the diversification of alpha sources benefits an investor, whether they want a portfolio that is mainly efficient-portfolio based, mainly exotic-beta based or one that maximizes the Sharpe ratio.

In addition, we introduce risk parity as a valid starting portfolio, and produce a methodology for using almost any well-diversified portfolio as a starting portfolio. This is useful when the capitalization-weighted portfolio is not an appropriate starting point.

Our extended Black–Litterman approach symbiotically unifies the Black–Litterman optimization, exotic betas and risk parity into a single, flexible framework that combines the various strengths of the three approaches to improve investors' portfolios. These results give a large number of investment professionals new tools for their investment tool box without throwing out everything they might previously have been using.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

REFERENCES

Anderson, R., Bianchi, S., and Goldberg, L. (2012). Will my risk parity strategy outperform? *Financial Analysts Journal* **68**(6), 75–93 (http://doi.org/b6p9).

Asness, C., Frazzini, A., and Pedersen, L. (2012). Leverage aversion and risk parity. *Financial Analysts Journal* **68**(1), 47–59 (http://doi.org/fxt9j9).

Asness, C., Moskowitz, T., and Pedersen, L. (2013). Value and momentum everywhere. *Journal of Finance* **68**(3), 929–985 (http://doi.org/bqw3).

Bevan, A., and Winkelmann, K. (1998). Using the Black–Litterman global asset allocation model: three years of practical experience. White Paper, Goldman Sachs.

Black, F., and Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal* **48**(5), 28–43 (http://doi.org/d79fhm).

Carhart, M. (1997). On persistence in mutual fund performance. *Journal of Finance* **52**(1), 57–82.

Carhart, M., Cheah, U.-W., De Santis, G., Farrell, H., and Litterman, R. (2014). Exotic betas revisited. *Financial Analysts Journal* **70**(5), 24–52 (http://doi.org/b6qb).

DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus naive diversification: how inefficient is the $1/N$ portfolio strategy? *Review of Financial Studies* **22**(5), 1915–1953 (http://doi.org/dzkc89).

Ellis, C. D. (1975). The loser's game. *Financial Analysts Journal* **31**(4), 19–26 (http://doi.org/dttwvf).

Fama, E., and French, K. (1992). The cross-section of expected stock returns. *Journal of Finance* **47**(2), 427–465.

Gorton, G., and Rouwenhorst, K. G. (2006). Facts and fantasies about commodity futures. *Financial Analysts Journal* **62**(2), 47–68 (http://doi.org/dtjqhw).

Jagannathan, R., and Ma, T. (2003). Risk reduction in large portfolios: why imposing the wrong constraints helps. *Journal of Finance* **58**(4), 1651–1684 (http://doi.org/c95xn4).

Ledoit, O., and Wolf, M. (2004). Honey, I shrunk the sample covariance matrix. *Journal of Portfolio Management* **30**(4), 110–119 (http://doi.org/b93gvf).

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* **47**(1), 13–37.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance* **7**(1), 77–91 (http://doi.org/bhzd).

Meucci, A. (2010). Black–Litterman approach. In *Encyclopedia of Quantitative Finance*, Cont, R. (ed), pp. 196–199. Wiley (http://doi.org/b3btkt).

Pafka, Sz., Potters, M., and Kondor, I. (2004). Exponential weighting and random-matrix-theory-based filtering of financial covariance matrices for portfolio optimization. Working Paper, arXiv (http://bit.ly/2qIvBYM).

Qian, E. (2006). On the financial interpretation of risk contribution: risk budgets do add up. *Journal of Investment Management* **4**, 41–51.

Roll, R. (1978). Ambiguity when performance is measured by the securities market line. *Journal of Finance* **33**(4), 1051–1069 (http://doi.org/b6qc).

Sharpe, W. F. (1964). Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance* **19**(4), 425–442.

Yang, Z., Han, K., Molyboga, M., and Molyboga, G. (2016). A new diagnostic approach to evaluating the stability of optimal portfolios. *Journal of Investing* **25**(1), 37–45 (http://doi.org/b6qd).