Long-Short Portfolio Management: An Integrated Approach

The real benefits of long-short are released only by an integrated portfolio optimization.

Bruce I. Jacobs, Kenneth N. Levy, and David Starer

Most investors focus on the management of long portfolios and the selection of "winning" securities. Yet the identification of winning securities ignores by definition a whole class of "losing" securities. The ability to sell short frees the investor to take advantage of the full array of securities and the full complement of investment insights by holding expected winners long and selling expected losers short.

A long-short portfolio, by expanding the scope of the investor's sphere of activity, can be expected to result in improved performance from active security selection vis-à-vis a long-only portfolio. But the benefits of long-short are to a large extent dependent on proper portfolio construction. Only an integrated optimization of long and short positions has the potential to maximize the value of investors' insights. The benefits that emerge from integrated optimization encompass not only freedom from the short-selling constraint but also freedom from the restrictions imposed by individual securities' benchmark weights.

Of course, these benefits do not come without some cost. Much of the incremental cost associated with a given long-short portfolio reflects the strategy's degree of leverage. Nevertheless, as we will see, long-short is not necessarily much costlier or, indeed, much riskier than long-only.

Although most existing long-short portfolios are constructed to be neutral to systematic risk, we will see that neutrality is neither necessary nor, in most cases,

BRUCE I. JACOBS and **KENNETH N. LEVY** are principals and **DAVID STARER** is a senior quantitative analyst at Jacobs Levy Equity Management in Roseland (NJ 07068)

optimal. Furthermore, we show that long-short portfolios do not constitute a separate asset class; they can, however, be constructed to include a desired exposure to the return (and risk) of virtually any existing asset class.

LONG-SHORT: BENEFITS AND COSTS

Consider a long-only investor who has an extremely negative view about a typical stock. The investor's ability to benefit from this insight is very limited. The most the investor can do is exclude the stock from the portfolio, in which case the portfolio will have about a 0.01% underweight in the stock, relative to the underlying market.[1] Those who do not consider this to be a material constraint should consider what its effect would be on the investor's ability to overweight a typical stock. It would mean the investor could hold no more than a 0.02% long position in the stock — a 0.01% overweight — no matter how attractive its expected return.

The ability to short, by increasing the investor's leeway to act on insights, has the potential to enhance returns from active security selection.[2] The scope of the improvement, however, will depend critically on the way the long-short portfolio is constructed. In particular, an integrated optimization that considers both long and short positions simultaneously not only frees the investor from the non-negativity constraint imposed on long-only portfolios, but also frees the long-short portfolio from the restrictions imposed by securities' benchmark weights. To see this, it is useful to examine one obvious (if suboptimal) way of constructing a long-short portfolio.

Long-short portfolios are sometimes constructed by combining a long-only portfolio, perhaps a preexisting one, with a short-only portfolio. This results in a long-plus-short portfolio, not a true long-short portfolio. The long side of this portfolio is identical to a long-only portfolio; hence it offers no benefits in terms of incremental return or reduced risk.

In long-plus-short, the short side is statistically equivalent to the long side, hence to the long-only portfolio.[3] In effect:

$$\alpha_L = \alpha_S = \alpha_{LO}$$

$$\omega_L = \omega_S = \omega_{LO}$$

That is, the excess return or alpha, α, of the long side of the long-plus-short portfolio will equal the alpha of the short side, which will equal the alpha of the long-only portfolio. Furthermore, the residual risk of the long side of the long-plus-short portfolio, ω, will equal the residual risk of the short side, which will equal the residual risk of the long-only portfolio.

These equivalencies reflect the fact that all the portfolios, the long-only portfolio and the long and short components of the long-plus-short portfolio, are constructed relative to a benchmark index. Each portfolio is active in pursuing excess return relative to the underlying index only insofar as it holds securities in weights that depart from their index weights. The ability to pursue such excess returns may be limited by the need to control the portfolio's residual risk by maintaining portfolio weights that are close to index weights. Portfolio construction is index-constrained.

Consider, for example, an investor who does not have the ability to discriminate between good and bad oil stocks, or who believes that no oil stock will significantly outperform or underperform the underlying benchmark in the near future. In long-plus-short, this investor may have to hold some oil stocks in the long portfolio and short some oil stocks in the short portfolio, if only to control each portfolio's residual risk.

The ratio of the performance of the long-plus-short portfolio to that of the long-only portfolio can be expressed as follows:[4]

$$\frac{IR_{L+S}}{IR_{LO}} = \sqrt{\frac{2}{1+\rho_{L+S}}}$$

where IR is the information ratio, or the ratio of excess return to residual risk, α/ω, and ρ_{L+S} is the correlation between the alphas of the long and short sides of the long-plus-short portfolio.

In long-plus-short, the advantage offered by the flexibility to short is curtailed by the need to control risk by holding or shorting securities in index-like weights. A long-plus-short portfolio thus offers a benefit over a long-only portfolio only if there is a less-than-one correlation between the alphas of its long and short sides. In that case, the long-plus-short portfolio will enjoy greater diversification and reduced risk relative to a long-only portfolio. A long-only portfolio can derive a similar benefit by adding a less than fully correlated asset with comparable risk and return, however, so this is not a benefit unique to long-short.

The Real Benefits of Long-Short

The real benefits of long-short emerge only when the portfolio is conceived of and constructed as a single, integrated portfolio of long and short positions. In this framework, long-short is not a two-portfolio strategy. It is a one-portfolio strategy in which the long and short positions are determined jointly within an optimization that takes into account the expected returns of the individual securities, the standard deviations of those returns, and the correlations between them, as well as the investor's tolerance for risk.

Within an integrated optimization, there is no need to converge to securities' benchmark weights in order to control risk. Rather, offsetting long and short positions can be used to control portfolio risk. This allows the investor greater flexibility to take active positions.

Suppose, for example, that an investor's strongest insights are about oil stocks, some of which are expected to do especially well and some especially poorly. The investor does not have to restrict the portfolio's weightings of oil stocks to index-like weights in order to control the portfolio's exposure to oil sector risk. The investor can allocate much of the portfolio to oil stocks, held long and sold short. The offsetting long and short positions control the portfolio's exposure to the oil factor.

Conversely, suppose the investor has no insights into oil stock behavior. Unlike the long-only and long-plus-short investors discussed above, the long-short investor can totally exclude oil stocks from the portfolio. The exclusion of oil stocks does not increase portfolio risk, because the long-short portfolio's risk is independent of any security's benchmark weight. The flexibility afforded by the absence of the restrictions imposed by securities' benchmark weights enhances the long-short investor's ability to implement investment insights.

Costs: Perception versus Reality

Long-short construction maximizes the benefit obtained from potentially valuable investment insights by eliminating long-only's constraint on short-selling and the need to converge to securities' index weights in order to control portfolio risk. While long-short offers advantages over long-only, however, it also involves complications not encountered in long-only management. Many of these complications are related to the use of short-selling.

Costs Related to Shorting. To engage in short-selling, an investor must establish an account with a prime broker. The broker clears all trades for the long-short portfolio and arranges to borrow stock for shorting. For some shares, especially those of the smallest-capitalization companies, borrowability may be problematic. Even when such shares are available for borrowing, they may pose a problem for the short-seller if they are later called back by the stock lender. In that case, the broker may not be able to find replacement shares, and the long-short investor will be subject to a "buy-in" and have to cover the short positions.

The financial intermediation cost of borrowing, which includes the costs associated with securing and administering lendable stocks, averages 25 to 30 basis points and may be higher for harder-to-borrow names. This cost is incurred as a "haircut" on the short rebate received from the interest earned on the short sale proceeds.

Short-sellers may also incur trading opportunity costs because exchange rules delay or prevent short sales. Securities and Exchange Commission Rule 10a-1, for example, states that exchange-traded shares can be shorted only at a price that is higher than the last trade price (an uptick) or the same as the last trade price if that price was higher than the previous trade (zero-plus-tick).

Such tick tests can be circumvented by the use of "principal packages" (traded outside U.S. markets) or the sale of call options, but the costs involved may be higher than the costs exacted by the rules themselves. For a long-short strategy that engages in patient trading, where the plan is to sell short only after a price rise, the incremental impact of uptick rules will be minimal.

Trading Costs. Some other costs of long-short may seem as though they should be high relative to long-only and are often portrayed as such. For example, a long-short portfolio that takes full advantage of the leverage allowed by Federal Reserve Board Regulation T (two-to-one leverage) will engage in about twice as much trading activity as a comparable unlevered long-only strategy. The differential, however, is largely a function of the portfolio's leverage. Long-short management does not require leverage. Given capital of $10 million, for example, the investor can choose to invest $5 million long and sell $5 million short; trading activity for the resulting long-short portfolio will be roughly equivalent to that for a $10 million long-only portfolio.[5]

Aside from the trading related to the sheer size of the investment in long-short versus long-only, the mechanics of long-short management may require some incremental trading not encountered in long-only. As security prices change, for example, long and short positions may have to be adjusted in order to maintain the desired degree of portfolio leverage and to meet collateralization requirements (including margin requirements and marks to market on the shorts). When a long-short portfolio is equitized by a position in stock index futures contracts, the need for such trading is reduced because price changes in the long futures positions will tend to offset marks to market on the short stock positions. (For some examples, see Jacobs [1998].)

Management Fees. Management fees for a long-short portfolio may appear to be higher than those for a comparable long-only portfolio. Again, the differential is largely a reflection of the degree to which leverage is used in the former and not in the latter. If one considers management fees per dollar of securities positions, rather than per dollar of capital, there should not be much difference between long-short and long-only.

Furthermore, investors should consider the amount of active management provided per dollar of fees. As noted, long-only portfolios must be managed with an eye to the underlying benchmark, as departures from benchmark weights introduce residual risk. In general, long-only portfolios have a sizable "hidden passive" component; only their overweights and underweights relative to the benchmark are truly active. By contrast, virtually the entire long-short portfolio is active. In terms of management fees per active dollars, then, long-short may be substantially less costly than long-only. Furthermore, long-short management is almost always offered on a performance-fee basis.

Risk. Long-short is often portrayed as inherently riskier than long-only. In part, this view reflects a concern for potentially unlimited losses on short positions. Although it is true that the risk of a short position is theoretically unlimited because there is no bound on a rise in the price of the shorted security, this source of risk is considerably mitigated in practice. It is unlikely, for example, that the prices of all the securities sold short will rise dramatically at the same time, with no offsetting increases in the prices of the securities held long. And the investor can guard against precipitous rises in the prices of individual shorted stocks by holding small positions in a large number of stocks, both long and short.

In general, a long-short portfolio will incur more risk than a long-only portfolio to the extent that it engages in leverage and/or takes more active positions. A long-short portfolio that takes full advantage of the leverage available to it will have at risk roughly double the amount of assets invested in a comparable unlevered long-only strategy. And, because it does not have to converge to securities' benchmark weights in order to control risk, a long-short strategy may take larger positions in securities with higher (and lower) expected returns compared with an index-constrained long-only portfolio.

But both the portfolio's degree of leverage and its "activeness" are within the explicit control of the investor. Furthermore, proper optimization should ensure that incremental risks, and costs, are compensated by incremental returns.

THE OPTIMAL PORTFOLIO

Here we consider what proper optimization involves, and what the resulting long-short portfolio looks like. There are some surprises. In particular, a rigorous look at long-short optimality calls into question the goals of dollar- and beta-neutrality — common practices in traditional long-short management.

We use the utility function:[6]

$$U = r_P - \frac{1}{2}\sigma_P^2 / \tau \qquad (1)$$

where r_P is the expected return of the portfolio over the investor's horizon, σ_P^2 is the variance of the portfolio's return, and τ is the investor's risk tolerance. This utility function, favored by Markowitz [1952] and Sharpe [1991], provides a good approximation of other, more general, functions and has the agreeable characteristics of providing more utility as expected return increases and less utility as risk increases.

Portfolio construction consists of two interrelated tasks: 1) an asset allocation task for choosing how to allocate the investor's wealth between a risk-free security and a set of N risky securities, and 2) a risky portfolio construction task for choosing how to distribute wealth among the N risky securities.

Let h_R represent the fraction of wealth that the investor specifically allocates to the risky portfolio, and let h_i represent the fraction of wealth invested in the i^{th} risky security. There are three components of capital that earn interest at the risk-free rate. The first is the wealth that the investor specifically allocates to the risk-free security, and this has a magnitude of $1 - h_R$. The second is the balance of the deposit made with the broker after paying for the purchase of shares long, and this has a magnitude of $h_R - \Sigma_{i \in L} h_i$, where L is the set of securities held long. The third is the proceeds of the short sales, and this has a magnitude of $\Sigma_{i \in S} |h_i| = -\Sigma_{i \in S} h_i$, where S is the set of securities sold short. (For simplicity, we assume no "haircut" on the short rebate.)

Summing these three components gives the total amount of capital h_F that earns interest at the risk-free rate as

$$h_F = 1 - \sum_{i=1}^{N} h_i$$

A number of interesting observations can be made about h_F. First, note that it is independent of h_R. Second, observe that, in the case of short-only management in which $\sum_{i=1}^{N} h_i = -1$, the quantity h_F is equal to two; that is, the investor earns the risk-free rate twice. Third, in the case of dollar-balanced long-short management in which $\sum_{i=1}^{N} h_i = 0$, the investor earns the risk-free rate only once.

Let r_F represent the return on the risk-free security, and let R_i represent the expected return on the i^{th} risky security. The expected return on the investor's total portfolio is

$$r_P = h_F r_F + \sum_{i=1}^{N} h_i R_i$$

Substituting the expression derived above for h_F into this equation gives the total portfolio return as the sum of a risk-free return component and a risky return component, expressed as $r_P = r_F + r_R$.

The risky return component is

$$r_R = \sum_{i=1}^{N} h_i r_i \tag{2-A}$$

where $r_i = R_i - r_F$ is the expected return on the i^{th} risky security in excess of the risk-free rate. The risky return component can also be expressed in matrix notation as

$$r_R = h^T r \tag{2-B}$$

where $h = [h_1, h_2, ..., h_N]^T$ and $r = [r_1, r_2, ..., r_N]^T$. It can be shown that the variance of the risky return component, σ_R^2, is

$$\sigma_R^2 = h^T Q h \tag{3}$$

where Q is the covariance matrix of the risky securities' returns. The variance of the overall portfolio is $\sigma_P^2 = \sigma_R^2$.

With these expressions, the utility function in Equation (1) can be expressed in terms of controllable variables. We determine the optimal portfolio by maximization of the utility function through appropriate choice of these variables. This maximization is performed subject to the appropriate constraints. A minimal set of appropriate constraints consists of 1) the Regulation T margin requirement, and 2) the requirement that all the wealth allocated to the risky securities is fully utilized. The solution (providing Q is non-singular) gives the optimal risky portfolio as

$$h = \tau Q^{-1} r \tag{4}$$

where Q^{-1} is the inverse of the covariance matrix. We refer to the portfolio in Equation (4) as the minimally constrained portfolio.

The optimal portfolio weights depend on predicted statistical properties of the securities. Specifically, the expected returns and their covariances must be quantities that the investor expects to be realized over the portfolio's holding period. As no investor knows the true statistical distribution of the returns, expected returns and covariances are likely to differ between investors. Optimal portfolio holdings will thus differ from investor to investor, even if all investors possess the same utility function.

The optimal holdings given in Equation (4) have a number of important properties. First, they define a portfolio that permits short positions because no non-negativity constraints are imposed during its construction. Second, they define a *single* portfolio that exploits the characteristics of *individual* securities in a single integrated optimization. Even though the single portfolio can be partitioned artificially into one sub-portfolio of only stocks held long and another sub-portfolio of only stocks sold short, there is no benefit to doing so. Third, the holdings need not satisfy any arbitrary balance conditions; dollar- or beta-neutrality is not required.

Because optimal portfolio weights are determined in a single integrated optimization, without regard to any index or benchmark weights, the portfolio has no inherent benchmark. This means that there exists no *inherent* measure of portfolio excess return or residual risk; rather, the portfolio will exhibit an absolute return and an absolute variance of return. This return can be calculated as the weighted spread between the returns to the securities held long and the returns to the securities sold short.

Performance attribution cannot distinguish between the contributions of the securities held long and those sold short; the contributions of the long and short positions are inextricably linked. Separate long and short alphas (and their correlation) are meaningless.

Neutral Portfolios

The flexibility afforded by the ability to short stocks allows investors to construct long-short portfolios that are insensitive to chosen exogenous factors. In practice, for example, most long-short portfolios are designed to be insensitive to the return of the equity market. This may be accomplished by constructing the portfolio so that the beta of the short positions equals and offsets the beta of the long positions, or (more problematically) the dollar amount of securities sold short equals the dollar amount of securities held long.[7]

Market neutrality, whether achieved through a balance of dollars or betas, may exact costs in terms of forgone utility. If more opportunities exist on the short than the long side of the market, for example, one might expect some return sacrifice from a portfolio that is required to hold equal-dollar or equal-beta positions long and short. Market neutrality could be achieved by using the appropriate amount of stock index futures, without requiring that long and short security positions be balanced.

Investors may nevertheless prefer long-short balances for "mental accounting" reasons. That is, investors may prefer to hold long-short portfolios that have no systematic risk, without requiring seemingly separate management of derivatives overlays. Even if separate managers are used for long-short and for derivatives, however, there is no necessity for long-short balance; the derivatives manager can be instructed to augment or offset the long-short portfolio's market exposure.

Imposing the condition that the portfolio be insensitive to the equity market return (or to any other factor) constitutes an additional constraint on the portfolio. The optimal neutral portfolio is the one that maximizes the investor's utility subject to all constraints, including that of neutrality.

This optimal neutral portfolio need not be, and generally is not, the same as the portfolio given by Equation (4) that maximizes the minimally constrained utility function. To the extent that the optimal neutral portfolio differs from the minimally constrained optimal portfolio, it will involve a sacrifice in investor utility. In fact, a neutral long-short portfolio will maximize the investor's minimally constrained utility function only under the very limited conditions discussed below.

Dollar-Neutral Portfolios. We consider first the conditions under which a dollar-neutral portfolio maximizes the minimally constrained utility function. By definition, the risky portfolio is dollar-neutral if the net holding H of risky securities is zero, meaning that

$$H = \sum_{i=1}^{N} h_i = 0 \tag{5}$$

This condition is independent of h_R, the fraction of wealth held in the risky portfolio. Applying the condition given in Equation (5) to the optimal weights from Equation (4), together with a simplifying assumption regarding the covariance matrix, it can be shown that the dollar-neutral portfolio is equal to the minimally constrained optimal portfolio when:[8]

$$H \propto \sum_{i=1}^{N} (\xi_i - \bar{\xi}) \frac{r_i}{\sigma_i} = 0 \tag{6}$$

where σ_i is the standard deviation of the return of stock i, $\xi_i = 1/\sigma_i$ is a measure of the stability of the return of stock i, and $\bar{\xi}$ is the average return stability of all stocks in the investor's universe. The term r_i/σ_i is a risk-adjusted return, and the term $\xi_i - \bar{\xi}$ can be regarded as an excess stability, or a stability weighting. Highly volatile stocks will have low stabilities, so their excess stabilities will be negative. Conversely, low-volatility stocks will have high stabilities, so their excess stabilities will be positive.

The condition shown in Equation (6) states that the optimal net holding of risky securities is proportional to the universe's net stability-weighted risk-adjusted expected return. If this quantity is positive, the net holding should be long; if it is negative, the net holding should be short. The optimal risky portfolio will be dollar-neutral only under the relatively unlikely condition that this quantity is zero.

Beta-Neutral Portfolios. We next consider the conditions under which a beta-neutral portfolio maximizes the minimally constrained utility function. Once the investor has chosen a benchmark, each security can be modeled in terms of its expected excess return α_i and its beta β_i with respect to that benchmark. Specifically, if r_B is the expected return of the benchmark, then the expected return of the i^{th} security is

$$r_i = \alpha_i + \beta_i r_B \tag{7}$$

The expected return of the portfolio can be modeled in terms of its expected excess return α_P and beta β_P with respect to the benchmark

$$r_P = \alpha_P + \beta_P r_B \tag{8}$$

where the beta of the portfolio is expressed as a linear combination of the betas of the individual securities, as follows:

$$\beta_P = \sum_{i=1}^{N} h_i \beta_i \tag{9}$$

From Equation (8), it is clear that any portfolio that is insensitive to changes in the expected benchmark return must satisfy the condition

$$\beta_P = 0 \tag{10}$$

Applying the condition given in Equation (10) to the optimal weights from Equation (4), together with the model given in Equation (7), it can be shown that the beta-neutral portfolio is equal to the optimal minimally constrained portfolio when:

$$\sum_{i=1}^{N} \frac{\beta_i r_i}{\omega_i^2} = 0 \tag{11}$$

where ω_i^2 is the variance of the excess return of security i.

Equation (11) describes the condition that a universe of securities must satisfy in order for an optimal portfolio constructed from that universe to be unaffected by the return of the chosen benchmark. The summation in Equation (11) can be regarded as the portfolio's net beta-weighted risk-adjusted expected return. Only under the relatively unlikely condition that this quantity is zero will the optimal portfolio be beta-neutral.

Optimal Equitization

Using various benchmark return vectors, one can construct an orthogonal basis for a portfolio's returns.[9] The portfolio can then be characterized as a sum of components along (or exposures to) the orthogonal basis vectors.

Consider a two-dimensional decomposition. The expected return of the chosen benchmark can be used as the first basis vector and an orthogonalized cash return as the second. The expected return of a beta-neutral portfolio is independent of the returns of the chosen benchmark. That is, its returns are orthogonal to the returns of the benchmark, and can therefore be treated as being equivalent to an orthogonalized cash component. In this sense, the beta-neutral portfolio appears to belong to a completely different asset class from the benchmark. It can be "transported" to the benchmark asset class by using a derivatives overlay, however.

A long-short portfolio can be constructed to be close to orthogonal to a benchmark from any asset class, and can be transported to any other asset class by use of appropriate derivatives overlays. But because long-short portfolios comprise existing underlying securities, they inhabit the same vector space as existing asset classes; they do not constitute a separate asset class in the sense of adding a new dimension to the existing asset class vector space.

Some practitioners nevertheless treat long-short portfolios as though they represent a separate asset class. They do this, for example, when they combine an optimal neutral long-short portfolio with a separately optimized long-only portfolio so as to optimize return and risk relative to a chosen benchmark. The long-only portfolio is in effect used as a surrogate benchmark to transport the neutral long-short portfolio toward the desired risk and return profile.

Although unlikely, it is possible that the resulting combined portfolio can optimize the investor's original utility function. It can do so, however, only if the portfolio h that maximizes that utility can be constructed from a linear combination of the long-only portfolio and the neutral long-short portfolio.

Specifically, if h_{LO} represents the holdings of the long-only portfolio and h_{NLS} those of the neutral long-short portfolio, the combined portfolio can be optimal if h belongs to the range of the transformation induced by vectors h_{LO} and h_{NLS}; that is, if

$$h \in R[h_{LO}\ h_{NLS}] \tag{12}$$

In general, however, there is nothing forcing the three portfolios to satisfy such a condition.

How, then, should one combine individual securities and a benchmark security to arrive at an optimal portfolio? The answer is straightforward: One includes the benchmark security explicitly in the formulation of the investor's utility function and performs a single integrated optimization to obtain the optimal individual security and benchmark security holdings simultaneously.

Consider the problem of maximizing a long-short portfolio's return with respect to a benchmark while simultaneously controlling for residual risk. The variables that can be controlled in this problem are h and the benchmark holding denoted by h_B. We make the simplifying assumption that benchmark holdings consume no capital. This is approximately true for benchmark derivatives, such as futures and swaps. The portfolio's expected excess return is thus

$$r_E = r_F + \sum_{i=1}^{N} h_i r_i + h_B r_B - r_B \tag{13}$$

It can be shown (see Jacobs, Levy, and Starer [1998]) that the optimal risky portfolio h in this case is:

$$h = (\phi + m\psi)\tau$$

where $\phi = Q^{-1}r$ is the standard portfolio that would be chosen by an idealized investor with unit risk tolerance who optimizes Equation (1) without any constraints; $\psi = Q^{-1}q$ is a minimum-residual risk (MRR) portfolio; $q = \text{cov}(r, r_B)$ is a vector of covariances between the risky securities' returns and the benchmark return; and m is the ratio of the expected excess return of the MRR portfolio to the variance of that return.

Clearly, as the expected excess return to the MRR portfolio increases, or the variance of that return decreases, the ratio m increases, and a larger proportion of the risky portfolio should be assigned to the MRR portfolio. Conversely, as m decreases, more of the risky portfolio should be assigned to the standard portfolio ϕ. As the investor's risk tolerance increases, the amount of wealth assigned to the risky portfolio increases.

The exposure to the benchmark that maximizes the investor's utility is

$$h_B = 1 - m\tau$$

This exposure decreases as the MRR portfolio becomes more attractive and as the investor's risk tolerance increases. The exposure may be negative, under which condition the investor sells the benchmark security short. Conversely, as the investor's risk tolerance or the attractiveness of the MRR portfolio decreases, the benchmark exposure should increase.

In the limit, as either m or τ tends toward zero, the optimal benchmark exposure reaches 100% of the invested wealth. An optimally equitized portfolio, however, will generally not include a full exposure to the benchmark security. In the limit, as m approaches zero (and h_B approaches one), the risky portfolio h becomes proportional to the standard portfolio; for this risky portfolio to be optimally beta- or dollar-neutral, the same conditions must be satisfied as those given in Equations (6) and (11) for the unequitized portfolio defined by Equation (4).

The risky part of the equitized portfolio is optimally dollar-neutral when

$$\sum_{i=1}^{N} (\xi_i - \bar{\xi}) \frac{r_i + mq_i}{\sigma_i} = 0 \tag{14}$$

The term on the left-hand side of Equation (14) can be interpreted as a net stability-weighted risk-adjusted expected return. The risky part of the optimally equitized portfolio should be net long if this quantity is positive and net short if it is negative. This is analogous to the condition given in Equation (6) for an unequitized long-short portfolio. The equitized case includes an additional term, mq_i, that captures the attractiveness of the MRR portfolio and the correlations between the risky securities' and the benchmark's returns.

Similarly, the risky part of the optimally equitized portfolio is beta-neutral when

$$\sum_{i=1}^{N} \frac{\beta_i}{\omega_i^2}(r_i + mq_i) = 0$$

This is analogous to the condition given in Equation (11) for an unequitized portfolio. Again, the condition for the equitized portfolio to be beta-neutral includes the additional term mq_i.

CONCLUSION

The freedom to sell stocks short allows the investor to benefit from stocks with negative expected returns as well as from those with positive expected returns. The advantages of combining long and short portfolio positions, however, depend critically on the way the portfolio is constructed. Traditionally, long-short portfolios have been run as two-portfolio strategies, where a short-only portfolio is added to a long-only portfolio. This is suboptimal compared with an integrated, single-portfolio approach that considers the expected returns, risks, and correlations of all securities simultaneously. Such an approach maximizes the investor's ability to trade off risk and return for the best possible performance.

Also generally suboptimal are construction approaches that constrain the short and long positions of the portfolio to be dollar- or beta-neutral. Only under very limited conditions will such a constrained portfolio provide the same utility as an unconstrained portfolio. In general, rather than using long-short balance to achieve a desired exposure (including no exposure at all) to a particular benchmark, investors will be better off considering benchmark exposure as an explicit element of their utility functions.

Long-short management is often perceived as substantially riskier or costlier than long-only management. Much of any incremental cost or risk, however, reflects either the long-short portfolio's degree of leverage or its degree of "activeness"; both of these parameters are under the explicit control of the investor. Additionally, proper optimization ensures that expected returns compensate the investor for risks incurred.

Given the added flexibility that a long-short portfolio affords the investor, it can be expected to perform better than a long-only portfolio based on the same set of insights.

ENDNOTES

The authors thank Clarence C.Y. Kwan for helpful comments, and Judith Kimball for editorial assistance.

[1]As the median-capitalization stock in the Russell 3000 index has a weighting of 0.01%.

[2]The ability to short will be particularly valuable in a market in which short-selling is restricted and investment opinion diverse. When investors hold diverse opinions, some will be more pessimistic than others. With short-selling restricted, however, this pessimism will not be fully reflected in security prices. In such a world, there are likely to be more profitable opportunities for selling overpriced stocks short than there are profitable opportunities for purchasing underpriced stock. See Miller [1977].

[3]This assumes symmetry of inefficiencies across attractive and unattractive stocks. It also assumes that portfolio construction proceeds identically and separately for the long and short sides as it does in long-only portfolio construction. Although these assumptions may appear unduly restrictive, they have often been invoked. See Jacobs, Levy, and Starer [1998] for a discussion of this literature and our counterpoints.

[4]In deriving the formula, it is assumed that the beta of the short side equals the beta of the long side.

[5]Furthermore, under Regulation T, a long-only portfolio can engage in leverage to the same extent as a long-short portfolio. Long-short has an advantage here, however, because purchasing stock on margin can give rise to a tax liability for tax-exempt investors. According to Internal Revenue Service Ruling 95-8, borrowing shares to initiate short sales does not constitute debt financing, so any profits realized when short positions are closed out do not give rise to unrelated business taxable income.

[6]For analytical tractability and expositional simplicity, we use the traditional mean-variance utility function, although it is only a single-period formulation and is not sensitive to investor wealth. Also, behavioral research may question the use of an analytic utility function in the presence of apparently irrational investor behavior. Nevertheless, we believe our conclusions hold for more elaborate descriptions of investor behavior.

[7]A dollar balance may appear to provide tangible proof of the market neutrality of the portfolio. But unless a dollar-balanced portfolio is also beta-balanced, it is not market-neutral.

[8]The simplifying assumption applied is the constant-correlation model of Elton, Gruber, and Padberg [1976].

[9]One could, for example, use the Gram-Schmidt procedure (see Strang [1988]).

REFERENCES

Elton, Edwin J., Martin J. Gruber, and Manfred W. Padberg. "Simple Criteria for Optimal Portfolio Selection." *Journal of Finance*, December 1976, pp. 1341-1357.

Jacobs, Bruce I. "Controlled Risk Strategies." In *ICFA Continuing Education: Alternative Assets*. Charlottesville, VA: Association for Investment Management and Research, 1998.

Jacobs, Bruce I., Kenneth N. Levy, and David Starer. "On the Optimality of Long-Short Strategies." *Financial Analysts Journal*, March/April 1998.

Markowitz, Harry. "Portfolio Selection." *Journal of Finance*, March 1952, pp. 77-91.

Miller, Edward M. "Risk, Uncertainty, and Divergence of Opinion." *Journal of Finance*, September 1977, pp. 1151-1168.

Sharpe, William F. "Capital Asset Prices with and without Negative Holdings." *Journal of Finance*, June 1991, pp. 489-509.

Strang, Gilbert. *Linear Algebra and Its Applications*, 3rd ed. New York: Harcourt Brace Jovanovich, 1988.