The Efficiency Gains of Long-Short Credit Strategies

FREDERICK E. DOPFEL AND SUNDER R. RAMKUMAR

FREDERICK E. DOPFEL is managing director with the Client Advisory Group, Barclays Global Investors, San Francisco, CA.
fred.dopfel@barclaysglobal.com

SUNDER R. RAMKUMAR is an associate with the Client Advisory Group, Barclays Global Investors, San Francisco, CA.
sunder.ramkumar@barclaysglobal.com

Most fixed-income investors constrain their managers' portfolios to be long-only instead of allowing for the shorting of securities. The practical impact of this restriction is that not all of the signal information generated by the managers' research processes can be used, and the resulting portfolio is suboptimal compared with an unconstrained portfolio. This observation has been made for the equity asset class and, in principle, for active management of any asset class. However, although a material impact has been found for equity portfolios, there has been no estimate of the potential impact for fixed-income portfolios. This issue is even more important today because of the evolution of the fixed-income marketplace. Recently, electronic trading, greater price transparency, and especially the emergence of the credit derivative market have enabled managers to more easily express negative views on individual credit securities. This development has created an opportunity to exploit the efficiency gains of long-short strategies for fixed income.

Measurement of the impact of the long-only constraint begins with an understanding of the performance of an unconstrained portfolio measured by expected return above the benchmark (called active return or alpha) and expected risk relative to the benchmark (called active risk or tracking error). To be meaningful, there must be a presumption of the manager's skill, measured by the information coefficient, and we must take into account the economic environment for fixed income, measured by the breadth of securities in the benchmark, security-specific risk, and the risk budgets of typical portfolios. Only then is it possible to consider the expected return and risk relationships of the unconstrained optimized portfolio and compare it with a portfolio that is optimized but constrained to be long only. This methodology has been developed and applied for equity portfolios by Grinold and Kahn [2000] and later expanded to more generalized constraints by Clarke, deSilva, and Thorley [2002]. For equities, the impact of the long-only constraint is estimated to be a 40–60% reduction in expected alpha for a typical active strategy compared with a long-short strategy. There is no *a priori* reason to believe that the impact of the long-only constraint should be more or less significant for fixed-income credit portfolios. It is a question to be settled by extending the current framework to include the key factors for fixed income that drive the impact of the long-only constraint.

We focus our analysis on security selection strategies, rather than tactical bets, as the long-only constraint is most binding on security selection at typical risk budgets. Tactical bets on industry and quality sectors are much less affected by the constraint since the average benchmark weights of the industry and quality sectors are fairly large compared with the weighting of an individual security within a benchmark of several hundred securities. For example, the high-quality and the telecommunications industry sectors represent approximately 20% and 8%, respectively, of the investment-grade credit benchmark, whereas an individual security represents less than 1% on average. Thus, tactical bets can be very long or short of benchmark and consume a material portion of the risk budget while still preserving a net long position (and not violating the long-only constraint). Therefore, the advantage of long-short is most critically associated with security selection strategies rather than tactical bets.

This article begins with a review of the framework and driving factors that determine the relative performance of constrained and unconstrained portfolios in three asset classes—equities as well as investment-grade and high-yield credit portfolios. The information for equities is included for purposes of comparison; the economic characteristics for investment-grade and high-yield credit are somewhat different, but they are more similar to each other than to those for equities. Next, the efficiency gains for long-short portfolios are estimated for the three asset classes and compared side by side, including sensitivity tests. Finally, the implications of these findings for the overall efficiency of fixed-income portfolios is examined with a comment on how investors can best exploit the efficiency gains of long-short investing within the fixed-income asset class.

FRAMEWORK AND KEY FACTORS

The return of any security or portfolio can be separated into two components: a systematic component that is correlated with the benchmark index returns and the residual return. Active management is concerned with the residual exposures that portfolio managers control by overweighting or underweighting securities relative to the benchmark weights. The manager's goal is to generate outperformance, or alpha, relative to the underlying benchmark while controlling the active risk or tracking error that results from these active positions. The residual exposures can be separated further into industry-sector, credit-rating, and duration factors as distinguished from pure security-specific exposures. The security-specific exposures, forming the major potential value-added from managers, may be significantly limited by the long-only constraint.

The long-only constraint is important since it directly limits the degree to which managers can underweight securities and indirectly limits the degree to which they can overweight securities. If a manager had a negative view on a security, he would underweight the security relative to the benchmark, but he is ultimately limited in expressing this view by an inability to go short. Furthermore, since overweights must be funded by corresponding underweights, the manager is limited in expressing positive views with large overweights. As a result, the manager cannot implement his entire signal information. We measure this loss in information by an *efficiency measure*—the ratio of expected alpha from a long-only portfolio to the expected alpha from a portfolio constructed with the same information at the same risk budget but without this constraint. The efficiency measure is derived within the framework of portfolio optimization in Appendix A, following an approach similar to that of Grinold and Kahn [2000], and is equivalent to the *transfer coefficient* defined in Clarke, de Silva, and Thorley [2002]. The principal empirical factors that are found to determine the loss in efficiency from the long-only constraint (or, stated positively, the efficiency gains from long-short investing) are

1. benchmark breadth and concentration
2. average security-specific risk
3. portfolio active risk budget

Although the analytical framework is the same, these factors may vary materially between equity portfolios and fixed-income portfolios, and within fixed-income portfolios between investment-grade and high-yield credit. The following section compares typical values for the driving factors across the three asset classes.

Benchmark Breadth and Concentration

Two characteristics of the benchmark distribution affect long-only efficiency: breadth and degree of concentration. Breadth in this context refers to the annual number of independent bets available in the benchmark, or the *maximum breadth potential*. In theory, a portfolio manager could change weightings on several thousand bonds comprising a fixed-income benchmark multiple times each year. However, bonds issued by the same source and actions taken more than once a year are likely to be highly correlated and thus are not independent bets. Whereas the residual exposures of same-issuer securities typically show a material correlation even after adjusting out industry-sector, credit-rating, and duration factors, the correlation of specific returns between different issuers is usually negligible. For this reason, the measure of breadth is approximated by the number of issuers, assuming annual forecasts, and throughout the text we use the terms "issuer" and "security" interchangeably.

As the breadth of securities (issuers) in the benchmark increases, the average weight of securities decreases. Lower weights decrease the size of permissible negative active positions before violation of the long-only constraint occurs. Thus, negative signals on securities cannot be fully used, thereby sacrificing expected alpha and reducing efficiency. Aside from breadth, benchmarks for which the market capitalization is concentrated among a few of the largest securities are most severely affected by the long-only constraint. This concentration implies very small benchmark weights for the large number of remaining securities, reducing the size of negative active positions that can be expressed on a significant proportion of the securities, in turn reducing long-only portfolio efficiency.

The most commonly used benchmarks for U.S. fixed-income portfolios are the Lehman Brothers Aggregate and the Lehman Brothers Universal. The investment-grade credit and high-yield components of these benchmarks are used as the credit benchmarks in this study and compared with the S&P 500 equity benchmark. The breadths of the three benchmarks are all of the same order of magnitude. The S&P 500 has 500 issuers, and the Lehman Credit and Lehman High Yield have approximately 600 and 800 issuers, respectively. The concentration of the benchmark distributions can be illustrated by Lorenz curves, as in Exhibit 1. The Lorenz curve shows the cumulative percentage of market capitalization of the benchmark by the cumulative percentage of the securities, ranked in descending order of market capitalization. A diagonal Lorenz curve indicates the lowest possible benchmark concentration (equally weighted), and Lorenz curves with higher initial steepness indicate a higher concentration. The Lorenz curve for equity is straddled by the slightly more concentrated investment-grade credit curve and the slightly less concentrated high-yield credit curve. However, the concentrations of securities within the three benchmarks appear fairly similar.

Average Security-Specific Risk

Security-specific risk is defined as the volatility (standard deviation) of the security return that is not explained by common systematic factors such as industry or quality sectors. The average specific risk of the securities determines the size of the active positions required to meet the desired portfolio active risk budget. For benchmarks with lower average specific risk, larger average active positions are required for the same risk budget. But since short positions are prohibited by the long-only constraint, managers cannot use their negative signals fully to generate alpha, which lowers efficiency. Equities have the highest median historical specific risk (25%), followed by high-yield (12%) and then investment-grade credit (3%), based on 10-year median numbers.

Security-specific risks for equities were estimated using Barra's USE3 model and were measured relative (residual) to the 13 risk factors and 55 industry factors used by Barra. On the other hand, fixed-income security-specific risk was measured residual to a standard set of common factors including 5 interest rate factors, 37 corporate sector factors, and the equity factors from Barra's USE3 model. Forecasts of security-specific risk were obtained using a parsimonious model based on historical specific risks and bond characteristics. This was calibrated using over 10 years of historical bond data obtained from Merrill Lynch, EJV, CSFB, and JP Morgan.

Exhibit 2 depicts how the estimated specific risk of securities varies according to the relative size of securities within each asset class as of June 30, 2004. In this sample, equities had an average specific risk of 22.7%, compared with 9.3% for high-yield credit and 2.5% for investment-grade credit. These risks are lower than historical averages, reflecting a low-volatility environment during that period. For each benchmark, the security-specific risk was regressed against the percentile size rank of the securities in the benchmark, ordered by descending market

EXHIBIT 1

Concentration of Benchmark Distributions (Lorenz Curves)

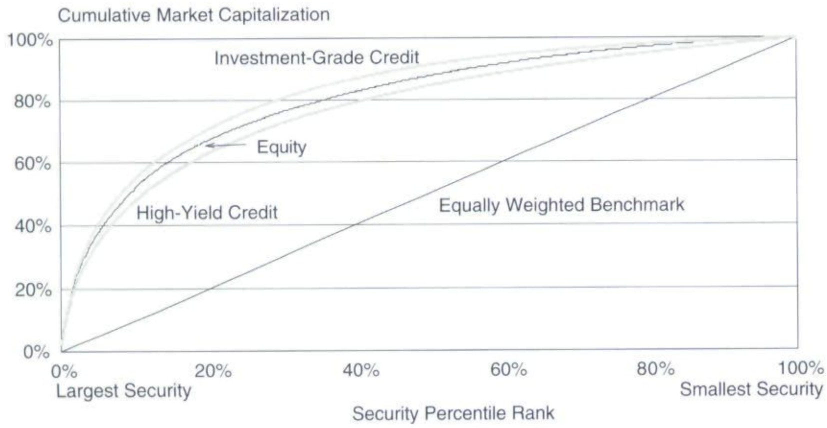

EXHIBIT 2
Benchmark Specific Risk Distributions (Linear Fit)

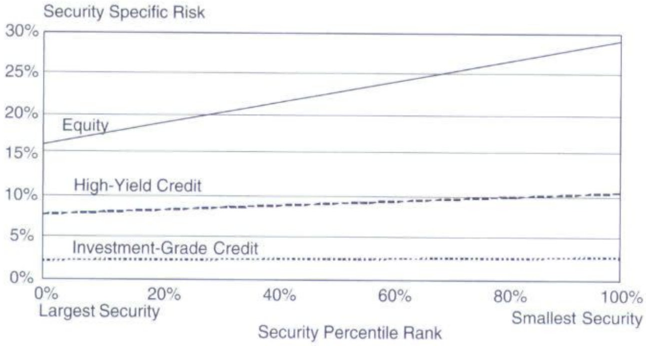

capitalization. The resulting linear fit highlights the much higher specific risk and size sensitivity for equities compared with high-yield credit and investment-grade credit. The relationship for equities supports the finding by Grinold and Kahn [2000], who calculated the long-run correlation between size (measured by the natural log of market cap) and specific risk to be approximately −0.55. For both investment-grade bonds and high-yield bonds, this correlation is less significant, at approximately −0.2, as spreads are more important than size in determining a bond's specific risk. The size distribution of risks has a slight secondary effect in determining the efficiency gains of long-short portfolios, as demonstrated in Appendix B.

Portfolio Active Risk Budgets

The portfolio active risk budget is the desired level of volatility or tracking error in a portfolio relative to its benchmark. Active fixed-income portfolios are typically run at lower active risk budgets than equity portfolios. Active equity managers may have a portfolio active risk budget of approximately 6.00%, though small-cap risk budgets tend to be higher and there are risk-controlled managers running at lower risk budgets. Credit managers have substantially lower portfolio active risk budgets than equity managers. A typical portfolio budget for investment-grade credit portfolios is 0.75% active risk, though some risk-controlled managers have even lower risk budgets, and managers who engage in low-breadth tactical bets (e.g., including views on the movement of the yield curve or macro-sector spreads) may have active risk budgets of more than 1.00%. High-yield managers usually have higher portfolio active risk budgets than investment-grade credit managers. A typical budget for high-yield credit portfolios is 2.50% active risk.

As the desired portfolio active risk budget increases, larger active positions are required to meet the portfolio risk target. For example, to get to a risk budget that is twice as high, a manager would want to double the active positions relative to benchmark weights. However, doubling one's bets would imply more frequent short positions, which are not permitted by the long-only constraint, resulting in a loss of expected alpha. Thus, the efficiency gains of long-short strategies increase with higher portfolio active risk budgets but decrease with higher average security-specific risk, as noted previously. In fact, it is the ratio of the average security-specific risk to the portfolio's active risk budget that determines the magnitude of efficiency gains for a given benchmark concentration and breadth. The lower this ratio is, the higher will be the potential efficiency gains of long-short portfolios.

Summary of Factors

Exhibit 3 summarizes the driving factors that determine long-only efficiency for equities, investment-grade credit, and high-yield asset classes. The breadth and concentration (summarized by the Gini coefficient[1]) of the respective benchmarks are similar, as noted previously. Typical credit portfolios have securities with lower specific risks, but this is counterbalanced by the lower typical active risk budgets of credit portfolios, with the ratio of security-specific risk to active risk budgets appearing fairly similar across the benchmarks. Hence, qualitatively, we might expect similar efficiency gains from long-short equity, investment-grade credit, and high yield portfolios. In the next section, the potential efficiencies are quantified.

ASSUMPTIONS AND FINDINGS FOR EFFICIENCY GAINS

In this section, reasonable assumptions for the key factors are set, and an efficiency measure is calculated to estimate the potential gains from long-short credit portfolios. The efficiency measure is calculated as the ratio of expected alpha that may be attained from the long-only portfolio to the expected alpha that may be attained from the unconstrained long-short portfolio. But to measure efficiency gains of long-short portfolios, it is first necessary to assume that portfolios, whether constrained or not, are constructed to optimize the portfolio's expected alpha against the

EXHIBIT 3
Summary of Equity, Investment-Grade Credit, and High-Yield Factors

	Equities	Investment-Grade Credit	High-Yield Credit
Benchmark	S&P500	Lehman Credit	Lehman High Yield
Breadth	500	602	799
Concentration–Gini Coefficient	0.62	0.67	0.56
Average Security-Specific Risk	25.00%	3.00%	12.00%
Portfolio Active Risk Budget	6.00%	0.75%	2.50%
Average Specific Risk/Active Risk Budget	4.2	4.0	4.8

portfolio's active risk budget based on the research and signal information available to the manager. Further, the constrained and unconstrained portfolios should be based on exactly the same information and assumptions.

Assumptions

The following assumptions are used:

1. *Benchmark breadth and concentration.* The actual benchmark distributions as of June 30, 2004, are used, consistent with Exhibits 1 and 3. The available data cover approximately 99% of the market value of the credit indexes.
2. *Security-specific risks.* The specific risks and correlations between issuers' securities, combined with holdings, determine the overall active risk of any proposed portfolio. Although the specific risk of securities varies, a constant specific risk is assumed for each asset class: 25% for equities, 3% for investment-grade credit, and 12% for high-yield credit, consistent with Exhibit 3. These represent (rounded) median specific risks for the issuers' securities in the three benchmarks over the past 10 years. As previously noted, specific return correlations between different issuers are typically negligible after controlling for industry-sector, credit-rating, and duration factors. Therefore, the forward-looking assumption is zero correlation between different issuers' security-specific returns.
3. *Portfolio active risk budgets.* Active risk budgets for the portfolios are allowed to vary from zero to a typical manager's active risk budget and then to include portfolio active risk budgets that would be considered higher than average.
4. *Expected security-specific return.* The expected alpha contributed by an individual security holding is assumed to be proportional to the security's specific risk and a randomly generated score that reflects the manager's view, as in Grinold [1994]. The score expresses a relative attractiveness ranking of the security and is simulated by a normal distribution with a zero mean and unit variance. This assumption is reasonable because, although *total* returns exhibit negative skewness, after removing industry-sector, credit-rating, and duration factors, the *security-specific* returns no longer show negative skewness. Using a zero mean ensures that the average cross-sectional alpha is also zero, consistent with the zero-sum nature of active returns. Further, assigning the signal a unit variance controls the scale of the alphas. The coefficient that scales the expected returns reflects the skill level of the forecaster, called the information coefficient (IC), and is the correlation between the manager's forecasts and the realizations. A constant positive IC is used across all securities in a benchmark, signifying manager(s) with investment skills that are above average and that do not vary for different sizes of securities within a benchmark. The magnitude of the constant IC does not change the efficiency measure, so long as it is applied consistently to both the constrained and unconstrained scenarios.

For an asset class simulation trial, a random score (signal) was generated for each security in the benchmark. Then, the optimal portfolio holdings were determined at a given portfolio active risk budget based on the set of signals. Optimal holdings for the long-short optimization follow from the first-order conditions for mean-variance optimization, as described in Appendix A. The long-only optimization was calculated under the same set of assumptions but with the constraint of no shorting. The optimal holdings of the long-only portfolio are the solution to an optimization problem with a non-linear objective function and linear inequality constraints, as expressed in Appendix A. This simulation was repeated 100 times for each asset class. Then, the efficiency measure was calculated as the ratio of the average expected alpha from the optimal long-only portfolios and the average expected alpha from the optimal long-short portfolios. This process was repeated across a range of typical manager active risk budgets.

Findings for Efficiency Gains

The resulting efficiency measures are plotted in Exhibit 4 for each asset class and for a range of portfolio active risk budgets. Results for the equities, showing an efficiency measure of approximately 42% at a typical manager risk budget of 6.00%, match the impact of the long-only constraint results in Grinold and Kahn [2000] and the transfer coefficient result for long-only equity in Clarke, de Silva, and Thorley [2002]. For investment-grade credit, the efficiency measure is approximately 38% for a typical risk budget of 0.75%. For high yield, the efficiency measure is 40% for a typical risk budget of 2.50%. Thus, for all strategies, the efficiency gain for long-short strategies is approximately 2.5 times (1/40%) before costs of shorting, confirming the hypothesis in the previous section that the efficiency gains would be fairly similar. This also means that skillful fixed-income managers not restricted by the long-only constraint can obtain much higher expected alpha per unit of active risk than typical fixed-income managers.

Note that the decline in the efficiency measure is steepest for investment-grade credit, followed by high yield and then equities. The incremental expected alpha that long-only fixed-income portfolios can generate with an increase in portfolio active risk is thus severely limited due to their inability to go short. This might explain why investment-grade credit and high-yield managers run portfolios at lower risk budgets than equity managers. In other words, managers across different asset classes may be setting their portfolio active risk budgets to moderate the penalty associated with the long-only constraint.

As expected, the impact of the long-only constraint is seen in the optimal holdings of securities, as shown in Exhibit 5 for an example of an investment-grade credit portfolio. The long-short portfolios allowed the signals to be implemented efficiently with overweights on securities with a positive expected alpha and underweights on those with a negative expected alpha. The magnitude of the active positions was proportional to the magnitude of the expected alpha. In contrast, the corresponding long-only portfolios had smaller underweights, limited by the long-only constraint, and there were large overweights to relatively few securities with the highest forecast alpha. The long-only portfolios also had the larger underweights concentrated in the larger securities (where these could be implemented without violating the constraint), which were then used to fund overweights.

Transaction Costs and Sensitivity Analysis

Our results thus far have been presented gross of transaction costs for simplicity and to build intuition. However, although both long-only and long-short portfolios are subject to transaction costs in the form of bid-ask spreads, long-short portfolios have an additional borrowing cost associated with shorting securities. The cost of shorting equities is approximately 0.25%, whereas the costs of shorting investment-grade and high-yield strategies are 0.30% and 0.40%, respectively, if physical bonds are used. However, with the advent of credit default swaps, managers can now short bonds for only a few basis points. This provides for a cost-efficient implementation of long-short portfolios but with the addition of basis risk and some counterparty risk. Thus typical implementations will involve a combination of physicals and derivatives. We estimate the magnitude of these additional shorting costs for typical portfolios by measuring the amount of shorting associated with the optimal long-short portfolio and the corresponding costs. As shown in Exhibit 6, these borrowing frictions lower the efficiency measures (or, equivalently, reduce the gains of long-short portfolios) by between 2% and 6% for skillful active managers. The performance drag from borrowing costs is greatest for

EXHIBIT 4

Efficiency Measure for Equity, Investment-Grade Credit, and High-Yield Credit Portfolios

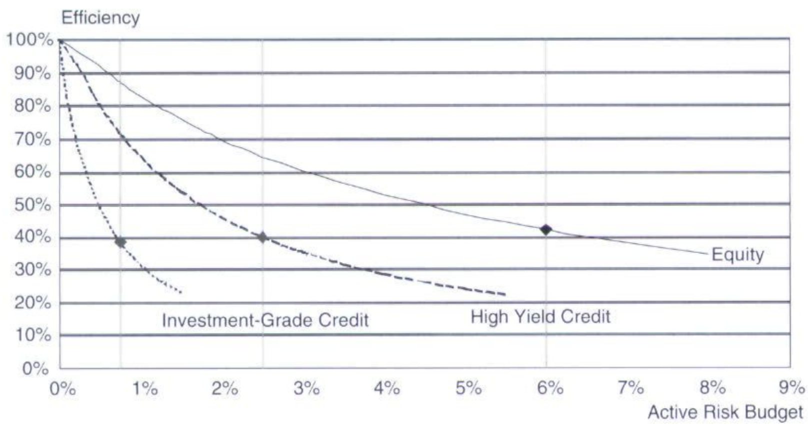

EXHIBIT 5

Comparison of Optimal Holdings for Long-Only and Long-Short Investment-Grade Credit Portfolios

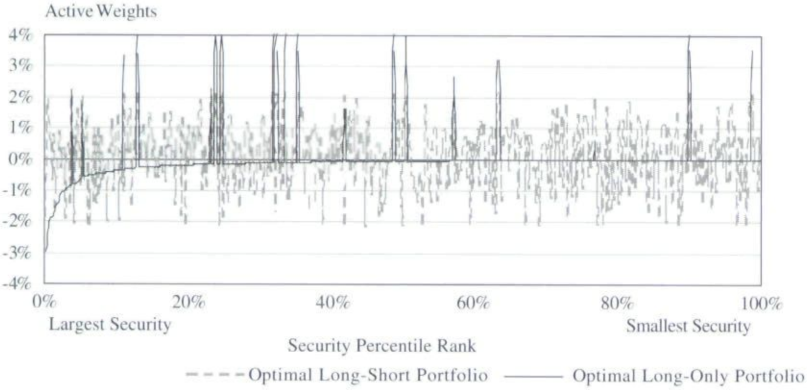

EXHIBIT 6

Impact of Borrowing Costs on Efficiency Measure at Typical Risk Budgets

Efficiency Measure	Equities	Investment-Grade Credit	High Yield Credit
Before Borrowing Costs (Exhibit 4)	42%	38%	40%
After Borrowing Costs	44%	45%	42%

investment-grade credit portfolios since we have assumed that investment-grade credit has higher transaction costs relative to expected portfolio alpha.[2]

Manager data are not publicly available to validate the efficiency measures calculated for credit strategies, and we suspect that we will be waiting for many years more. Further, true efficiency effects would be masked by variation in manager skill and other factors, preventing a precise measurement of the long-only impact from empirical data. In the absence of empirical data, we test our base case results for several sensitivities, documented in Appendix B. Instead of assuming a constant specific risk for all securities (Assumption 1), we run simulations using a recent actual distribution of specific risks. This captures some of the heterogeneity inherent in securities. Next, we test our assumption of a constant IC (Assumption 4) for all securities by allowing changes in the IC by security, which in turn allows the quality of the manager's forecast to vary across securities. Both sensitivity tests are described in detail in Appendix B and yield efficiency measures that are similar to the base case results at typical risk budgets.

IMPLICATIONS FOR FIXED-INCOME PORTFOLIOS

An important practical question is how the efficiency gains of long-short investing can be used to improve the actual performance of fixed-income portfolios while maintaining best practices for risk control. This question can be viewed as one of manager structure optimization.[3] The objective is to maximize expected alpha for the entire portfolio (the overall expected return in excess of the benchmark) against the active risk (or tracking error) relative to the benchmark. A long-short manager with a positive expected alpha may be considered one of many active managers who may potentially contribute to portfolio performance, but there are a few additional considerations that distinguish long-short strategies—benchmarks, risk, and expected return assumptions.

In the United States, the most common benchmark for fixed-income portfolios is the Lehman Brothers Aggregate Bond Index, and many managers have this as a stated benchmark or use a custom benchmark that also includes high-yield and global securities. In contrast, most long-short strategies have a stated benchmark of cash, and this disparity in benchmarks causes any allocation to long-short strategies to create tracking error, even before consideration of the long-short manager's own specific risk. To remedy this, the return stream of long-short strategies must be adapted to more closely resemble the desired portfolio benchmark, a process sometimes called *bondization*. In practice, swaps or a combination of futures that approximates the benchmark may be used as a remedy. The cost of adjusting the systematic exposures should be a trade-off against the benefit of reducing tracking error. This could be a minor issue in the case of small allocations to long-short strategies but becomes more important for larger allocations.

To manage fixed-income portfolios, it is essential to establish meaningful forward-looking assumptions for manager performance relative to the benchmark. Rather than targets or historical performance, these expectations should be set with an understanding that active management is difficult, and, but for the presence of skill, expected alpha is zero for *both* long-only and long-short strategies. However, assuming all things equal, including an equal level of skill, analysis has shown that the expected information ratio for a long-short strategy exceeds that for a typical long-only strategy by about 2.5 times (before costs of shorting) at average portfolio active risk levels. For example, we may believe that a long-only manager with upper-quartile skill has an ability to produce an expected information ratio of 0.5, which is typical of upper-quartile performance over the long run. Taking another manager for whom we have the same very favorable assessment of skill, but this time a long-short manager, the expected information ratio should be much higher. This suggests that in forming our assumptions for manager structure we ought to explicitly take into account the efficiency gains of long-short investing. At the same time, we cannot automatically assume that a long-short strategy will produce a positive expected alpha without first establishing the manager's above-average skill level. Successful active management requires managers who have the skill to produce alpha as well as investors who have the skill to select good managers.

If we are truly exploiting the efficiencies of long-short investing, it is likely we are considering strategies with above-average risk. This poses a concern that incorporating such a strategy into a risk-controlled portfolio of managers would result in volatility beyond the norm. To address this question, consider an example of a portfolio with several managers who, in combination, produce a portfolio tracking error of 0.75%. Now consider a long-short strategy with high, 10% active risk. Assuming the strategies are uncorrelated, then a (bondized) 5% allocation to this strategy increases tracking by only 0.12% to 0.87%, still well within the range of tracking error experienced by typical institutional portfolios.[4] Rather than increasing the portfolio risk budget, it is usually possible to maintain the original risk budget by adjusting the mix of managers in the original portfolio, depending on the size of the allocation to the higher-risk long-short strategy. Ultimately, the size of the allocation to the long-short strategy depends on the expected alpha relative to the additional risk it contributes, with the terms of the trade-off assessed consistently with the investor's risk preferences.

Perhaps the most profound implication of long-short strategies for fixed-income investing is how this affects our thinking about the role of the credit portfolio within the total portfolio. The returns derived from long-short credit may be thought of as a source of alpha separate from the systematic returns of traditional strategies, that is, "portable alpha." This enables investors to better engineer fixed-income portfolios that combine cheaper sources of systematic exposures with more efficient sources of active returns.

CONCLUSION

The same factors that drive the efficiency of long-short equity investing relative to long-only strategies—benchmark characteristics, security-specific risk, and portfolio risk budgets—also drive the potential efficiency gains of fixed-income credit long-short strategies. The most common credit benchmarks (both investment grade and high yield) and equity benchmarks have very similar breadth and concentration. Further, although the security-specific risks of credit securities are observed to be much less than for equities, when considered in relation to the typical ranges of their respective portfolio active risk budgets, the relative risk is very similar. In other words, there is little difference in our empirical estimates of the key factors for fixed-income credit compared with equity portfolios. So it is not surprising that the estimated net impact of the long-only constraint for fixed-income credit portfolios is very similar to what has been previously estimated for equity portfolios.

These results apply most directly to comparisons of portfolios in which the active risk budget is derived primarily from pure selection bets as opposed to style biases and tactical bets, for which the long-only constraint is less material. As most fixed-income managers and most investors have broader benchmarks than credit and consume a material proportion of their risk budget on style biases and tactical bets, the overall impact of the long-only constraint would be diluted. However, the greatest prospect for exploiting the efficiency gains of long-short credit strategies is to view them as an alpha overlay on the broader portfolio, with appropriate allocations to manage risk, and with adjustments so as to leave the overall strategic asset allocation policy unchanged. As for any active strategy, the overriding consideration is the assessment of the manager's skill in producing a positive expected alpha, lest the potential efficiency gains of long-short credit investing cut in the opposite direction as more efficient losses.

APPENDIX A
Derivation of Key Factors

Active portfolios, whether long-only or long-short, are assumed to be mean-variance efficient. The mean-variance-efficient portfolio can be described by the optimization of the following objective function:

$$U = \alpha_p - \lambda\omega_p^2 = \sum(h_i - h_{bi})\alpha_i - \lambda\sum(h_i - h_{bi})^2\omega_i^2 \tag{A-1}$$

where α_p is the portfolio expected alpha, ω_p is the portfolio active risk, α_i and ω_i are the expected specific return and volatility of the *i*th issuer's securities (here, simply called the "*i*th security"), λ is a risk-aversion coefficient, h_i is the portfolio weight of the *i*th security, h_{bi} is the benchmark weight of the *i*th security and the summation is over N (issuers) securities in the benchmark. The specific returns are assumed to be uncorrelated. The optimization of (A-1) is subject to $\Sigma h_i = 1$, the full investment condition, and, in the case of long-only portfolios, the constrained optimization of (A-1) is also subject to $h_i \geq 0$, the long-only constraint. Following Grinold and Kahn [2000], the first-order conditions for the unconstrained optimization of (A-1) require

$$h_i - h_{bi} = \frac{\alpha_i}{2\lambda\omega_i^2} \quad \text{(A-2)}$$

According to Grinold [1994], the expected specific return of an individual security is assumed to be proportional to the security's specific risk and a randomly generated score z as follows:

$$\alpha_i = IC \cdot \omega_i \cdot z_i \quad \text{(A-3)}$$

where IC is the information coefficient, the correlation between the manager's forecasts and the realizations. The IC is assumed here to be constant for all securities in an asset class. The distribution of z is assumed to be standard unit normal. Combining (A-2) and (A-3),

$$h_i - h_{bi} = \frac{\mathrm{IC} \cdot z_i}{2\lambda\omega_i} \quad \text{(A-4)}$$

But, as $\omega_p^2 = \sum (h_i - h_{bi})^2 \omega_i^2$ it follows from (A-4) that $\omega_p \approx \mathrm{IC}\sqrt{N}/2\lambda = \mathrm{IR}/2\lambda$ as the information ratio (IR) of a portfolio can be expressed as $\mathrm{IR} = \mathrm{IC}\sqrt{N}$ from the fundamental law of active management in Grinold [1989]. Then, by substitution,

$$h_i - h_{bi} = \frac{\omega_p}{\omega_i} z_i / \sqrt{N} \quad \text{(A-5)}$$

Equation (A-5) is used to determine the unconstrained, long-short optimal holdings and associated alpha. For portfolios subject to the long-only case, the constraint is binding if $h_i < 0$, or

$$z_i < -h_{bi}\frac{\omega_i}{\omega_p}\sqrt{N} \quad \text{(A-6)}$$

Interpreting (A-6), we see that the long-only constraint is binding whenever the score for a particular security falls below this minimum value. This is more likely to happen at higher risk budgets and for lower-breadth (narrow) benchmarks. It is also more likely to be binding for securities that have lower benchmark weights h_{bi} or lower specific risk ω_i. For such securities, the portfolio manager is unable to take advantage of strong negative signals by taking on corresponding large negative active positions. The largest securities usually have the lowest specific risks, making the constraint somewhat equally binding on all the securities. If so, then $h_{bi}\omega_i \approx \overline{\omega}/N$ for the "average security," where $\overline{\omega}$ is the average specific risk per security. Then, the probability that an average security is constrained may be estimated as $\Phi(-(\overline{\omega}/\omega_p)/\sqrt{N})$, where Φ is the standard cumulative normal distribution. The estimate provides a rough approximation of efficiency but is biased upward. In addition to constraining large negative active positions, the long-only constraint also reduces the ability to utilize positive signal information completely since positive active positions need to be funded by corresponding negative active positions.

APPENDIX B
Sensitivity Analysis

Distribution of Security-Specific Risks

The base case model assumed a constant security-specific risk for each security equal to the long-run average security-specific risk, as shown in Exhibit 3. However, as noted in the text, there may be a variation in specific risk over time and according to security size, with securities that have a smaller weight in the benchmark having higher specific risk. Exhibit 2 showed a sample of specific risks according to size rank in the benchmark as of June 30, 2004, an environment that was relatively less volatile than average. The sensitivity of the efficiency measure was tested by changing the specific risk from constant values used in the base model to the actual distribution of security-specific risks as of that date.

Exhibit A1 shows how the efficiency measure for investment grade and high yield changes based on this test at typical portfolio risk budgets. The results indicate similar, but slightly higher, efficiency measures at typical ranges of risk budgets from using the actual distribution. Higher efficiency measures result because the optimal unconstrained active positions in the smaller securities are less extreme due to their higher risk, and therefore less likely

EXHIBIT A1
Sensitivity of Efficiency Measure to Specific Risk Distribution

	Investment-Grade Credit Efficiency			High Yield Credit Efficiency	
Portfolio Risk Budget	Constant Specific Risk Assumption	Variable Specific Risk Assumption	Portfolio Risk Budget	Constant Specific Risk Assumption	Variable Specific Risk Assumption
0.25%	62%	61%	0.50%	75%	72%
0.50%	45%	45%	1.50%	47%	48%
0.75%	*34%*	*37%*	*2.50%*	*33%*	*38%*
1.00%	27%	32%	3.50%	26%	32%
1.25%	23%	28%	4.50%	21%	28%

EXHIBIT A2
Information Coefficient Distribution

Investment-Grade Credit			High Yield Credit		
Market Value Size Quartile	Number of Bonds	Information Coefficient (IC)	Market Value Size Quartile	Number of Bonds	Information Coefficient (IC)
Q1	13	0.015	Q1	21	0.060
Q2	37	0.045	Q2	69	0.050
Q3	94	0.075	Q3	176	0.040
Q4	458	0.060	Q4	533	0.030

to be constrained, and because average security-specific volatility was less than that assumed in the base case.

The benchmark concentration and number of issuers has been fairly stable, although the average security-specific risks of investment-grade and high-yield bonds have varied from 2 to 4% and from 7 to 17%, respectively, over the past 10 years. In general, high-volatility environments produce smaller gains from long-short investing and low-volatility environments produce larger gains from long-short investing compared with the base case.

Distribution of Information Coefficients

The base case model assumed a constant IC for all bonds within a benchmark. In reality, it is possible that an investment manager would be relatively more or less skilled within subgroups of securities based on quality, sector, or other factors related to size. To test sensitivity for a variable IC assumption, the investment-grade and high-yield securities were each divided into quartiles according to cumulative market value (e.g., quartile 1 contains 25% of the total market capitalization of the benchmark), and each quartile was assigned a different information coefficient as shown in Exhibit A2. The tiered IC schedule assigns the hypothetical manager its highest ability to forecast within the middle range of investment-grade bonds and assigns a declining ability to forecast the lower-quartile range of high-yield bonds. The relative IC assumptions were controlled in order to make a fair comparison of constant and variable IC assumptions.

Exhibit A3 shows how the efficiency measure for investment grade and high yield changes based on this test at typical portfolio risk budgets. The results indicate similar efficiency measures for investment-grade credit and slightly higher efficiency measures for high yield based on variable IC measures. A lower IC for the smaller securities would imply lower expected alpha and hence less extreme optimal unconstrained active positions in these securities. The resulting portfolios are less likely to be constrained, resulting in higher efficiency measures. This explains why the predicted efficiency gains are lower for high yield given our assumptions.

EXHIBIT A3
Sensitivity of Efficiency Measure to IC Distribution

	Investment-Grade Credit Efficiency			High Yield Credit Efficiency	
Portfolio Risk Budget	Constant IC Assumption	Variable IC Assumption	Portfolio Risk Budget	Constant IC Assumption	Variable IC Assumption
0.25%	66%	65%	0.50%	80%	86%
0.50%	49%	49%	1.50%	54%	59%
0.75%	*38%*	*40%*	*2.50%*	*40%*	*44%*
1.00%	31%	33%	3.50%	31%	35%
1.25%	26%	28%	4.50%	26%	29%

ENDNOTES

The authors gratefully acknowledge insights provided by Ronald Kahn and contributions by colleagues in the Fixed Income Group at Barclays Global Investors, especially Brian Zalaznick, James Gilliland, Peter Wilson, Shawn Steel, Nicholas Baturin, Jonathan Sadowsky, Victor Wong, and Peter Knez.

[1]The degree of concentration can be summarized by a statistic called the Gini coefficient, which varies between 0 and 1. Higher values of the Gini coefficient indicate greater concentration. It can be calculated as

$$G = \left| 1 - \sum_{i=0}^{N-1} (Y_{i+1} + Y_i)(X_{i+1} - X_i) \right|$$

where Y_i represents the cumulative market cap of the *i*th largest asset and X_i represents the cumulative market cap of the *i*th largest assets if the benchmark were equally weighted.

[2]In practice, investment managers will optimize transaction costs. An optimization of transaction costs was performed on the base case model and showed negligible changes in the efficiency measure for all three asset classes. Performing the optimization on a more complex model could generate a different impact, but the results are difficult to generalize. The (non-optimized) efficiency measure after borrowing costs (Exhibit 6) may be viewed as a conservative upper bound on the efficiency measure, accounting for the level of shorting frictions.

[3]See Grinold and Kahn [2000] for the general foundations of active management and Waring et al. [2000] for the development of manager structure optimization. See Dopfel [2004] for application of these approaches to the fixed-income asset class.

[4]This calculation assumes that the active component of the long-short strategy is uncorrelated to the active component of the initial long-only portfolio. As the risks are orthogonal, the portfolio active risk is calculated $0.87\% = \sqrt{(0.95 \star 0.75\%)^2 + (0.05 \star 10\%)^2}$.

REFERENCES

Clarke, Roger, Harindra de Silva, and Steven Thorley. "Portfolio Constraints and the Fundamental Law of Active Management." *Financial Analysts Journal*, Vol. 58, No. 5 (September/October 2002), pp. 48–66.

Dopfel, Frederick E. "Fixed-Income Style Analysis and Optimal Manager Structure." *Journal of Fixed Income*, Vol. 14, No. 2 (September 2004), pp. 32–43.

Grinold, Richard C. "Alpha is Volatility Times IC Times Score." *The Journal of Portfolio Management*, Vol. 20, No. 4 (1994), pp. 9–16.

——. "The Fundamental Law of Active Management." *The Journal of Portfolio Management*, Vol. 15, No. 3 (1989), pp. 30–37.

Grinold, Richard C., and Kahn, Ronald N. "Efficiency Gains of Long-Short Investing." *Financial Analysts Journal*, Vol. 54, No. 6 (November/December 2000), pp. 40–53.

Waring, M. Barton, Duane Whitney, John Pirone, and Charles Castille. "Optimizing Manager Structure and Budgeting Manager Risk." *The Journal of Portfolio Management*, Vol. 26, No. 3 (Spring 2000), pp. 90–104.