Factor Neutrality

Maintaining it can be suboptimal.

Daniel Y. Lam and Wai Lee

Daniel Y. Lam
is head of Quantitative Research at Third Wave Global Investors, LLC, in Greenwich, CT.
daniel.lam@thirdwaveglobal.com

Wai Lee
is head of Quantitative Investments at Lehman Brothers Asset Management in New York City.
wai.lee@lehman.com

To manage a portfolio against an explicitly specified benchmark, there are good reasons to aim at benchmark neutrality, in the sense of having a target of unity for the portfolio's beta with respect to the benchmark. For instance, if equities are only one portion of a portfolio, a particular equity benchmark is likely to be chosen according to its perceived volatility and correlations with other asset classes. In this case, deviation of the active equity portfolio from a unit beta exposure to the equity benchmark may change the risk profile of the whole portfolio.

Yet implementation and verification of benchmark neutrality is far from easy. Ex post return-based analysis of benchmark neutrality relies on estimation of the average beta exposure to the benchmark during the sample period, which is subject to the dynamic beta problem identified by Grinblatt and Titman [1989]. Ex ante benchmark neutrality depends on the quality of the managers' covariance matrix. While they may be able to remain benchmark-neutral according to estimated betas, there is no guarantee that the ex post beta of the resulting portfolio with respect to the benchmark will be at one.

The issue becomes more complex if one moves beyond one factor to the multifactor world. For quantitative portfolio managers who tilt portfolios toward factors that are predicted to deliver better returns, factor neutrality is clearly suboptimal. In a sense, factor neutrality unwinds most, if not all, of the information advantage gained by such managers.

Also questionable is the practice of factor neutrality by portfolio managers who claim to be focused on stock fundamentals but who do not rely on a purely quantitative approach. Apparently they define stock-picking as predicting the idiosyncratic return components in the total stock returns that are independent of all factors.

It is true that a natural extension of the one-factor framework into a multifactor framework would be to risk-adjust alpha, so that only the return component of the portfolio that is independent of all factors should be perceived as adding value. In this framework, returns resulting from increased exposure to factors would not be viewed as due to investment skill.[1]

While we are somewhat sympathetic to this view, we see major issues that make excluding factors from alpha potentially suboptimal. First, while benchmark neutrality may be supported by a solid investment rationale, we do not see similar arguments concerning other factors. Portfolio managers are typically evaluated against their predetermined benchmarks. A manager is said to have outperformed if portfolio returns are higher than benchmark returns, and if the manager has delivered a positive information ratio, defined as alpha over tracking error. If the manager has superior information concerning the timing of factors, it is entirely legitimate, in our view, to construct the portfolio to have net exposures to some factors relative to the benchmark to capitalize on this advantage.

Second, as mentioned above, many portfolio managers like to use screens to reduce their investment universe from thousands to hundreds before conducting in-depth fundamental analysis of each stock. These screens may include minimum market capitalization, valuation ratios such as price-to-earnings or book value, profitability, past earnings growth, quality, leverage, and recent stock price performance, to name a few. It is highly likely that the criteria these managers use are themselves correlated with common risk factors in stock returns. As a result, a factor-neutral constraint during the portfolio construction stage will likely unwind some of the information advantage embedded in the investment process. The degree to which performance is hurt will vary case by case.

Third, there is little agreement on the number and exact representation of factor models, or as to whether particular factors are priced into securities. It is a common practice in the industry to take a risk model from an outside vendor as a given. Consequently, the impact of factor neutrality will depend on how the vendor defines and models these factors, something that is probably not integrated into the portfolio manager's investment process.

Clarke, de Silva, and Thorley [2002] analyze the impact of constraints on portfolio performance in the context of Grinold's [1989] fundamental law of active management. Our approach differs in that we do not adopt the market model assumption that market beta is the only source of comovement between stocks, but instead allow information on other risk factors to be explicitly incorporated into the investor's portfolio selection. When the investor *knows* that all factors are unpriced and have zero expected returns, we show that factor neutrality can be suboptimal even in this case.

ANALYTICAL FRAMEWORK

We describe an analytical framework to help clarify to what degree investment performance can be affected by factor neutrality under various circumstances. To illustrate, we use hypothetical examples and simulations that are close approximations to real investment situations.

Let R be the random $N \times 1$ vector of excess returns of stocks over the risk-free rate, R_f, the random $K \times 1$ vector of factor returns with covariance matrix F, and ε the random N vector of stock-specific returns with diagonal covariance matrix D. We have:

$$R = Xf + \varepsilon \tag{1}$$

where X is an $N \times K$ matrix of factor exposures.

The covariance matrix of stock returns is then:

$$\Sigma = XFX' + D \tag{2}$$

Assume that the expected excess return of stocks is given by an $N \times 1$ vector μ. Adopting a Bayesian viewpoint, we treat μ as a random variable representing the investor's knowledge about expected returns. The investor selects a portfolio using as expected returns the posterior mean of μ, conditional on all available information.

We assume throughout that the investor chooses a portfolio that is mean-variance efficient in a relative sense, where risk means tracking error against a known benchmark (which may be cash in the case of an absolute-return product). This is what is meant when we say a portfolio is *efficient* or *optimal.*

We need to make some assumptions about the process that generates the investor's private information. The investor may receive information on both factor returns and stock-specific returns:

$$\mu = X\bar{f} + \bar{\varepsilon} \tag{3}$$

where $\bar{f}$ and $\bar{\varepsilon}$ are expected factor returns and expected stock-specific returns, respectively. We assume that, unconditionally, $\bar{f}$ and $\bar{\varepsilon}$ are random variables drawn from multivariate Gaussian distributions:

$$\begin{aligned} \bar{f} &\sim N(0, \tau_f^2 F) \\ \bar{\varepsilon} &\sim N(0, \tau_\varepsilon^2 D) \end{aligned} \tag{4}$$

These distributions summarize the prior information available to the investor before additional conditioning information about factors and stocks. We have made the simplifying assumption that, unconditionally, none of the risk factors is priced and all stock-specific expected returns are zero. Thus, no bets at all would be made in the absence of conditioning information; the investor has no prior bias in favor of or against any factor or any stock.

The parameters τ_f and τ_ε control the typical strength of the conditioning information received on factor returns and stock-specific returns, respectively. They are the typical magnitudes (in a root mean squared sense) of the conditional Sharpe ratios of factors and of stock-specific returns, respectively.

Suppose that conditioning information arrives in the form of a (possibly noisy) observation of the expected returns of ℓ portfolios:

$$Q = P\mu + \nu \tag{5}$$

where Q is an $\ell \times 1$ vector of expected portfolio returns, P is an $\ell \times N$ matrix representing the portfolios, and $\nu \sim N(0, \Omega)$. Ω is the covariance matrix of the noise term ν. Calling Q our investment *views*, Ω then represents the uncertainties in these views.

For example, in the case of an investment view that the first stock would outperform the second stock by 3% with some degree of uncertainty, Q will be just a scalar equal to 3%, while P will be a $1 \times N$ vector with a 1 as the first element, −1 as the second element, and zeroes as all other elements.

Rearranging, we get:

$$Q = P(X\bar{f} + \bar{\varepsilon}) + \nu \tag{6}$$

Thus, we can express the joint distribution of factor returns, stock-specific returns, and investment views by the multivariate normal distribution:

$$\begin{bmatrix} \bar{f} \\ \bar{\varepsilon} \\ Q \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \tau_f^2 F & 0 & \tau_f^2 FX'P' \\ 0 & \tau_\varepsilon^2 D & \tau_\varepsilon^2 DP' \\ \tau_f^2 PXF & \tau_\varepsilon^2 PD & P\left(\tau_f^2 XFX' + \tau_\varepsilon^2 D\right)P' + \Omega \end{bmatrix} \right) \tag{7}$$

In Appendix A, we prove that the conditional expectations of factor returns and stock-specific returns, given the investment views, are:

$$\begin{aligned} E\left(\bar{f} \mid Q\right) &= \tau_f^2 FX'P'\left[P\left(\tau_f^2 XFX' + \tau_\varepsilon^2 D\right)P' + \Omega\right]^{-1} Q \\ E\left(\bar{\varepsilon} \mid Q\right) &= \tau_\varepsilon^2 DP'\left[P\left(\tau_f^2 XFX' + \tau_\varepsilon^2 D\right)P' + \Omega\right]^{-1} Q \end{aligned} \tag{8}$$

What this matrix algebra shows is that views on returns, Q, reflect some views on conditional factor returns, $\bar{f}$, as well as on stock-specific returns, $\bar{\varepsilon}$. Simple algebra will show that any non-trivial investment view Q will have no information on factor returns if and only if conditions as follows are satisfied:

$$PX = 0 \tag{9}$$

In other words, any investment views on stock returns can be decomposed as views on factor returns and views on stock-specific returns, unless these investment views are specified only for factor-neutral portfolios. The extent to which the information advantage falls between factor returns and stock-specific returns depends on the relative size of the control parameters, τ_f and τ_ε.

EMPIRICAL EXAMPLES

We can present evidence that the imposition of factor neutrality in the context of a realistic and commonly used multifactor risk model can lead to a significant deterioration in investment performance. This may be true even if the investor has only stock-specific information and none of the common risk factors is priced. It is especially true when there are too few stocks or not a variety of stocks in the portfolio to form good factor-mimicking portfolios for the purpose of hedging factor risks in maintaining factor neutrality.

We adopt Barra's E3 model for U.S. equities, among the most widely used risk models in the investment industry. In this model, stock returns exhibit a factor structure with $K = 68$ common factors, 13 of them risk indexes and 55 industries. We use the estimated X, F, and D as of June 30, 2004, as provided by Barra.

Exhibit 1

Example 1—Expected Sharpe Ratios of Factors and Stock-Specific Returns

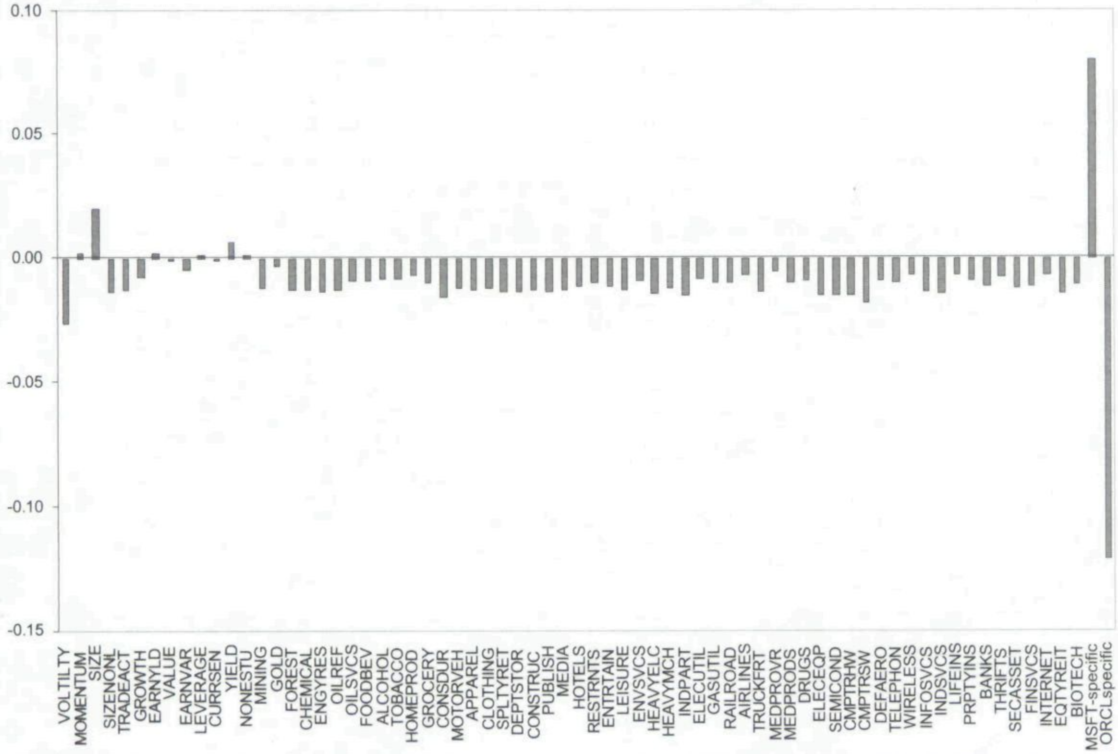

To simplify matters, in the first two examples below, we restrict the control parameters, τ_f and τ_ε, to be equal to τ so that the strength of information on factor returns and stock-specific returns is the same, all else equal. This restriction is relaxed in the third example, where we investigate the effects on investment performance of changing the relative size of these two parameters.

Example 1: Buy Microsoft, Sell Oracle

Suppose an equity analyst who follows the software industry has only one investment view: that Microsoft (MSFT) is expected to outperform Oracle (ORCL) by 5%. All else equal, this investment view will lead to an increased weighting of MSFT in the portfolio, funded by a sale of ORCL shares.

To put this in our analytical framework, we assume that the analyst has no uncertainty about this view. That is, Ω is zero. We set τ at a value of 0.25, which generates resulting expected returns that appear to be reasonable. Using Equation (8), we compute the expected factor returns and expected stock-specific returns conditional on this single investment view.

Exhibit 1 plots the expected Sharpe ratios of the conditional factor returns and stock-specific returns. Overall, only the two specific returns have meaningful expected Sharpe ratios, while the other factors have ratios that are negligible.

A clear message from this illustration is that this is the relatively bullish view on MSFT's specific return compared with that of ORCL. This result should not be surprising, as risks of individual stocks are dominated by their idiosyncratic risks. Therefore, a pair trade of this sort should also be dominated by the stocks' specific returns—precisely what our analytical framework shows.

The negligible expected Sharpe ratios of the 13 Barra risk indexes indicate that exposures to these indexes are hedged away to a large extent, as the two stocks have comparable characteristics, likely because they are both in the same industry, with similar business profiles. Expected Sharpe ratios for all 55 industry returns are slightly negative. This is because, according to the Barra covariance matrix, MSFT has a lower beta than ORCL. Consequently, a long/short MSFT/ORCL trade will have a negative net beta, which shows up as negative expected returns to all industries.

In short, maintaining factor neutrality in this exam-

EXHIBIT 2

Example 2—Expected Sharpe Ratios of Factors and Stock-Specific Returns

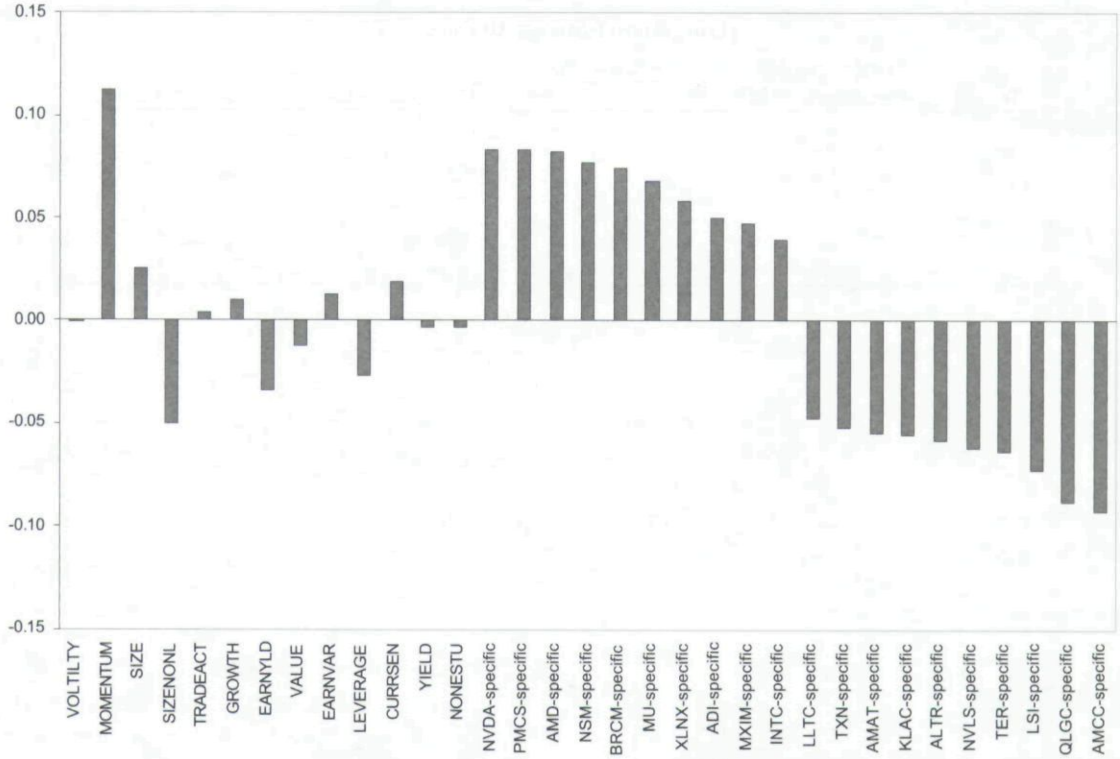

ple appears to have a negligible impact on investment performance.

Example 2: Buy/Sell Basket of Semiconductor Stocks

Suppose now that an equity analyst covers 20 stocks in the semiconductor industry, assigning the stocks to two groups: stocks expected to deliver higher-than-average returns, and stocks expected to deliver lower-than-average returns. This investment view will lead to an increase in the weighting of stocks in the first group, funded by the sale of stocks in the second group.

To put this into our analytical framework, we assume that the information can be summarized by a single investment view, that an equal-weighted portfolio of ten stocks in the first group is expected to outperform an equal-weighted portfolio of the other ten stocks in the second group. We maintain the same assumptions, that Ω is zero and τ is 0.25, as in Example 1. Once again, using Equation (8), we compute the expected factor returns and expected stock-specific returns conditional on this single investment view.

Exhibit 2 plots the expected Sharpe ratios of the factor returns and stock-specific returns. To make the graph more readable, we exclude all 55 industries, all found to have insignificant expected Sharpe ratios. We plot only the expected Sharpe ratios of the 13 Barra risk indexes and the 20 stock-specific returns.

Again, all 20 stock-specific returns have meaningful expected Sharpe ratios, suggesting that to a significant extent this investment view is betting on the outperformance of stock-specific returns of one portfolio of stocks versus another. Yet it is interesting to see that some risk indexes also have expected Sharpe ratios that are not negligible, meaning that exposures to these factors are not hedged away in this investment view.

In other words, the returns of these factors clearly affect the performance and risk characteristics of this trade. The momentum factor has the highest expected Sharpe ratio among all other factors and stock-specific returns. As a result, maintaining factor neutrality in this case ignores some useful information on returns, and thus will hurt investment performance to a significant extent.

EXHIBIT 3
Simulation Results in S&P 500 With and Without Factor Neutrality

$\tau_s = 0.05$		**Long-Short Portfolio, Universe = S&P 500**		
τ_f	Unconstrained expected info ratio	Factor-neutral expected info ratio	Expected % Loss	Standard Error
0.00	1.07	1.04	3.3%	0.01%
0.05	1.12	1.04	7.0%	0.01%
0.10	1.24	1.04	15.9%	0.02%
0.15	1.41	1.04	26.2%	0.03%
0.20	1.63	1.04	35.9%	0.04%
0.25	1.86	1.04	44.0%	0.04%

Standard error is the standard deviation of the sampling error in the estimated expected % loss.

Example 3: General Case

Let $\hat{\mu}$ be the posterior mean of μ conditional on all information available to the investor. We show in Appendix B that, in the absence of constraints (including borrowing and short-sale constraints), any mean-variance efficient portfolio has the information ratio:[2]

$$IR_{\text{unconstrained}} = \sqrt{\hat{\mu}'\Sigma^{-1}\hat{\mu}} \tag{10}$$

while any portfolio that is mean-variance optimal subject to the constraint of factor neutrality has the information ratio:[3]

$$IR_{\text{factor-neutral}} = \sqrt{\hat{\mu}^{*\prime}\Sigma^{-1}\hat{\mu}^{*}} \tag{11}$$

where

$$\hat{\mu}^{*} = \left(I - X\left(X'\Sigma^{-1}X\right)^{-1}X'\Sigma^{-1}\right)\hat{\mu} \tag{12}$$

For any given $\hat{\mu}$, we can compute $IR_{\text{unconstrained}}$ and $IR_{\text{factor-neutral}}$. The expected fractional loss of investment performance caused by factor neutrality is then given by:

$$1 - \frac{IR_{\text{factor-neutral}}}{IR_{\text{unconstrained}}} \tag{13}$$

Given values of τ_f and τ_ε, we can estimate the cost of factor neutrality by running Monte Carlo simulations, making draws of $\bar{f}$ and $\bar{\varepsilon}$ from their unconditional distributions in (4) and thereby sampling $\hat{\mu}$, and then computing the optimal information ratio with and without factor neutrality for each draw. In the simulations, we fix τ_ε = 0.05 and vary τ_f between zero and 0.25 in steps of 0.05. We argue that it is reasonable to have a τ_f that is up to several times the magnitude of τ_ε, as it is generally reasonable to believe that the return of a diversified portfolio, such as a factor-mimicking portfolio, is potentially more predictable than the specific return of a single stock.

The results for the universe of all stocks in the S&P 500 index are reported in Exhibit 3. Not surprisingly, there is a greater expected loss from factor neutrality when the investor has stronger information on factor returns. Under factor neutrality, the expected information ratio is 1.04, whatever the strength of the investment manager's information on factor returns, since taking advantage of this information by taking factor risk is not permitted. Without the constraint of factor neutrality, a higher τ_f, which means more information on factor returns, always leads to a higher expected information ratio. Thus, it is no surprise to see that the cost of factor neutrality is greater when τ_f is higher.

For example, at $\tau_f = 0.25$, the expected loss in information ratio is 44.0%, from the unconstrained information ratio of 1.86 to the factor-neutral constrained information ratio of 1.04.

What is perhaps a little surprising is that even when the investor has absolutely no information on factor returns, that is, $\tau_f = 0$, factor neutrality still leads to a small loss in the expected information ratio, of 3.3%, from 1.07 to 1.04. This can be understood as follows. With the absence of information on factor returns at $\tau_f = 0$, we would like to have exposures only to stock-specific risks and not to any factor risks at all. Yet to gain exposure to a particular stock, we must buy or sell the stock, including all the factor exposures that come with it. In order to hedge these factor exposures, we must buy and sell some other stocks, including the stock-specific exposures that come with them as well.

EXHIBIT 4

Simulation Results in Top 100 Stocks With and Without Factor Neutrality

$\tau_s = 0.05$		Long-Short Portfolio, Universe = Top 100		
τ_f	Unconstrained expected info ratio	Factor-neutral expected info ratio	Expected % Loss	Standard Error
0.00	0.45	0.32	29.0%	0.05%
0.05	0.50	0.32	35.3%	0.05%
0.10	0.62	0.32	47.5%	0.06%
0.15	0.77	0.32	57.8%	0.05%
0.20	0.94	0.32	65.6%	0.05%
0.25	1.13	0.32	71.2%	0.04%

Standard error is the standard deviation of the sampling error in the estimated expected % loss.

EXHIBIT 5

Simulation Results in Top 50 Stocks With and Without Factor Neutrality

$\tau_s = 0.05$		Long-Short Portfolio, Universe = Top 50		
τ_f	Unconstrained expected info ratio	Factor-neutral expected info ratio	Expected % Loss	Standard Error
0.00	0.31	0.12	62.2%	0.10%
0.05	0.35	0.12	66.6%	0.09%
0.10	0.45	0.12	73.8%	0.08%
0.15	0.58	0.12	79.2%	0.06%
0.20	0.72	0.12	83.2%	0.05%
0.25	0.87	0.12	86.2%	0.04%

Standard error is the standard deviation of the sampling error in the estimated expected % loss.

We can hedge away factor risks without any distortions of the desirable stock-specific risks only if we can trade factor-mimicking portfolios that are so diversified as to have negligible stock specific risks. This is possible only if we are prepared to have positions in many stocks from a wide universe within which the covariance matrix is estimated. Otherwise, if the universe of stocks permitted in the portfolio is restricted, the hedging of factor exposures that is required for factor neutrality inevitably leads to stock-specific exposures that would be suboptimal in the absence of factor neutrality.

The Barra covariance matrix is estimated with a universe of more than 2,000 stocks, including all stocks in the S&P 500 index. Therefore, a universe limited to only the constituents of the S&P 500 index may not be sufficient to span all the estimated factors in the Barra universe.

While holding positions in all S&P 500 constituents limits the expected loss in information ratio to a small 3.3% (from 1.07 to 1.04) when $\tau_f = 0$, the loss can be more serious in the cases of more concentrated portfolios. This is particularly relevant as a typical active equity portfolio, traditional long-only or long-short hedge funds, holds about 50 to little more than 100 positions.

Exploring this matter further, Exhibit 4 reports the simulation results when the universe is restricted to the 100 largest stocks in the S&P 500 index. In this smaller universe, the expected loss in information ratio is 29.0% when $\tau_f = 0$, from 0.45 to 0.32. That is, one can expect the portfolio to lose 29% of its alpha when factor neutrality is imposed, even though all factors are not priced and have zero expected returns. When $\tau_f = 0.25$, the expected loss of information ratio from factor neutrality is a stunning 71.2%, from 1.13 to 0.32.

Finally, Exhibit 5 reports the simulation results when the universe is restricted further, to the 50 largest stocks in the S&P 500 index. The expected losses in information ratios are 62.2% from 0.31 to 0.12 when $\tau_f = 0$, and 86.2% from 0.87 to 0.12 when $\tau_f = 0.25$. That is, within such a small universe, 60% to almost 90% of alphas will be lost by the factor neutrality constraint imposed, whether or not the investor has information on factor returns.

To summarize all these results, Exhibit 6 plots the information ratios with and without factor neutrality constraints with 500, 100, and 50 stocks as a function of

EXHIBIT 6

Example 3—Effects of Factor Neutrality on Information Ratios of Long/Short Portfolios

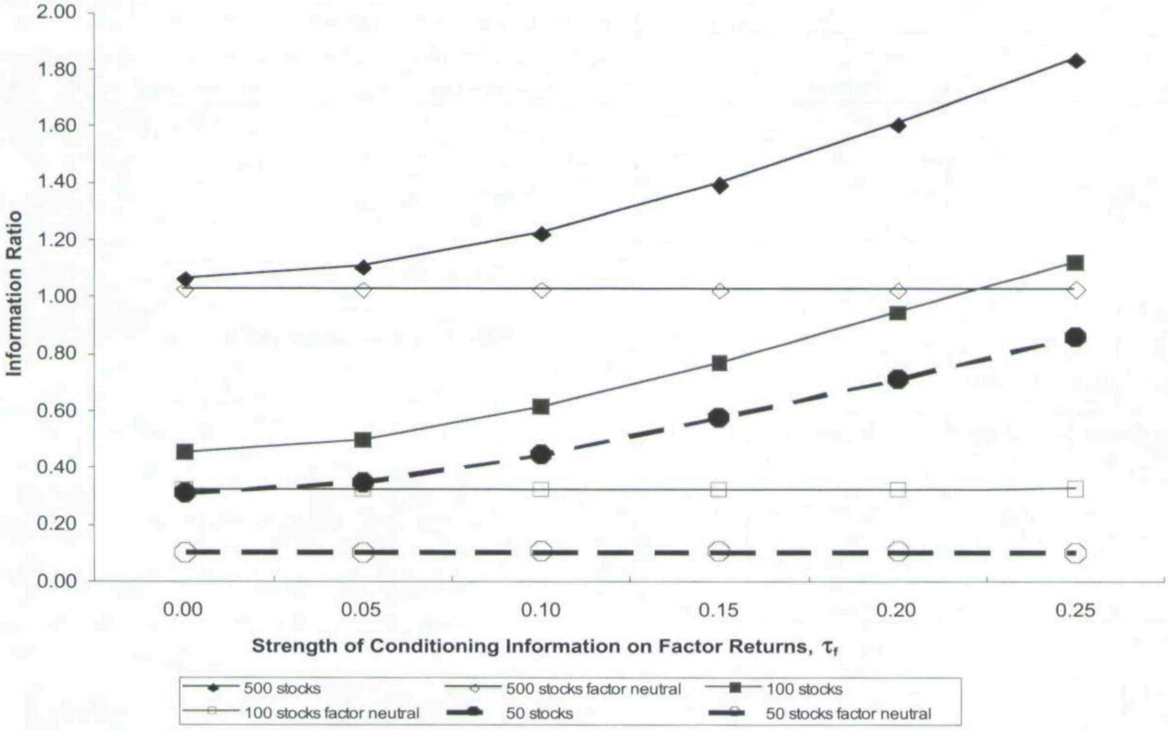

the strength of conditioning information on factor returns, τ_f. The gaps between the pairs of two lines at $\tau_f = 0$ for the three cases correspond to the deadweight costs of factor neutrality due to the differences in universe for factor structure estimation and construction of portfolios, even though all information is on stock-specific returns only. The fewer the stocks in the portfolio, the higher the deadweight cost of factor neutrality.

A natural question is what happens in cases other than the extremes of unconstrained factors and strict neutrality we have considered so far. To address this, we impose an upper bound on the percentage of portfolio variance due to exposures to risk factors, so that the vector of asset weights ω relative to benchmark is constrained to satisfy:

$$\frac{\omega' XFX' \omega}{\omega' \Sigma \omega} \leq p \tag{14}$$

for some $0 \leq p \leq 1$. With this constraint, the optimal information ratio is no longer given by simple analytical expressions such as Equations (10) or (11), but can still be evaluated using numerical optimization.[4]

Exhibits 7-9 report the results of simulations with an upper bound on the percentage of portfolio variance due to factor exposures ranging from 10% ($p = 0.1$) to 50% ($p = 0.5$). When the investor has no information on factor returns ($\tau_f = 0$), allowing up to 10% of the risk to come from factor exposures is enough to eliminate the deadweight loss of information ratio that would otherwise occur with factor neutrality, even for a small universe of the top 50 stocks. Thus, alleviating the cost of factor neutrality does not necessarily require taking massive factor bets, but rather only a degree of flexibility in deviating from strict factor neutrality.

CONCLUSIONS

Prudent investors should take intentional risks, and hedge away risks that are unintentional. In theory, if investors have no information advantage on a particular return factor, they should construct portfolios so that active exposure to this factor relative to the benchmark is negligible. In practice, however, the exact factor structure and the number of factors are usually unknown. Instead, the factor structure has to be estimated on the basis

EXHIBIT 7

Simulation Results—S&P 500 with Upper Bound on Factor Risk

$\tau_s = 0.05$	Upper bound on factor risk %					
τ_f	0%	10%	20%	30%	40%	50%
0.00	3.3%	0.0%	0.0%	0.0%	0.0%	0.0%
0.05	7.0%	0.0%	0.0%	0.0%	0.0%	0.0%
0.10	15.9%	1.2%	0.0%	0.0%	0.0%	0.0%
0.15	26.2%	5.2%	1.1%	0.1%	0.0%	0.0%
0.20	35.9%	10.1%	3.7%	0.8%	0.0%	0.0%
0.25	44.0%	14.9%	6.8%	2.4%	0.4%	0.0%

EXHIBIT 8

Simulation—Top 100 Stocks with Upper Bound on Factor Risk

$\tau_s = 0.05$	Upper bound on factor risk %					
τ_f	0%	10%	20%	30%	40%	50%
0.00	29.0%	0.1%	0.0%	0.0%	0.0%	0.0%
0.05	35.3%	1.6%	0.0%	0.0%	0.0%	0.0%
0.10	47.5%	7.6%	1.7%	0.2%	0.0%	0.0%
0.15	57.8%	14.4%	5.3%	1.3%	0.1%	0.0%
0.20	65.6%	19.9%	8.6%	2.9%	0.5%	0.0%
0.25	71.2%	24.1%	11.4%	4.4%	1.1%	0.1%

EXHIBIT 9

Simulation—Top 50 Stocks with Upper Bound on Factor Risk

$\tau_s = 0.05$	Upper bound on factor risk %					
τ_f	0%	10%	20%	30%	40%	50%
0.00	62.2%	0.5%	0.0%	0.0%	0.0%	0.0%
0.05	66.6%	3.5%	0.3%	0.0%	0.0%	0.0%
0.10	73.8%	11.2%	3.5%	0.7%	0.1%	0.0%
0.15	79.2%	18.3%	7.6%	2.6%	0.6%	0.1%
0.20	83.2%	23.6%	11.0%	4.4%	1.2%	0.2%
0.25	86.2%	27.4%	13.6%	5.9%	1.9%	0.3%

of a subset of the investment universe, typically a universe of stocks with a long enough history of high-quality data.

In fact, investment managers typically outsource the estimation of a factor model to vendors, and take the resulting risk model as exogenously given. Therefore, pure stock-picking, which consists of generating stock-specific returns that are independent of all true underlying factors, is close to an absolute impossibility. Consequently, investment opinions really reflect views on both factor returns and stock-specific returns to various degrees.

We have developed an analytical framework to understand the deterioration in investment performance when maintaining factor neutrality. Our results suggest that the degree of loss in information ratio due to factor neutrality can be substantial, particularly in more realistic cases when there are 100 or fewer stocks in the portfolio. This is true even when all information is on stock-specific returns as a result of the different universes used to estimate the unknown factor structure and construct the portfolio.

Our results support the idea that the common practice among many investment managers of maintaining factor neutrality can be a suboptimal process. We suggest it is advisable to allow for modest factor exposures in equity portfolios, even when one has no explicit views on factor returns.

APPENDIX A

Proof of Equation (8)

For a bivariate normal distribution:

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \sim N\left(\begin{bmatrix} \overline{x_1} \\ \overline{x_2} \end{bmatrix}, \begin{bmatrix} S_{11} & S_{12} \\ S_{12}' & S_{22} \end{bmatrix}\right) \tag{A-1}$$

the conditional distribution of one on another is given by

$$x_1 \mid x_2 \sim N\left(\overline{x_1} + S_{12}S_{22}^{-1}\left(x_2 - \overline{x_2}\right), S_{11} - S_{12}S_{22}^{-1}S_{12}'\right) \tag{A-2}$$

Applying these results and Equation (7) gives:

$$\begin{bmatrix} \bar{f} \mid Q \\ \bar{\varepsilon} \mid Q \end{bmatrix} \sim N\left(\begin{bmatrix} \tau_f^2 FX'P'\left[P\left(\tau_f^2 XFX' + \tau_\varepsilon^2 D\right)P' + \Omega\right]^{-1} Q \\ \tau_\varepsilon^2 DP'\left[P\left(\tau_f^2 XFX' + \tau_\varepsilon^2 D\right)P' + \Omega\right]^{-1} Q \end{bmatrix}, \begin{bmatrix} V_{11} & V_{12} \\ V_{12}' & V_{22} \end{bmatrix}\right) \tag{A-3}$$

where

$$V_{11} = \tau_f^2 F - \tau_f^4 FX'P'\left[P\left(\tau_f^2 XFX' + \tau_\varepsilon^2 D\right)P' + \Omega\right]^{-1} PXF \tag{A-4}$$

$$V_{22} = \tau_\varepsilon^2 D - \tau_\varepsilon^4 DP'\left[P\left(\tau_f^2 XFX' + \tau_\varepsilon^2 D\right)P' + \Omega\right]^{-1} PD \tag{A-5}$$

$$V_{12}' = \tau_f^2 \tau_\varepsilon^2 DP'\left[P\left(\tau_f^2 XFX' + \tau_\varepsilon^2 D\right)P' + \Omega\right]^{-1} PXF \tag{A-6}$$

APPENDIX B

Proof of Equations (10) and (11)

In the mean-variance framework, we maximize returns with a penalty to portfolio variance risk. That is, we solve:

$$\max_{\omega} \; \omega'\hat{\mu} - \frac{1}{2\lambda}\omega'\Sigma\omega \tag{B-1}$$

subject to

$$C\omega = c \qquad \text{(B-2)}$$

where ω is an N-by-1 vector of portfolio weights, $\hat{\mu}$ is an N-by-1 vector of returns in excess of the risk-free rate, λ is a risk tolerance parameter, Σ is the N-by-N covariance matrix, C is a j-by-N matrix, c is a j-by-1 vector of constrained values, and j is the number of constraints.

The Lagrangian is:

$$L = \omega'\hat{\mu} - \frac{1}{2\lambda}\omega'\Sigma\omega - \gamma'\left(C\omega - c\right) \qquad \text{(B-3)}$$

where γ is a j-by-1 vector of multipliers. Differentiating the Lagrangian, we get the first-order conditions:

$$\omega = \lambda\Sigma^{-1}\left(\hat{\mu} - C'\gamma\right) \qquad \text{(B-4)}$$

and

$$C\omega = c \qquad \text{(B-5)}$$

Combining the two, we get:

$$\lambda C\Sigma^{-1}\left(\hat{\mu} - C'\gamma\right) = c \qquad \text{(B-6)}$$

or

$$\gamma = \frac{C\Sigma^{-1}\hat{\mu}}{C\Sigma^{-1}C'} - \frac{1}{\lambda}\frac{c}{C\Sigma^{-1}C'} \qquad \text{(B-7)}$$

Substituting (B-7) into (B-4) and rearranging terms, we get:

$$\begin{aligned}\omega = {} & \Sigma^{-1}C'\left(C\Sigma^{-1}C'\right)^{-1}c + \\ & \lambda\Sigma^{-1}\left[I - C'\left(C\Sigma^{-1}C'\right)^{-1}C\Sigma^{-1}\right]\hat{\mu}\end{aligned} \qquad \text{(B-8)}$$

When there are no constraints, the optimal portfolio weight can be simplified as:

$$\omega_{unconstrained} = \lambda\Sigma^{-1}\hat{\mu} \qquad \text{(B-9)}$$

As a result, the unconstrained information ratio can be expressed as

$$IR_{unconstrained} = \frac{\omega'_{unconstrained}\hat{\mu}}{\sqrt{\omega'_{unconstrained}\Sigma\omega_{unconstrained}}} = \sqrt{\hat{\mu}'\Sigma^{-1}\hat{\mu}} \qquad \text{(B-10)}$$

which is Equation (10).

In the case of factor neutrality constraints, we set $C = X'$ where X is the N-by-K matrix of factor exposures, and $c = 0$, a K-by-1 vector of zeros. Thus, the optimal portfolio weights are given by:

$$\omega_{factor-neutral} = \lambda\Sigma^{-1}\left[I - X\left(X'\Sigma^{-1}X\right)^{-1}X'\Sigma^{-1}\right]\hat{\mu} \qquad \text{(B-11)}$$

where

$$\hat{\mu}^* = \left[I - X\left(X'\Sigma^{-1}X\right)^{-1}X'\Sigma^{-1}\right]\hat{\mu} \qquad \text{(B-12)}$$

Multiplying both sides by Σ^{-1} gives

$$\begin{aligned}\Sigma^{-1}\hat{\mu}^* &= \Sigma^{-1}\hat{\mu} - \Sigma^{-1}X\left(X'\Sigma^{-1}X\right)^{-1}X'\Sigma^{-1}\hat{\mu} \\ \Sigma^{-1}\hat{\mu} &= \Sigma^{-1}\hat{\mu}^* + \Sigma^{-1}X\left(X'\Sigma^{-1}X\right)^{-1}X'\Sigma^{-1}\mu \\ \hat{\mu}^{*\prime}\Sigma^{-1}\hat{\mu} &= \hat{\mu}^{*\prime}\Sigma^{-1}\hat{\mu}^* + \hat{\mu}^{*\prime}\Sigma^{-1}X\left(X'\Sigma^{-1}X\right)^{-1}X'\Sigma^{-1}\hat{\mu} \\ &= \hat{\mu}^{*\prime}\Sigma^{-1}\hat{\mu}^*\end{aligned} \qquad \text{(B-13)}$$

since

$$X'\Sigma^{-1}\hat{\mu}^* = \frac{1}{\lambda}X'\omega_{factor-neutral} = 0$$

Thus the information ratio with the factor neutrality constraint is:

$$\begin{aligned}IR_{factor-neutral} &= \frac{\omega_{factor-neutral}'\hat{\mu}}{\sqrt{\omega_{factor-neutral}'\Sigma\omega_{factor-neutral}}} \\ &= \frac{\hat{\mu}^{*\prime}\Sigma^{-1}\hat{\mu}}{\sqrt{\hat{\mu}^{*\prime}\Sigma^{-1}\hat{\mu}^*}}\end{aligned} \qquad \text{(B-14)}$$

using (B-13). This proves Equation (11).

ENDNOTES

The authors thank Eric Remole, Ingrid Tierens, and Alexandre Da Silva for comments, and John Geer and Philip Maher for editorial assistance.

[1]This is related to the structural bet issue in tactical asset allocation (TAA) as discussed in detail in Lee [2000]. A TAA manager may deliver a positive information ratio over a reasonable period of time simply by adding a constant structural overweight on, for example, stocks against bonds. The issue is considerably much more complicated in equity investing than in TAA. While there is compelling evidence that stocks com-

mand a risk premium over bonds, evidence on factor pricing in the equity universe is less convincing. Besides, because of a limited universe in asset allocation, it is easy to evaluate performance due to a constant structural bet of one asset against another. In the large universe of equity investing, though, unknown factors have first to be defined and estimated, before any positive premium can be examined. Furthermore hedging exposures to factors involve trades, both long and short, on a large number of stocks in the factor-mimicking portfolios.

[2]The absence of constraints implies that the efficient frontier is a straight line passing through the benchmark asset, and hence that all mean-variance efficient portfolios have the same information ratio, so Equation (10) holds whatever the level of risk aversion or targeted risk. If there are long-only constraints, our results will hold exactly only when the no-shorting constraints are not binding. This occurs at the limit as the tracking error of the portfolio approaches zero when all assets in the universe are held in the benchmark. These conditions approximate the case of enhanced indexing with very low tracking error. Thus our results cast doubt on the common perception that strict factor neutrality is especially appropriate for enhanced-indexing strategies.

[3]A portfolio is factor-neutral if and only if any scaled version of itself (with all the active weights relative to the benchmark multiplied by an arbitrary real number) is factor-neutral. Thus the efficient frontier under factor neutrality remains a straight line, and Equation (11) holds whatever the level of risk aversion or targeted risk.

[4]Constraint (14) is scale-invariant in the sense that a portfolio with weights ω satisfies the constraint if and only if any scaled portfolio $a\omega$ where a is any non-zero real number, also satisfies the constraint. Thus the imposition of this constraint does not cause the efficient frontier to change from a straight line, and therefore the optimal information ratio is independent of the level of risk aversion or targeted risk.

REFERENCES

Clarke, Roger, Harindra de Silva, and Steven Thorley. "Portfolio Constraints and the Fundamental Law of Active Management." *Financial Analysts Journal*, September/October 2002, pp. 48-66.

Grinblatt, Mark, and Sheridan Titman. "Portfolio Performance Evaluation: Old Issues and New Insights." *The Review of Financial Studies*, Vol. 2, No. 3 (1989), pp. 393-421.

Grinold, Richard C. "The Fundamental Law of Active Management." *The Journal of Portfolio Management*, Spring 1989, pp. 30-37.

Lee, Wai. *Theory and Methodology of Tactical Asset Allocation*. New York: John Wiley & Sons, 2000.