Journal Of Investment Management, Vol. 14, No. 2, (2016), pp. 71–85
© JOIM 2016

www.joim.com

FACTOR MISALIGNMENT AND PORTFOLIO CONSTRUCTION

Jose Menchero

In recent years, there has been heightened interest among practitioners in the topic of factor misalignment; this term refers to the practice of employing mean-variance optimization to construct portfolios when the alpha signal is not contained within the set of risk model factors. In this paper, we employ a realistic simulation framework to study the efficiency of optimized portfolios under a variety of conditions. In particular, we study the case in which the alpha factor contains true systematic risk, and the case in which it does not. We also consider two risk models: one that contains the alpha factor, and the other that omits it. We find some evidence to support a modest increase in portfolio information ratio when the alpha factor is included in the risk model, provided two conditions hold: (1) the alpha factor must include true systematic risk, and (2) the factor correlations must be estimated with sufficient precision. If the alpha factor does not contain true factor risk, we find that including the alpha signal in the risk model is detrimental to portfolio information ratio. Finally, we conduct an empirical analysis of portfolio efficiency in the US stock market and find that the results are in excellent agreement with our simulations.

1 Introduction

Mean–variance optimization is widely used as a portfolio construction tool by quantitative asset managers. The technique was pioneered by Markowitz (1952) and provided for the first time a formal structure for balancing the trade-off between risk and return. The required inputs to the Markowitz optimization problem include: (a) the expected returns, or alphas, of every asset, (b) the asset covariance matrix, and (c) a set of investment constraints.

The alphas are typically generated using a proprietary set of alpha signals, or factors. The primary purpose of an alpha factor is to outperform a benchmark. Hence, alpha factors are selected based on their ability to explain the mean, or the *first* moment, of the portfolio return distribution.

By contrast, the asset covariance matrix is typically obtained using a multi-factor risk model, often provided by a third-party vendor. The

This work was conducted at MSCI. Jose Menchero is currently head of portfolio analytics research at Bloomberg.

Not for Distribution

factors used in such models are selected based on their ability to explain cross-sectional variation in stock returns, with the primary aim being to predict portfolio volatility. That is, risk factors are used to estimate the standard deviation, or *second* moment, of the portfolio return distribution.

Given these divergent objectives, it is not surprising that differences will exist between the two types of factors. In general, unless the alpha factor is explicitly included in the risk model, it is not possible to represent alpha as a linear combination of risk model factors. When this occurs, we say that there is *misalignment* between the alpha factor and the risk factors.

In general, alpha signals can always be decomposed into two components by performing a cross-sectional regression of stock alphas against the risk model factors. The part of alpha that is explained by the risk factors is known as the *spanned alpha*. The second component is orthogonal to the risk factors and represents the *residual alpha*.

The risk model views the residual alpha as diversifiable idiosyncratic risk. By contrast, the spanned alpha represents non-diversifiable factor risk. Hence, on a risk-adjusted basis, the residual component appears far more attractive. This causes the optimizer to tilt the portfolio holdings in the direction of the residual alpha. If one makes the rather heroic assumption that the alpha signal and risk model are *known without error*, then such tilting would always be desirable as it maximizes the *ex ante* portfolio Information Ratio.

In reality, of course, neither the alpha signal nor the risk model can be known with certainty. Instead, these quantities must be estimated from available data. This process inevitably leads to errors in the factor exposures, the factor returns, and the covariance matrix. Such errors may have detrimental effects on portfolios constructed using mean–variance optimization.

Lee and Stefek (2008) were among the first to investigate the effects of factor misalignment in portfolio construction. They considered the example of a momentum alpha signal constructed using trailing 12-month returns with a 1-month lag, with a momentum risk factor constructed using trailing 12-month returns with no lag. In this example, stocks that performed well 13 months ago contribute positively to alpha, but have no associated factor risk. Hence, these stocks appear particularly attractive. By contrast, stocks that performed well last month contain factor risk, but have no associated alpha. Indeed, Lee and Stefek showed that the optimized portfolio tends to heavily overweight stocks that performed well 13 months ago, while underweighting those that performed well over the previous month. These would likely represent unintentional bets by the portfolio manager, and may be detrimental to performance.

Ceria *et al.* (2012) highlighted other problems that may arise from factor misalignment. One such problem is the underestimation of risk of optimized portfolios. This occurs when the residual alpha contains factor risk that was omitted from the risk model. In this case, the optimizer will align the portfolio in the direction of the missing factor, thus leading to an underestimation of risk. Ceria *et al.* also discussed the reduction in Information Ratio that can result from suboptimal allocation of the risk budget due to factor misalignment.

In this paper, we investigate the impact on portfolio Information Ratio of including alpha factors in risk models. We consider two types of risk models, denoted Model A and Model B. These two models are identical in every respect except that Model A includes the alpha factor, whereas it is

Not for Distribution

omitted from Model B. We also consider two distinct cases, denoted Case 1 and Case 2. In the first case, the residual alpha contains true factor risk, whereas it is purely idiosyncratic in the second case.

By considering Case 1 and Case 2 jointly with Model A and Model B, we naturally segment the world into four distinct possibilities. When the residual alpha contains true factor risk (Case 1), Model A will be properly specified since it has included a valid risk factor, whereas Model B will be misspecified for failing to include the factor. By contrast, when the residual alpha contains no factor risk (Case 2), Model A will be misspecified for including a spurious risk factor, whereas Model B will be properly specified for correctly omitting the factor.

Our objective is to compare the efficiency of optimized portfolios constructed using Model A and Model B. To conduct our analysis, we begin with a simulation framework designed to realistically mimic the behavior of the US equity market. Within our simulations, we know the true return-generating process. This enables us to determine the true risk and return characteristics of the portfolios we study. The simulation framework also enables us to study the effects of sampling error in a controlled setting. Next, we show that our simulated results are in excellent agreement with out-of-sample empirical results for the US equity market. We conclude with a discussion on the sources of factor misalignment and compare various potential treatments.

2 Alpha factors and risk factors

Mathematically, factors are represented as $N \times 1$ column vectors, where N is the number of stocks. If we consider K such factors, the factor exposure matrix has dimensionality $N \times K$. We estimate factor returns by multivariate cross-sectional regression of stock returns against the factor exposures. As discussed by Menchero (2010), factor returns are properly interpreted as the returns of pure factor portfolios that have unit exposure to the factor in question, and zero exposure to the other factors in the regression.

Of course, not every $N \times 1$ column vector represents an interesting factor from an investment viewpoint. The factors that are of financial interest can be categorized into alpha factors and risk factors. Before proceeding with our analysis, it is essential to first define these terms precisely.

In this paper, we define an *alpha* factor to be one that is effective at explaining the first moment of stock returns. More specifically, if a pure factor portfolio has non-zero expected return, it represents an alpha factor. In other words, alpha factors capture sources of directional drift. A "good" alpha factor is one that has a high Information Ratio, meaning that the expected return of the pure factor portfolio is large relative to its volatility. An equivalent way of stating this is that the pure alpha factor portfolios will have significant t-statistics in their return time series.

We define a *risk* factor to be one that is effective at explaining the second moment of the return distribution. That is, risk factors identify the sources of cross-sectional variation in stock returns. Stocks with positive exposure to a given factor will earn a return contribution opposite in sign to those with negative exposure. Since the pure factor portfolio takes net-long positions in stocks with positive exposure and net-short positions in stocks with negative factor exposure, a "good" risk factor will have high volatility. This implies that the factor returns tend to be large relative to the sampling error in the return estimate. In other words, a good risk factor will have a high percentage of significant t-statistics in the cross-sectional regressions.

Not for Distribution

It is important to recognize that alpha factors and risk factors are not mutually exclusive categories. For instance, some factors may simultaneously represent both alpha factors and risk factors. Momentum is a good example of such a factor, as it tends to exhibit both positive drift and high volatility.

Most risk factors, however, cannot rightly be considered as alpha factors. For example, industries and countries are often excellent risk factors, as they capture important sources of equity return co-movement. Nonetheless, most investors do not view them as persistent sources of alpha. Similarly, the beta factor portfolio is extremely volatile, thus making it an excellent risk factor. By contrast, it tends not to exhibit persistent drift, thus making it a poor alpha factor.

In principle, there may exist alpha factors that are *not* sources of systematic risk, although such factors may be exceedingly difficult to find in practice. Nonetheless, consider a signal derived from the *unique* insights of a portfolio manager. Such a factor is unlikely to be a risk factor, since no other investors would be actively trading the signal. On the other hand, if the portfolio manager's insights are valid, it may represent a good alpha factor.

Finally, according to our definitions, it is trivial to construct a "factor" that is neither alpha factor nor risk factor. Such a factor could be produced, for instance, by simply drawing the factor exposures at random from a standard normal distribution.

3 Simulated factor portfolios

A simulation exercise is helpful for illustrating the behavior of the different factor types. The Cholesky decomposition, as described for example by Greene (2000), is a common statistical technique used to simulate returns that are drawn from a given multivariate normal distribution. In our simulation, we suppose that the "true" return-generating process for the US stock market is governed by the multivariate normal distribution from the Barra USE4 asset covariance matrix.[1] We then use the Cholesky decomposition to simulate 240 periods of stock returns (representing 20 years of monthly observations) for all constituents of the MSCI USA IMI index.

The first step in our simulation exercise is to construct a pure factor portfolio for a given USE4 style factor. This is accomplished in the usual way by performing a multivariate cross-sectional regression of stock returns against the factor exposures. To compute the actual return of the pure factor portfolio, we must provide the simulated stock returns as inputs.

To illustrate the various factor types, we generate three distinct sets of simulated stock returns. The first set of simulations is designed to illustrate the behavior of a risk factor that is not an alpha factor. These stock returns were generated using the Cholesky decomposition with the factor returns and specific returns drawn from mean-zero distributions. The second set of simulations illustrates the behavior of a factor that is both an alpha factor and a risk factor. In this case, the stock returns were generated using the same Cholesky decomposition but with the mean of the factor return distribution shifted slightly to produce a true Information Ratio of 1. The final set of simulations illustrates the comportment of an alpha factor that is not a risk factor. To generate these stock returns, we set the volatility of the factor under consideration to zero, while inducing a drift by making the mean specific returns proportional to the factor exposures. The constant of proportionality was calibrated to make the true Information Ratio equal to 1.

In Figure 1, we plot the cumulative performance of the pure factor portfolio using the three sets of simulated stock returns. The dashed red line

Not for Distribution

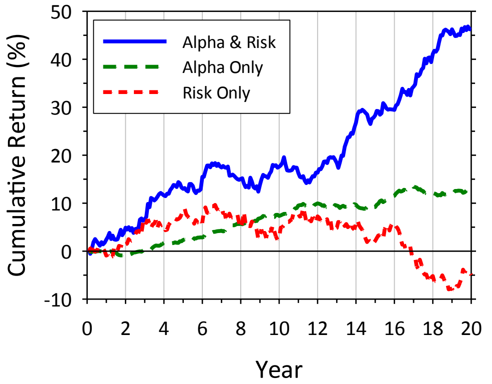

Figure 1 Simulated cumulative returns for different factor types.

represents the cumulative return of a pure factor portfolio corresponding to a risk factor that is not an alpha factor. We see that the factor portfolio tends to drift sideways, consistent with having zero expected returns. Even so, such a factor portfolio may, by chance, exhibit several years of positive or negative performance. The true *ex ante* volatility of this factor portfolio is 2.86 percent. Most of the volatility of the factor portfolio is due to the underlying factor, although there is also a small contribution from the stock-specific component.

The solid blue line in Figure 1 represents the cumulative return of a factor that is both an alpha factor and a risk factor. Over a long time, we in fact observe a persistent drift. Even so, we see that there were also extended periods with mediocre performance. That is, even skillful managers may sometimes underperform for periods spanning several years. The *ex ante* volatility of the factor portfolio is again 2.86 percent.

The dashed green line in Figure 1 represents the cumulative performance of an alpha factor that is not a risk factor. We see that the factor portfolio indeed exhibits positive drift over the long term. Note also that the volatility of the factor portfolio is much lower than in the other two cases. Even though the *true* factor return was zero every period, the *estimated* factor returns were always non-zero due to sampling error. The true *ex ante* volatility of this portfolio is 76 bps and is completely due to idiosyncratic effects.

In the next section, we describe the analytic framework that we employ in our study. Our framework is used to compare the efficiency of portfolios constructed using different risk models.

4 Analytic framework

We begin by supposing that the true alpha signal and the true asset covariance matrix are known in advance by the investor. As described by Grinold and Kahn (2000), it is a simple exercise to construct the holdings vector h_T of the true unconstrained optimal portfolio

$$h_T = \frac{V_T^{-1}\alpha}{\alpha' V_T^{-1}\alpha}, \tag{1}$$

where V_T is the $N \times N$ true asset covariance matrix and α is the $N \times 1$ vector of true stock alphas. Portfolio h_T has the lowest risk of any portfolio with unit exposure to the alpha factor, and hence the highest Information Ratio.

In the real world, of course, investors do not have the luxury of knowing either the true alpha signal or the true covariance matrix. The purpose of this paper is to compare the efficiency of optimized portfolios constructed using two different risk models: Model A which includes the alpha factor, and Model B which excludes it. Each model is estimated on two sets of simulated stock returns. In Case 1, the alpha factor contains true systematic risk; in Case 2 the alpha factor contains no systematic risk. While the two optimizations use different risk models, they utilize the same alpha signal. We take as our alpha signals the 12 style factors of the Barra USE4 risk model.

As before, we use the Cholesky decomposition to generate L periods of simulated stock returns

Not for Distribution

consistent with the true covariance matrix. For simulation purposes, we assume that the true covariance matrix is given by a multivariate normal distribution based on one of the two variants of the Barra USE4 risk model. The first variant, corresponding to Case 1, is the actual USE4 model without modification. The second variant, corresponding to Case 2, is derived from the USE4 model by simply deleting the alpha factor under consideration.

We then use the simulated stock returns to estimate two risk models. Model A is estimated using a factor exposure matrix X_A taken directly from the USE4 model. Model B is estimated using a factor exposure matrix X_B derived from USE4 by deletion of the alpha factor.

Let $r_{nt}^{(1)}$ be the simulated return for stock n and period t under Case 1,

$$r_{nt}^{(1)} = \sum_k X_{nk}^A f_{kt}^A + u_{nt}^A, \qquad (2)$$

where X_{nk}^A is the exposure of stock n to factor k taken from matrix X_A, f_{kt}^A is the true return for factor k during period t, and u_{nt}^A is the true specific return for the stock. The *estimated* factor returns $\hat{f}_{kt}^{A(1)}$ and estimated specific returns $\hat{u}_{nt}^{A(1)}$ under Model A are obtained by cross-sectional regression,

$$r_{nt}^{(1)} = \sum_k X_{nk}^A \hat{f}_{kt}^{A(1)} + \hat{u}_{nt}^{A(1)}. \qquad (3)$$

The estimated factor covariance matrix elements $\hat{F}_{ij}^{A(1)}$ for Model A are directly computed as

$$\hat{F}_{ij}^{A(1)} = \frac{1}{L-1} \sum_{t=1}^{L} (\hat{f}_{it}^{A(1)} - \bar{f}_i^{A(1)}) \times (\hat{f}_{jt}^{A(1)} - \bar{f}_j^{A(1)}), \qquad (4)$$

where $\bar{f}_i^{A(1)}$ is the time-series mean of the estimated returns for factor i and L is the number of periods used to estimate the model. Finally, the specific variance forecasts are given by

$$\hat{\Delta}_n^{A(1)} = \frac{1}{L-1} \sum_{t=1}^{L} (\hat{u}_{nt}^{A(1)} - \bar{u}_n^{A(1)})^2, \qquad (5)$$

where $\bar{u}_n^{A(1)}$ is the time-series mean of the estimated specific return for stock n. The estimated asset covariance matrix for Model A and Case 1 is therefore given by

$$\hat{V}_A^{(1)} = X_A \hat{F}_A^{(1)} X_A' + \hat{\Delta}_A^{(1)}, \qquad (6)$$

where $\hat{F}_A^{(1)}$ is the $K \times K$ estimated factor covariance matrix whose elements are given by Equation (4), and $\hat{\Delta}_A^{(1)}$ is the diagonal matrix of estimated specific variances whose elements are given by Equation (5).

After estimating Model A, we repeat the exercise for Model B, still using the same set of simulated returns from Case 1. This leads to a new asset covariance matrix, denoted $\hat{V}_B^{(1)}$. Next, we use the estimated risk models $\hat{V}_A^{(1)}$ and $\hat{V}_B^{(1)}$ to construct two optimal portfolios. The holdings of the portfolio constructed using Model A are given by

$$h_A^{(1)} = \frac{(\hat{V}_A^{(1)})^{-1}\alpha}{\alpha'(\hat{V}_A^{(1)})^{-1}\alpha}, \qquad (7)$$

with a corresponding expression for portfolio $h_B^{(1)}$, constructed using Model B.

Finally, we repeat this entire exercise using the simulated returns for Case 2. The returns in this case are given by

$$r_{nt}^{(2)} = \sum_k X_{nk}^B f_{kt}^B + u_{nt}^B. \qquad (8)$$

That is, we simulate stock returns by removing the systematic risk associated with the alpha factor. We use the simulated stock returns to estimate asset covariance matrices denoted $\hat{V}_A^{(2)}$ and $\hat{V}_B^{(2)}$. Finally, we use these covariance matrices to construct optimized portfolios, denoted $h_A^{(2)}$ and $h_B^{(2)}$.

Not for Distribution

The objective of this paper is to compare the efficiency of optimized portfolios constructed using different risk models. To study this question, we must first review the concept of the *Transfer Coefficient*.

5 The transfer coefficient

For clarity, the previous section used superscripts (1) and (2) to denote Case 1 and Case 2, respectively. In this section, for simplicity, we suppress such notation. The reader should be aware, however, that we consider both cases in our analysis.

To study the efficiency of portfolios h_A and h_B, we use the Transfer Coefficient. Let h_T be the true optimal portfolio, with holdings given by Equation (1). The true Information Ratio of portfolio h_T is given by

$$IR_T = \frac{E[R_T]}{\sigma_T}, \tag{9}$$

where $E[R_T]$ is the true expected return of portfolio h_T and σ_T is the true volatility. The expected returns are easily obtained using the true alphas, and the volatility is computed using true asset covariance matrix V_T.

In a similar fashion, the true Information Ratio for portfolio h_A is given by

$$IR_A = \frac{E[R_A]}{\sigma_A}, \tag{10}$$

where $E[R_A]$ is the true expected return of portfolio h_A and σ_A is the true volatility (measured using V_T). A central result of Modern Portfolio Theory is that the expected return of *any* portfolio is given by the true beta of the portfolio with respect to the true optimal portfolio h_T, multiplied by the expected return of portfolio h_T. In other words,

$$E[R_A] = \beta_A E[R_T]. \tag{11}$$

Substituting this into Equation (10), we obtain

$$IR_A = \frac{\beta_A E[R_T]}{\sigma_A}. \tag{12}$$

Now, using Equation (9) and the standard relation $\beta_A = \rho_A \sigma_A / \sigma_T$, we obtain

$$IR_A = \rho_A IR_T, \tag{13}$$

where ρ_A is the true correlation (i.e., computed using V_T) between portfolio h_A and portfolio h_T. Equation (13) states that the Information Ratio of portfolio h_A is equal to the Information Ratio of the true optimal portfolio h_T, multiplied by the predicted correlation between the two portfolios. This special correlation (i.e., with the true optimal portfolio) is given the name *Transfer Coefficient*,

$$TC_A = \frac{h'_A V_T h_T}{\sqrt{h'_A V_T h_A}\sqrt{h'_T V_T h_T}}. \tag{14}$$

A similar expression TC_B holds for portfolio h_B. The Transfer Coefficient quantifies the drop in Information Ratio due to holding the actual portfolio instead of the true unconstrained optimal portfolio. In other words, the Transfer Coefficient provides a direct measure of portfolio efficiency.

6 Simulated results

Every month we generated 200 periods of simulated stock returns with the Barra USE4 model using the two distinct return-generating processes given by Equations (2) and (8), respectively. We then used the simulated stock returns to estimate Model A and Model B. We computed the Transfer Coefficient for each of the 12 Barra USE4 style factors every month and averaged them over a 15-year period from December 1998 through December 2013. Finally, we repeated the simulation 10,000 times and averaged across the trial runs.

Not for Distribution

Table 1 Simulated mean Transfer Coefficient (averaged across time) for Barra USE4 optimized factor portfolios. Results are presented for Case 1, in which the alpha factor contains true systematic risk. Sample period is from Dec-1998 to Dec-2013.

	Sample Average	
Alpha Factor	TC(A)	TC(B)
Growth	0.85	0.88
Size	0.86	0.81
Non-linear Size	0.87	0.84
Dividend Yield	0.86	0.87
Book-to-Price	0.85	0.87
Earnings Yield	0.86	0.85
Beta	0.84	0.53
Residual Volatility	0.85	0.71
Non-linear Beta	0.86	0.85
Momentum	0.87	0.86
Leverage	0.86	0.84
Liquidity	0.84	0.80
Average (ex Beta)	0.86	0.83

In Table 1, we present mean Transfer Coefficients (averaged across time) for Case 1, in which the alpha signal contains true factor risk. By and large, we see that the Transfer Coefficients for the two models were quite similar. For a few factors (growth, dividend yield, and book-to-price), Model B actually produced marginally higher Transfer Coefficients, *even though the model was misspecified*. For most factors, however, Model A produced slightly more efficient portfolios. The major exception was the beta factor, which saw a major increase in Transfer Coefficient when it was included in the risk model.

It is worth considering the beta factor in more detail. The beta factor is unique among style factors in that it forms a near-perfect hedge with another risk factor. More specifically, the return correlation in USE4 between the beta factor and the country factor tends to be extremely high, typically around 0.90. To understand the origin of this correlation we must consider the pure factor portfolios. As discussed by Menchero (2010), the country factor portfolio essentially represents the cap-weighted US market portfolio. The beta factor portfolio, by contrast, is a dollar-neutral portfolio that is long high-beta stocks and short low-beta stocks. Over the short run, when the return of the market portfolio is positive, there is a strong tendency for high-beta stocks to outperform low-beta stocks. This explains the high correlation between the two factors. Hence, a short position in the country factor will act as a near-perfect hedge for the beta factor. However, if the beta factor is omitted from the model, then it is impossible to exploit this hedge. This leads to the large drop in Transfer Coefficient seen for Model B in Table 1.

In Figure 2, we plot the mean Transfer Coefficient (averaged across factors) versus time for Case 1. The average Transfer Coefficient for Model A was quite stable. For Model B, by contrast, we find considerably more time variation. In particular, it appears that there was greater benefit to including the alpha factor during the Internet Bubble and the Global Financial Crisis. This is likely because

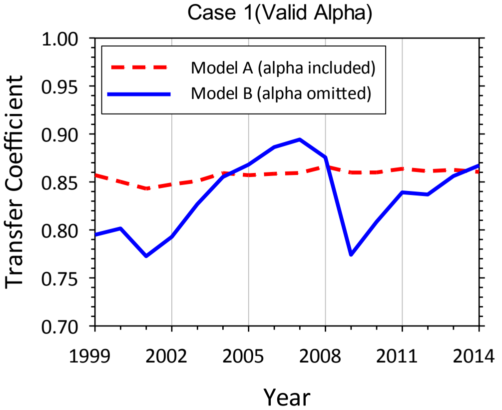

Figure 2 Simulated mean Transfer Coefficient (averaged across factors) and plotted versus time for Case 1, in which the alpha factor contains true systematic risk.

Not for Distribution

Table 2 Simulated mean Transfer Coefficient (averaged across time) for Barra USE4 optimized factor portfolios. Results are presented for Case 2, in which the alpha factor contains no systematic risk. Sample period is from Dec-1998 to Dec-2013.

	Sample Average	
Alpha Factor	TC(A)	TC(B)
Growth	0.79	0.99
Size	0.78	0.98
Non-linear Size	0.79	0.98
Dividend Yield	0.78	0.98
Book-to-Price	0.79	0.98
Earnings Yield	0.80	0.98
Beta	0.78	0.97
Residual Volatility	0.77	0.98
Non-linear Beta	0.81	0.99
Momentum	0.80	0.98
Leverage	0.79	0.98
Liquidity	0.78	0.98
Average (all factors)	0.79	0.98

the USE4 factor correlations were higher during these crisis periods, making the factor hedges more effective.[2]

In Table 2, we report the mean Transfer Coefficient (averaged across time) for Case 2, in which the alpha factor contains no systematic risk. The returns of the pure alpha factor portfolio are now 100 percent idiosyncratic in nature. In other words, the estimated alpha factor returns in Model A are pure noise. We now find an average Transfer Coefficient of 0.79 by improperly including the factor (Model A) versus an average Transfer Coefficient of 0.98 by properly excluding the alpha factor from the risk model (Model B). In this case, including a spurious factor in the risk model significantly reduces portfolio efficiency. This is a consequence of the fact that the *in-sample* correlations between the alpha factor and the other risk factors are always non-zero. The optimizer interprets these in-sample correlations as hedging opportunities, thereby producing non-zero exposures to risk factors. However, since the *true* correlations are zero, the exposures to the risk factors *add* risk to the portfolio, rather than reducing it.

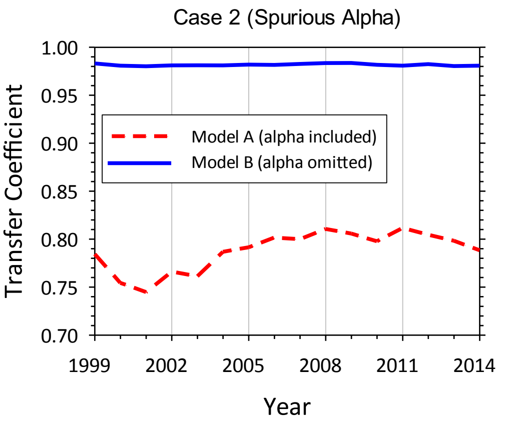

Figure 3 Simulated mean Transfer Coefficient (averaged across factors) and plotted versus time for Case 2, in which the alpha factor contains no systematic risk.

In Figure 3, we plot the mean Transfer Coefficient (averaged across factors) versus time for Case 2. The mean Transfer Coefficient for Model B was consistently high and quite stable. For Model B, the optimizer focuses on diversifying specific risk since the alpha factor does not contain any systematic risk that needs to be hedged. The high Transfer Coefficient for Model B implies that 200 periods of observations are sufficient to effectively diversify specific risk. For Model A, by contrast, the mean Transfer Coefficient is consistently lower due to the spurious factor hedges described above.

7 The role of sampling error

The results of the previous section were based on using 200 periods to estimate the risk models. Since our simulation environment is stationary, we always expect to find improvement by using

Not for Distribution

more observations. The amount of sampling error in the factor covariance matrix, of course, depends on the number of observations in the sample. To study the role of sampling error, we repeat the same exercise as before except we now vary the estimation window length from 100 periods to 800 periods. For each length of estimation window, we again perform 10,000 simulations. Due to the computational expense of the simulations, however, we consider only a single analysis date of December 31, 2013.

We first consider Case 1, in which the alpha factor contains true systematic risk. In Figure 4, we plot the mean Transfer Coefficient for Model A and Model B (averaged across factors) versus length of estimation window. For Model B, the Transfer Coefficient does not materially depend on the length of the estimation window. This is because the residual alpha component receives most of the risk budget, hence factor correlations play little role in determining the Transfer Coefficient. By contrast, for Model A, we see a large increase in average Transfer Coefficient when we expand the length of the estimation window. This is due to more reliable factor hedges as the correlations are estimated with greater precision.

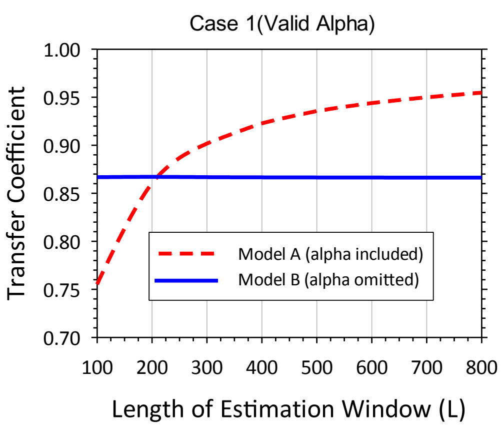

Figure 4 Simulated mean Transfer Coefficient (averaged across factors) versus length of estimation window for Case 1. The analysis date is December 31, 2013.

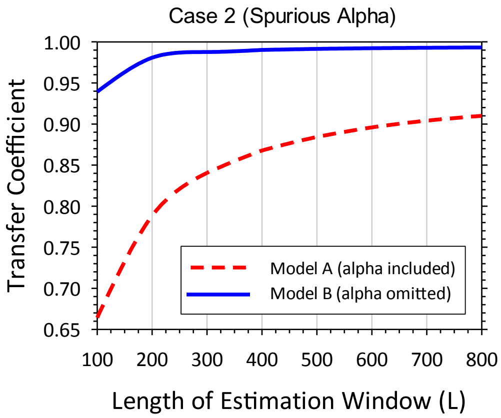

Figure 5 Simulated mean Transfer Coefficient (averaged across factors) versus length of estimation window for Case 2. The analysis date is December 31, 2013.

In Figure 5, we plot the mean Transfer Coefficient versus length of estimation window for Case 2, in which the alpha factor contains no systematic risk. The average Transfer Coefficient of Model B is close to 1 and dominates Model A for any length of estimation window. However, the Transfer Coefficient of Model A increases dramatically as we expand the length of the estimation window. This is because with more observations, Model A is able to estimate correlations with sufficient precision to avoid placing spurious hedges.

8 The long-only constraint

Up to now, we have only considered unconstrained optimal portfolios, which are determined analytically using Equation (1). In practice, portfolio managers often impose several investment constraints, such as the long-only constraint, the full-investment constraint, monthly turnover constraints, or constraints on the number of stocks held in the portfolio. In this paper, we focus on the long-only constraint. If the true risk model is used in the portfolio construction process, any constraint will lower the *ex ante* Information Ratio by reducing the ability of the optimizer to diversify specific risk or find effective factor hedges.

Not for Distribution

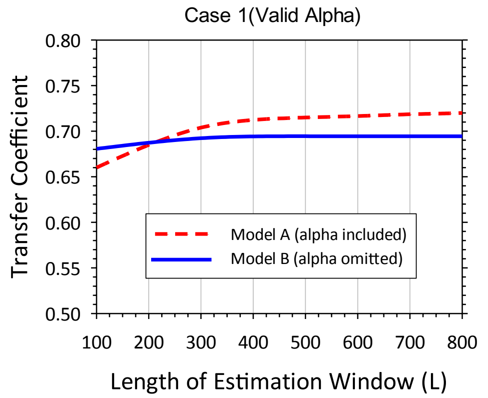

Figure 6 Simulated mean Transfer Coefficient (averaged across factors) versus length of estimation window for Case 1, with long-only constraint. The analysis date is December 31, 2013.

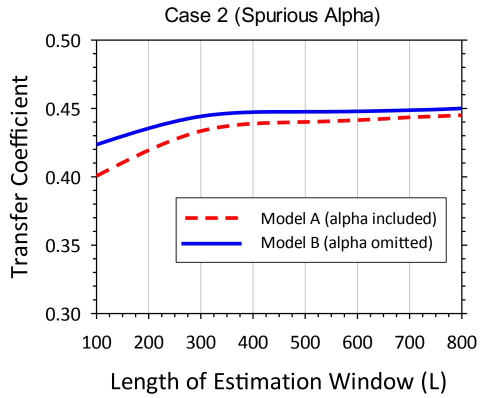

Figure 7 Simulated mean Transfer Coefficient (averaged across factors) versus length of estimation window for Case 2, with long-only constraint. The analysis date is December 31, 2013.

In Figure 6, we plot the mean Transfer Coefficient versus length of estimation window for Case 1 under the long-only constraint. Again we perform 10,000 simulations for analysis date of December 31, 2013. Compared with Figure 4, we see that the long-only constraint has the effect of lowering the Transfer Coefficient, as expected. The results in Figures 4 and 6 are also qualitatively similar in that Model A outperforms for large window length, whereas Model B outperforms for short window lengths. What is striking, however, is the tight compression observed in Figure 6 for the *difference* in Transfer Coefficient between Model A and Model B. That is, whereas the differences in Transfer Coefficient between Model A and Model B are quite significant in Figure 4, the differences are much smaller when the long-only constraint is imposed. For example, even with 800 observations in Figure 6, Model A has only a slightly higher Transfer Coefficient than Model B (0.72 versus 0.69). In other words, even when the alpha factor is a valid risk factor, we obtain only a 4 percent improvement in Information Ratio under the long-only constraint by including alpha in the risk model.

In Figure 7, we plot the mean Transfer Coefficient versus length of estimation window for Case 2 under the long-only constraint. Again, we performed 10,000 simulations for analysis date of December 31, 2013. Three points are worth highlighting. First, the long-only constraint dramatically lowers portfolio efficiency for Case 2. For instance, the Transfer Coefficient for Model B at 800 observations drops from 0.99 in Figure 5 to 0.45 in Figure 7. Second, we again find that the curves under the long-only constraint (Figure 7) are qualitatively similar to the unconstrained result (Figure 5), with Model B dominating across the entire spectrum. Lastly, we see that the *differences* in Transfer Coefficients have again narrowed considerably. For example, using 100 periods, Model B has a Transfer Coefficient of 0.94 in Figure 5 versus 0.66 for Model A. In Figure 7, by contrast, we see that this sizeable difference is now cut to 0.42 for Model B versus 0.40 for Model A. That is, we find only a modest gain in Information Ratio under the long-only constraint when the alpha factor has been properly omitted from the model.

9 Empirical analysis

Simulation exercises are extremely useful for investigating model behavior as they allow the

Not for Distribution

researcher to "peek behind the curtain" in a controlled environment and observe the *true* characteristics of the portfolios. Effects related to sampling error, therefore, can be studied by varying the length of the estimation window. Furthermore, since simulations can be repeated thousands of times, we are able to obtain high statistical confidence in the results.

Of course, simulations are subject to their own limitations. No matter how realistic they may be, simulations must always make idealized assumptions that do not quite reflect reality. For instance, in our simulations, we assumed a stationary return-generating process. Real financial markets, however, are not stationary. The skeptical reader may thus question the validity of our simulated results.

To investigate this question, we conducted the following experiment using the Barra USE4 risk model. We treat each USE4 style factor as an alpha factor, one at a time. Since the USE4 factors were specifically selected due to their strength as risk factors, we regard them all as having true systematic risk. We then consider two risk models. The first is the actual USE4 risk model, denoted Model A since it includes the alpha factor. The second model is a modified version of the USE4 model, which is obtained by simply setting the volatility of the alpha factor under consideration to zero.[3] This is denoted Model B since the alpha factor has been effectively omitted from the risk model.

We then form the unconstrained optimal portfolios, denoted h_A and h_B, using Model A and Model B, respectively. The investment universe is limited to the MSCI USA IMI index. Since portfolios h_A and h_B, by construction, have the same unit exposure to alpha, the portfolio with lower volatility will have the higher Information Ratio. In fact, it is simple to show that the ratio of out-of-sample volatilities is inversely proportional to the ratio of corresponding Transfer Coefficients,

$$\frac{TC_A}{TC_B} = \frac{\sigma_B}{\sigma_A}. \tag{15}$$

Table 3 Empirical out-of-sample volatility ratio for USE4 unconstrained optimized factor portfolios. The sample period is from Dec-1998 to Dec-2013.

	Realized Risk (%)		Volatility
Alpha Factor	Model A	Model B	Ratio
Growth	1.56	1.74	1.11
Size	2.57	2.82	1.10
Non-linear Size	2.76	2.59	0.94
Dividend Yield	1.76	1.79	1.02
Book-to-Price	1.82	1.90	1.04
Earnings Yield	2.79	3.05	1.09
Beta	3.09	5.90	1.91
Residual Volatility	2.66	3.44	1.30
Non-linear Beta	1.36	1.71	1.26
Momentum	6.26	6.43	1.03
Leverage	1.64	2.16	1.32
Liquidity	1.76	1.96	1.11
Average (ex Beta)	2.45	2.69	1.11

In Table 3, we report the out-of-sample volatilities σ_A and σ_B of the unconstrained optimal portfolios for each style factor in the USE4 model. The sample period is from December 1998 through December 2013. Empirically, the beta factor was again an outlier due to the near-perfect hedge with the USE4 country factor. The out-of-sample volatility of beta using Model B was 5.9 percent, versus 3.1 percent for Model A. All factors except non-linear size had higher volatility—hence, lower efficiency—when the factor was omitted from the model. The average volatility ratio across all factors, excluding beta, was 1.11.

In order to make a quantitative comparison with the simulated results, we must first estimate the effective number of observations in the USE4 model. The USE4 model uses daily factor returns, with factor correlations estimated

Not for Distribution

using exponentially weighted averages with a two-year half-life (504 trading days). The USE4 model also uses the methodology of Newey and West (1987) with two lags for treating serial correlation in the factor returns. Without lags, the effective number of observations would be approximately three times the half-life parameter, or roughly 1500. However, the Newey–West procedure, while reducing biases in the correlation forecasts, also increases the amount of sampling error. To a first approximation, the Newey–West procedure with two lags is equivalent to aggregating returns over a three-day horizon. This reduces the effective number of observations by a factor of 3. Hence, the effective number of observations in USE4 is approximately 500.

Given that our simulations suggest we always benefit by increasing the estimation window length, it is natural to ask why the USE4 model does not simply use the longest possible estimation window. The answer lies in the fact that real financial markets are non-stationary, in contrast to the simulations. To minimize sampling error, a long estimation window is desirable. However, long windows include much stale data that have little to do with current market conditions. To capture the most relevant data, using a short estimation window is desirable. The ideal window length is found by optimally balancing these two effects.

From Figure 4, with a window length of 500 periods, we find the Transfer Coefficient using Model A is 0.94 versus 0.87 for Model B. Using Equation (15), therefore, the simulations suggest a volatility ratio of 1.08. This is in excellent agreement with the observed volatility ratio of 1.11 shown in Table 3.

Next, we repeat the entire exercise under the long-only constraint. In Table 4, we present the out-of-sample volatilities for each factor. The mean volatility ratio, averaged across factors, was

Table 4 Empirical out-of-sample volatility ratio for USE4 optimized factor portfolios, with long-only constraint. The sample period is from Dec-1998 to Dec-2013.

	Realized Risk (%)		Volatility
Alpha Factor	Model A	Model B	Ratio
Growth	2.94	3.21	1.09
Size	5.15	5.15	1.00
Non-linear Size	8.98	8.93	0.99
Dividend Yield	2.65	2.75	1.04
Book-to-Price	3.06	3.09	1.01
Earnings Yield	4.69	4.64	0.99
Beta	6.42	6.78	1.06
Residual Volatility	4.49	5.16	1.15
Non-linear Beta	3.42	3.54	1.04
Momentum	6.73	6.79	1.01
Leverage	2.24	2.38	1.06
Liquidity	4.57	4.79	1.05
Average (all factors)	4.45	4.58	1.04

1.04. Now, comparing with the simulated results in Figure 6, we see that Model A has a Transfer Coefficient of 0.72 versus 0.69 for Model B at 500 periods. Hence, the simulations predict a volatility ratio of 1.04, in essentially perfect agreement with empirical observations.

A final point worth highlighting is that the volatility ratio for the beta factor was 1.06 under the long-only constraint (Table 4) versus 1.91 in the unconstrained case (Table 3). The reason for this major reduction in volatility ratio is that the long-only constraint forces the active exposure to the country factor to be zero. As a result, the beta factor is no longer able to exploit the hedge that the country factor provides.

10 Sources of factor misalignment

The degree of factor misalignment is related to the relative magnitudes of the spanned and residual

Not for Distribution

components of alpha. When the residual component dominates, the alpha factor is weakly collinear with the risk factors. Conversely, when the spanned component dominates, the alpha factor is strongly collinear with one or more risk factors.

Weak collinearity typically occurs when the alpha signal represents an attribute that is fundamentally distinct from any of the risk model factors. In our study, we considered the alpha signals to be the style factors from the Barra USE4 model. Since each Barra risk factor is designed to capture a distinct fundamental attribute, these factors tend to be weakly collinear.

Strong collinearity typically occurs when the alpha signal represents essentially the same attribute as one of the risk model factors. Momentum is a leading example of such a factor. Investors who follow a momentum strategy will typically use a proprietary definition of momentum as the alpha signal. These proprietary signals are likely to be very similar, but not identical, to the Barra USE4 momentum factor.

Based on the foregoing discussion, the reader will recognize that our paper was limited to the case of weak collinearity. It is possible that factor misalignment effects are more pronounced in the case of strong collinearity. One reason for this is that as the degree of collinearity increases, the residual alpha component becomes less stable, leading to possible increases in portfolio turnover. Further research is required to better understand these effects.

11 Treatments for factor misalignment

We now briefly discuss several techniques employed by some practitioners for treating factor misalignment in the case of strong collinearity. The first treatment is to simply *delete* the collinear factor from the risk model. For example, in the case of a momentum strategy, the Barra momentum factor would be zeroed out from both the factor exposure matrix and the factor covariance matrix. This essentially transforms the problem into one of weak collinearity. Although this technique is easy to apply, it suffers from two significant drawbacks. First, since the resulting portfolio is closely aligned with the missing factor, the risk model may significantly underpredict portfolio volatility. Second, with the factor deleted from the model, there is no mechanism for the optimizer to enhance portfolio efficiency by exploiting correlations with other risk factors. This results in a sub-optimal allocation of the risk budget.

Yet another approach for treating factor misalignment is to build a custom risk model that contains the alpha factor. This ensures consistency between the factor covariance matrix and the factor exposure matrix. Nevertheless, more research is necessary to understand the full implications of this approach.

12 Summary

We have investigated the effects of factor misalignment on portfolio construction. We considered the effect of using two distinct risk models, one including the alpha factor, the other excluding it. We also considered two distinct cases. In the first case, the residual alpha component contained true systematic risk, whereas it had no factor risk in the second case. We found that even when the residual alpha contained true factor risk, it did not necessarily imply that including the alpha factor in the risk model would increase the Information Ratio. This occurs when the correlations are estimated with low precision due to insufficient observations. However, as the length of the estimation window increases, we found a benefit to including the alpha factor in the risk model. When the residual alpha did not contain true factor risk,

Not for Distribution

we found a consistent reduction in the Information Ratio by improperly including the alpha factor in the risk model.

We also examined the impact of the long-only constraint. As expected, we found that this constraint reduced portfolio efficiency. More interesting was the finding that the differences in portfolio efficiency narrowed considerably with the introduction of the long-only constraint. More specifically, the Transfer Coefficients were nearly identical regardless of which risk model was used to construct the portfolios.

In addition, we conducted empirical back-tests in the US market to study the gain in Information Ratio that could be achieved by including alpha factors in the risk model. Using the Barra USE4 risk model, we found an average increase in Information Ratio of approximately 11 percent if no constraints were imposed. With the long-only constraint, however, the gain in Information Ratio dropped to 4 percent. Both empirical findings were in excellent agreement with the simulated results.

Finally, we discussed sources of factor misalignment. Depending on whether or not the alpha signal represents an attribute already captured by one of the risk factors, the factor misalignment can be classified as either weakly collinear or strongly collinear. We concluded with a brief discussion of several techniques employed by practitioners to treat the case of strong collinearity.

Notes

[1] The Barra USE4 model comes in two versions: short horizon and long horizon. Throughout this paper, we use the short-horizon version, which is calibrated for a 1-month prediction horizon.

[2] These higher correlations may partially be an artifact of sampling error. Since sampling error tends to be higher during volatile periods, it is more likely to result in large in-sample correlations.

[3] Note that removing a factor will cause all of the other factor returns to change. We also performed the more rigorous technique of projecting the factor covariance matrix to the new basis with one factor deleted. However, we found no material difference between the two results. Consequently, in this paper, we adopt the simpler method of simply setting the factor volatility to zero.

References

Ceria, S., Saxeena, A., and Stubbs, R. (2012). "Factor Alignment Problems and Quantitative Portfolio Management," *Journal of Portfolio Management* **38**(2) (Winter), 29–43.

Greene, W. (2000). *Econometric Analysis* (Fourth Edition). Prentice Hall.

Grinold, R. and Kahn, R. (2000). *Active Portfolio Management* (Second Edition). McGraw-Hill.

Lee, J. and Stefek, D. (2008). "Do Risk Factors Eat Alphas?" *Journal of Portfolio Management* **34**(4) (Summer), 12–25.

Markowitz, H. (1952). "Portfolio Selection," *Journal of Finance* **7**(1), 77–91.

Menchero, J. (2010). "The Characteristics of Factor Portfolios," *Journal of Performance Measurement* **15**(1) (Fall), 52–62.

Newey, W. and West, K. (1987). "A Simple, Positive Semi-Definite, Heteroskedasticity and Auto-Correlation Consistent Covariance Matrix," *Econometrica* **55**(3), 703–708.

Keywords: Factor misalignment; portfolio construction; portfolio optimization

Not for Distribution