www.joim.com

Journal Of Investment Management, Vol. 13, No. 3, (2015), pp. 57–79

AUGMENTED RISK MODELS TO MITIGATE FACTOR ALIGNMENT PROBLEMS

Anureet Saxena and Robert A. Stubbs†*

Construction of optimized portfolios entails a complex interaction between three key entities, namely, the risk factors, the alpha factors and the constraints. The problems that arise due to mutual misalignment between these three entities are collectively referred to as Factor Alignment Problems (FAP). Examples of FAP include risk underestimation of optimized portfolios, undesirable exposures to factors with hidden and unaccounted systematic risk, consistent failure in achieving ex-ante performance targets, and inability to harvest high quality alphas into above-average IR. In this paper, we give a detailed analysis of FAP and discuss solution approaches based on augmenting the user risk model with a single additional factor y. For the case of unconstrained mean–variance optimization (MVO) problems, we develop a generic analytical framework to analyze the ex-post utility function of the corresponding optimal portfolios, derive a closed-form expression of the optimal factor volatility value and compare the solutions for various choices of y culminating with a closed-form expression for the optimal choice of y. Augmented risk models not only correct for risk underestimation bias of optimal portfolios but also push the ex-post efficient frontier upward thereby empowering a portfolio manager (PM) to access portfolios that lie above the traditional risk–return frontier. We corroborate our theoretical results by extensive computational experiments, and discuss market conditions under which augmented risk models are likely to be most beneficial.

1 Introduction

Factor models play an integral role in quantitative equity portfolio management. Their applications extend to almost every aspect of quantitative investment methodology including construction of alpha models, risk models, portfolio construction, risk decomposition, and performance attribution. Given their pervasive presence in the field and the natural trend toward specialization, it comes as no surprise that different groups of researchers are often involved in developing factor models for each one of the aforementioned applications.

For instance, a team of quantitative portfolio managers (PM) can develop an in-house model for

*Vice President, Analyst, Allianz Global Investors, 600 West Broadway, San Diego, CA 92101, USA.

†Vice President, Research, Axioma, Inc., 400 Northridge Road, Suite 850, Atlanta, GA 30350, USA.

alpha generation, and procure a factor model for the purposes of risk management from a third-party risk model vendor. Subsequently, they can combine the two models within the framework of Markowitz mean–variance optimization (MVO) framework to construct optimal portfolios. A completely different factor model can then be used for the purposes of performance attribution to identify the key drivers and detractors of performance. Notably, the choice of factors in each one of these factor models need not be identical, thereby introducing incongruity in the portfolio management process. To further complicate the matters, the constraints in the quantitative strategy can introduce additional systematic risk exposures that are not captured by the risk model. The problems that arise due to the interaction between the alpha model, the risk model, and constraints in an MVO framework are collectively referred to as Factor Alignment Problems (FAP). Detailed theoretical investigation of FAP leading to a solution methodology in the form of augmented risk models constitute the emphasis of this paper. Next we give a brief survey of existing research on FAP, and highlight our key contributions.

The primary purpose of portfolio optimization is to create a portfolio having an optimal risk–return tradeoff. If a portion of systematic risk exposure of the portfolio is inadequately captured by the risk model then the resulting portfolio cannot be expected to be optimal ex-post, its ex-ante optimality notwithstanding. In other words, FAP symbolize the difficulties that a PM faces in ensuring the ex-post optimality of a portfolio that is deemed to be optimal ex-ante in the MVO framework. Examples of FAP include risk underestimation of optimized portfolios, undesirable exposures to factors with hidden and unaccounted systematic risk, consistent failure in achieving ex-ante performance targets, and inability to harvest high quality alphas into above-average IR.

Several authors have examined FAP recently and have proposed various solution techniques. Saxena and Stubbs (2013) conducted an empirical case study to understand the risk underestimation problem, a prominent symptom of FAP. The authors used real-world data and a battery of backtests to demonstrate the perverse and pervasive nature of FAP. They demonstrated that all optimized portfolios share a common property, namely, they have exposure to certain kinds of latent systematic risk factors that are uncorrelated with factors of the risk model that was used to generate them. Ceria *et al.* (2012) examine potential sources of the mentioned systematic risk factors and suggest that proprietary definitions of certain style (B/P, E/P, etc.) and technical factors can introduce them. Lee and Stefek (2008) illustrate a similar idea by using two different definitions of a momentum factor to define alpha and risk factors, and argue that the optimizer is likely to load up on the difference between the two thereby taking unintended bets. Saxena and Stubbs (2012) discuss a detailed empirical case study on FAP using the USER model (see Guerard *et al.*, 2012a); among other things, they quantify the portion of unaccounted systematic risk that can be attributed to the constituent factors of the USER model and constraints present in the strategy. Unlike previous studies which have investigated FAP from an empirical standpoint, we pursue a theoretical exploration of this topic. We go back to the roots of mean–variance optimization, and demonstrate analytically the tendency of the optimizer to adversely exploit inconsistencies between the alpha and risk models, thereby compromising the efficiency of the resulting portfolios. Our analysis not only yields diagnostic tools to identify the presence of FAP, but also provides a natural remedy to FAP in the form of augmented risk models. The rest of this paper is organized as follows.

Section 2 discusses a prototypical quantitative strategy with the aim of identifying some of the common symptoms of FAP, such as the risk underestimation problem, and undesirable and unintended exposure to systematic risk factors which are not captured by the risk model. The aim of this section is to provide a simple but comprehensive practical example which can be used to put the theoretical results presented in the later part of the paper in context. We return to the example discussed in Section 2 again in Section 7 wherein we show how our proposed methodology address FAP, and improves risk-adjusted returns. Section 3 lays out the theoretical model which is used in the rest of the paper. We discuss the assumptions we make in our analytical derivations, and provide theoretical and empirical justifications to them. Among other things, we introduce the notion of an augmented risk model which is used throughout the paper to remedy FAP. Given an arbitrary asset–asset covariance matrix Q, and a risk factor y, an augmented risk model is defined by the covariance matrix $Q_y = Q + \nu yy^{\mathrm{T}}$. In other words, if Q was derived from a factor model, Q_y is derived from an enhanced factor model that has all the original factors, and an additional augmenting factor y which is assumed to be uncorrelated with the original suite of risk factors.

In Section 4 we focus on the unconstrained MVO model, and develop a generic analytical framework to analyze the ex-post utility function of the corresponding optimal portfolios, derive a closed-form expression of the optimal factor volatility and compare the solutions for various choices of y culminating with a closed-form expression for the optimal choice of y. Among other things, we show that using an augmented risk model with an appropriately chosen volatility parameter ν not only solves the risk underestimation problem but also improves the ex-post utility function. The key result in this paper, referred to as the "Pushing Frontier Theorem" shows that using an augmented risk model shifts the ex-post efficient frontier upward, thereby allowing the PM to access portfolios that are not attainable using the traditional MVO approach. In Section 5, we extend these results to a constrained MVO model. We employ the concept of implied alpha to allow us to make a smooth transition from unconstrained MVO to its constrained counterpart. Recall that implied alpha is obtained by tilting the alpha in the direction of binding constraints, and acts as the de facto alpha for constrained MVO models.

Application of augmented risk models requires two key parameters, namely, the choice of an augmenting factor y and an estimate of its volatility ν. In Section 6, we discuss an exponentially weighted moving average (EWMA) volatility model to calibrate the volatility of augmenting factors. The emphasis in this section is on simplicity and practicality of the proposed approach, and we discuss a model that meets both of these criteria. Section 7 has a threefold emphasis. First, we give extensive empirical results that corroborate the theoretical findings discussed in the preceding sections. Second, we introduce the notion of "frontier spreads" to capture improvements in risk-adjusted returns that result due to the application of augmented risk models. Subsequently, we study the impact of various strategic parameters (turnover limits, asset bounds, etc.) on the frontier spreads, and also seek to identify market regimes where using an augmented risk model is most likely to yield significant improvements. Finally, we present computational results with a wide variety of alpha models to attest the robustness of the proposed approach. Section 8 concludes the paper with some closing remarks.

2 A practical active strategy

The focus of this section is twofold. First, we want to use a very simple value momentum strategy to illustrate a classic symptom of FAP, namely, the risk underestimation problem. Second, we show how the risk underestimation problem can be traced to certain hidden systematic risk factors that are not captured by the factor structure of the base risk model (BRM). Among other things, this sets the stage for the theoretical model discussed in the following section which assumes the existence of systematic risk factors missing from the BRM.

We used the following strategy in our experiments.[®,1]

$$\begin{array}{ll} maximize & Expected\ Return \\ s.t. & \\ & Fully\ invested\ long\text{-}only\ portfolio \\ & Active\ GICS^{®}\ sector\ exposure \\ & \quad constraint\ (\pm 20\%) \\ & Active\ GICS\ industry\ exposure \\ & \quad constraint\ (\pm 10\%) \\ & Active\ asset\ bounds\ constraint \\ & \quad (\pm 2\%) \\ & Turnover\ (two\text{-}way)\ constraint \\ & \quad (16\%) \\ & Active\ Risk\ constraint\ (\sigma\%) \\ & Benchmark = Russell^{®}\ 3000. \end{array}$$

We used a fundamental risk model as our BRM in defining the active risk constraint.[2] The expected returns were derived using an equal weighted combination of the BP variable in the USER model (Guerard *et al.*, 2012b) and the medium-term momentum factor in the BRM. We ran monthly backtests based on the above strategy in the 1999–2009 time period for various values of σ chosen from $\{1.0\%, 1.1\%, \ldots, 5.0\%\}$.

We use the notion of the bias statistic to identify statistically significant biases in risk prediction. If the ex-ante risk prediction is unbiased, then the bias statistic should be close to 1.0 (see Saxena and Stubbs, 2013 for more details). A bias statistic value which is significantly above (below) 1.0 indicates downward (upward) biases in risk prediction. Figure 1 reports the bias statistics of the portfolios for various risk target levels. Clearly, the bias statistics are significantly above the 95% confidence interval, thereby confirming the statistical significance of the downward bias in predicted risk estimates. We next focus on optimal portfolios that were generated when a risk target of 3.0% was employed.

At $\sigma = 3.0\%$, the optimal portfolios had realized active risk of 3.81%. The bias statistic for these portfolios was 1.27 which clearly lies outside the 95% confidence interval $[0.87, 1.13]$. Figure 2 further corroborates this phenomenon by showing the time series of realized risk of the optimal portfolios computed using 24 period realized returns on a rolling horizon basis; we also show the predicted risk of the portfolios

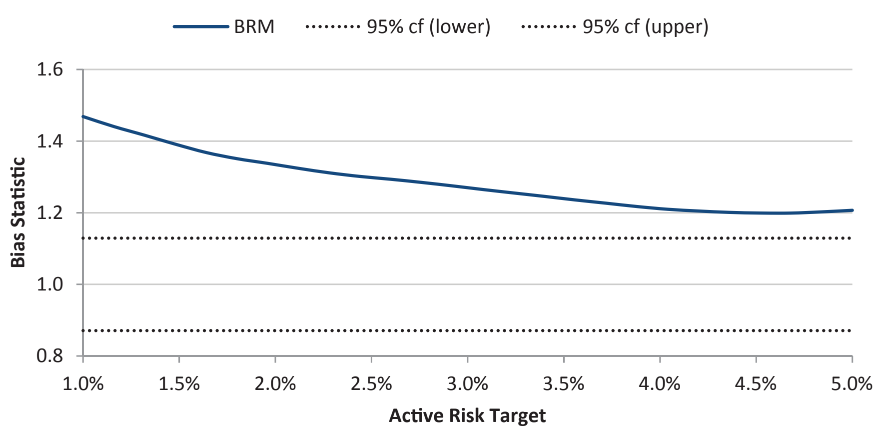

Figure 1 Bias statistic (active risk).

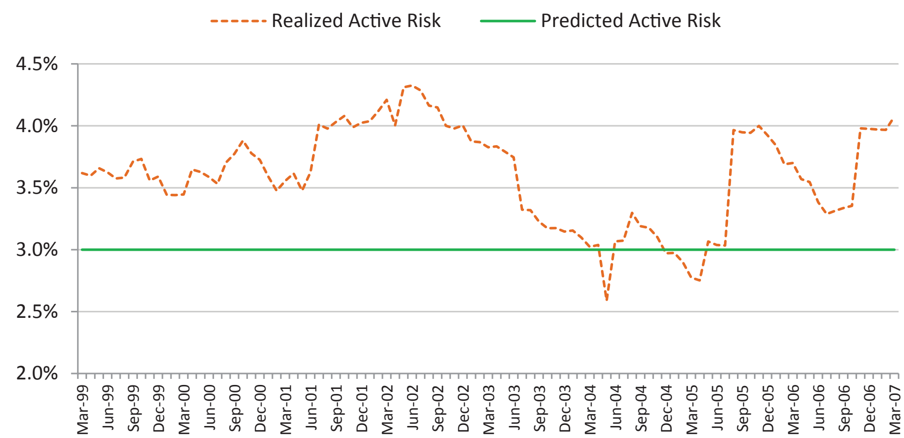

Figure 2 Time series of realized (24-period rolling) and predicted risk of optimal portfolios constructed using our BRM.

for the sake of comparison. While the degree of under-prediction might have varied, the realized risk was consistently above the predicted risk in most of the periods. It is tempting to believe that the risk models used in construction of the optimized portfolios were themselves biased, and the risk underestimation problem is simply an artifact of the bias in the BRM. Saxena and Stubbs (2013) examined this issue and demonstrated that the base risk model, in fact, produces unbiased risk estimates for random portfolios. Consequently, the bias depicted in Figure 1 is peculiar to optimized portfolios.

We introduce two additional concepts to assist us in tracing the sources of the risk underestimation problem. Let X denote the $n \times m$ exposure matrix associated with the BRM; n denote the number of assets; and m denote the number of factors in the BRM. Given an arbitrary factor α, consider the following linear regression model that regresses α against factors in the base risk model,

$$\alpha = Xu + \alpha_{\perp};$$

the residual $\alpha_{\perp}$ in the above regression model is referred to as the *orthogonal* component of α, whereas $\alpha_X = Xu$ is referred to as the *spanned* component of alpha. Mathematically,

$$\alpha_X = X(X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}\alpha; \quad \alpha_{\perp} = \alpha - \alpha_X.$$

By virtue of being uncorrelated with all the factors included in the BRM, $\alpha_{\perp}$ has no systematic risk exposure with respect to the BRM; in other words, the BRM assumes that $\alpha_{\perp}$ has only idiosyncratic risk. This assumption can be problematic if certain systematic risk factors are missing from the BRM and $\alpha_{\perp}$ has overlap with some of the missing systematic risk factors. As we will soon discover, this indeed turns out to be the case, thus providing a key insight into the risk underestimation problem.

Next we introduce the notion of augmented regressions which can be used to determine if the orthogonal component of a given factor has overlap with systematic risk factors missing from the BRM. Consider a linear regression model that regresses asset returns against factors in the BRM, represented by the matrix X, and the normalized orthogonal component $y = \frac{1}{\|\alpha_{\perp}\|}\alpha_{\perp}$ of α. If $\alpha_{\perp}$ was truly idiosyncratic in nature then the factor returns associated with y, denote by f_y, in the above regression model should not be statistically significant. Alternatively, if f_y is indeed statistically significant and has nontrivial volatility then we can be assured that $\alpha_{\perp}$ has overlap with systematic risk factors missing from the base risk model.

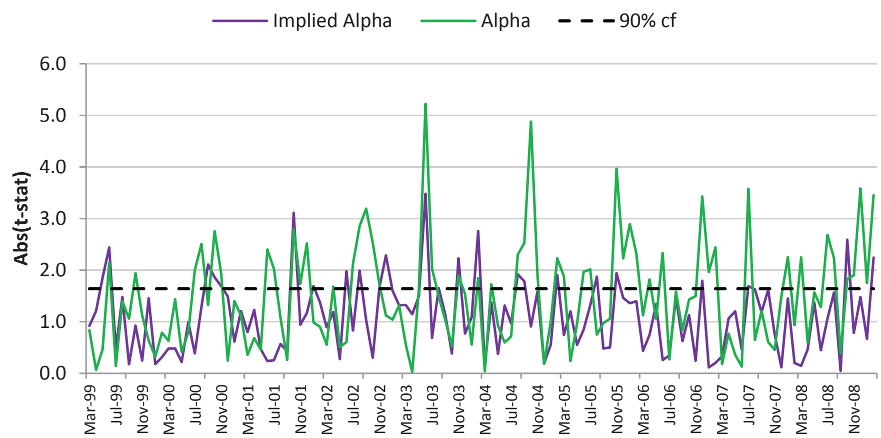

Figure 3 Time series of t-statistics (absolute values) associated with the orthogonal component.

Figure 3 shows the time series of (absolute) t-statistics associated with the orthogonal component of alpha (α) and implied alpha ($\tilde{\alpha}$) in the corresponding augmented regression model; $\alpha_{\perp}$ ($\tilde{\alpha}_{\perp}$) was found to be statistically significant (90% cf) in 40% (20%) of the periods. Given that a median factor in our BRM is statistically significant in about 20–30% of the periods, these statistics imply that $\alpha_{\perp}$ and $\tilde{\alpha}_{\perp}$ are as significant as half of the factors in the BRM. Figure 4 reports the annualized volatility of factor returns associated with $\alpha_{\perp}$ and $\tilde{\alpha}_{\perp}$ computed using a rolling 24-period window. As evident from the chart, not only are $\alpha_{\perp}$ and $\tilde{\alpha}_{\perp}$ statistically significant, but their factor returns also exhibit significant volatility. To put these numbers in perspective, note that a median normalized[3] factor in the BRM has annualized volatility of roughly 30%.

All of these results indicate that the orthogonal component of α and $\tilde{\alpha}$ does carry a significant amount of systematic risk which is not accounted for during the process of portfolio construction. The section that follows builds on these observations; specifically we propose a theoretical model that explicitly accounts for systematic risk factors missing from the BRM, and use it to assess the marginal cost of FAP.

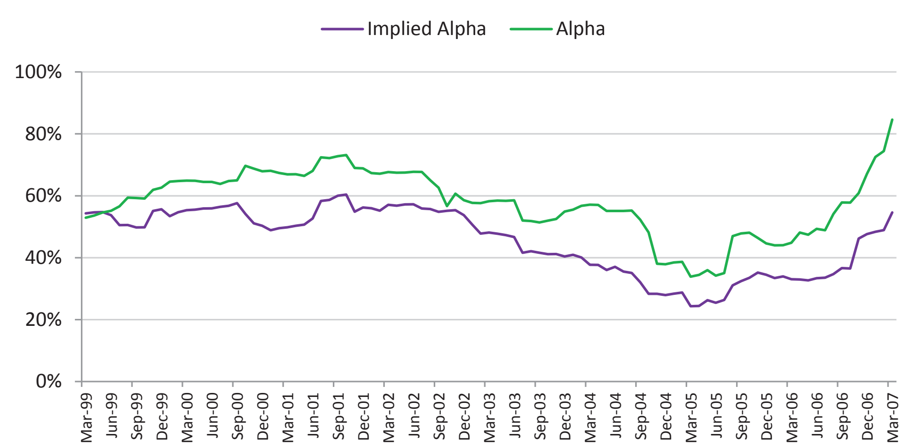

Figure 4 Time series of realized systematic risk of the orthogonal component computed using a rolling 24-period window.