QUANTITATIVE INVESTMENT MANAGEMENT

Risk Modelling in Optimized Factor Portfolios.

Practical research

Introduction

In a typical factor investing process, factor models feature heavily in alpha generation, risk management and performance attribution. A number of problems can arise from the discrepancies in the specification of factor models used in different stages of the factor investing process. Practical constraints faced by portfolio managers further complicate the matter.

Solution to these problems lies in customized risk modelling which augments generic risk models with return factors that are used in alpha generation, complements fundamental risk models with Principal Component Analysis, and applies suitable look-back period and weighting scheme that reflect the investment horizon and opportunity set of the factor investing strategy.

We first highlight the problem arising from the misalignment of return and risk factors with a mathematical representation of a typical optimization framework. To illustrate the point empirically, a case study on ex-ante risk decomposition is followed by a simulation of our Value strategy spanning more than 25 years of history.

Next we show the benefit of optimizing Momentum portfolio with risk model based on Principal Component Analysis, which is more adaptive to shifts in market regimes and results in tighter match between ex-ante and ex-post risk in simulation.

When it comes to deciding on the horizon and relative importance of historical data in estimating risk, we argue that they need to be customized to reflect the building blocks of the investment thesis. Comparing to optimization using generic risk models, a customized risk model with long-term look-back helps preserve the alpha of a Low Volatility portfolio that would otherwise have been completely eaten away by transaction costs.

Another area of customization lies in the weighting scheme used in Weighted Least Squares regressions that estimate factor returns. We propose a liquidity weighting scheme which origins from practical portfolio construction constraints and more realistically proxies for the theoretical weights to be used in the regression.

With theory backed by empirical evidence, we believe customized risk modelling offers practical solution to improve efficiency in factor portfolio construction.

Factor Alignment Problems

Misalignment in factor specification

In portfolio construction, the optimizer trades off between risk and return by rewarding a stock’s exposure to return factors and penalizing its contribution to risk. While most off-the-shelf multi-factor equity risk models have a wide range of style factors on top of market and industry factors, the specification of factors is usually different from that of the return factors used in alpha generation. Such a misalignment of factors is particularly problematic for factor investing strategies as it induces the optimizer to adversely exploit inconsistencies between the alpha and risk models, thereby compromising the efficiency of the resulting portfolios.

As summarized by Ceria et al. (2012) and Saxena and Stubbs (2015), 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. Examples of factor alignment problems include risk underestimation of optimized portfolios, undesirable exposures to factors with hidden and unaccounted systematic risk and consistent failure in achieving ex-ante performance targets.

For example, Lee and Stefek (2008) showed that the optimizer takes unintended bets by loading up on the difference between the definition of alpha and risk factors. When the momentum factor is constructed using price momentum from month T-13 to T-1 but the momentum factor in the risk model used past 12-month price momentum, the optimizer sees return but no factor risk in month T-13 and places a large bet. On the other hand, it sees factor risk but no return for month T-1 momentum and places a negative bet here. The combination of these two effects distorts the investment thesis.

To mitigate the problem of misalignment in factor specification, we replace a number of style factors in a generic Axioma fundamental risk model with our return factors such as Value, Quality, Momentum and Low Volatility. Now, this customized fundamental risk model consists of 3 types of factors. The first type is return factors that are exactly the same ones to which the optimizer tries to maximize exposures. They are followed by other style factors from Axioma that are not part of alpha generation. Lastly we have other risk factors from the generic Axioma fundamental risk model, such as country and industry factors. The latter two types of factors are risk factors in the customized risk model as we assume they don't have expected alpha.

Optimization in the presence of return and risk factors

Here we illustrate why an optimizer using risk model that doesn't include return factors fails to control for factor risk.

Without differentiating between return and risk factors, an unconstrained optimization can be specified as: $\max(X^T\alpha)^T w - 0.5\lambda w^T\Sigma w$, where X is the factor exposure matrix, α is a vector representing the expected alpha of factorS, w is portfolio weights, λ is risk aversion parameter, and Σ is the asset covariance matrix.

The optimal weight w^* can be solved as: $w^* = \frac{1}{\lambda}\Sigma^{-1}X^T\alpha$, which can be understood as a stock's exposure to alpha scaled down by its risk.

The above representation of portfolio construction is overly simplistic. Not every factor that helps explain stock return variance is expected to earn alpha. In fact, most factor strategies only aim to harvest alpha from a small number of return factors. Now we define the first j factors in a factor model as pure risk factors that do not have expected alpha and the next k factors as return factors that have expected alpha, such that the first j elements of the α vector are all zeros, and the rest k elements are positive.

Then let's consider two risk models for Σ. Under Model 1, $\Sigma_1 = X^T F_1 X + D_1$, where F_1 is the covariance matrix of factors and D_1 is the diagonal matrix of stock specific return variance. Model 1 only uses the j risk factors to estimate factor covariance though. Therefore, the $j + k$ by $j + k$ area of F_1 only has non-zero values for the top left j by j area and the rest of its elements are 0. Under Model 2 which uses both risk factors and return factors to estimate factor covariance, $\Sigma_2 = X^T F_2 X + D_2$ and all of the $j + k$ by $j + k$ area of F_2 is filled with factor variance and covariance terms. It's worth noting that due to the different factor specification under Model 1 and 2, D_1 and D_2 are not exactly the same, but similar when the number of factors and assets are large.

Optimal weights under both models are: $w^* = \frac{1}{\lambda}(X^T F X + D)^{-1} X^T\alpha$.

However, since the factor covariance matrix F_1 under Model 1 is filled with zeros beyond the first j rows and columns, and the α vector's first j elements are zeros, the optimal weights under Model 1 can be simplified to:

$$w_1^* = \frac{1}{\lambda}{D_1}^{-1}X^T\alpha$$

Here we can see that under Model 1, optimal weights are exposures to alpha scaled down only by stock specific risk D_1. In other words, factor risk is not being penalized.

The optimal weights under Model 2 are:

$$w_2^* = \frac{1}{\lambda}(X^T F_2 X + D_2)^{-1} X^T\alpha$$

Here, a stock's exposures to alpha are scaled down by both factor risk, as represented by $X^T F_2 X$, and stock specific risk D_2.

So far we've seen that while in both cases the portfolios are constructed by selecting stocks with high ex-ante risk-adjusted returns, Model 1 adjusts expected alpha only by stock specific risk and could lead to excessive amount of factor risk in the portfolio. We show this point with empirical examples in the following sections.

Ex-ante risk case study

For a portfolio that is constructed using a generic risk model that doesn't include return factors, we observe that it has high factor risk contribution when measured under a customized risk model that does include return factors. Rerunning the optimization using the customized risk model leads to tighter control of factor risk.

The below table shows ex-ante risk predictions from the generic Axioma fundamental factor risk model (Generic Model) and our customized fundamental risk model (Customized Model) on a Value portfolio that has been optimized by the generic model versus Russell 3000 index. The generic model predicts active risk to be 1.25%, which is the standard deviation of the predicted active return of the portfolio. The pie chart shows that in variance terms, where 45% of that risk comes from factors and 55% comes from stock specific risk. A stock's specific risk is assumed to be uncorrelated with that of other stocks, and also uncorrelated with factor risks. However, given that the generic model doesn't include our return factors, it is possible that some of the stock specific returns are positively correlated to those return factors but the resulting covariance among stock specific returns and covariance between stock specific returns and factor returns have not been accounted for under the generic model. Putting the same portfolio through the lenses of our customized risk model results in higher active risk prediction, which is now 75% driven by factor risk in variance terms.

US All Cap Value 2019-12-31	Portfolio optimized using generic risk model	
Risk Measured by	Generic Model	Customized Model
Active Risk	1.25%	1.77%
Active Factor Risk	0.84%	1.53%
Active Specific Risk	0.92%	0.89%
Active Factor Risk (% of Variance)	45%	75%
Active Specific Risk (% of Variance)	54%	25%

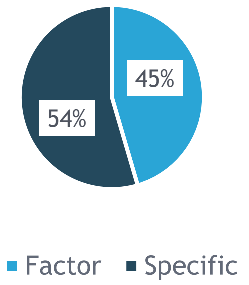

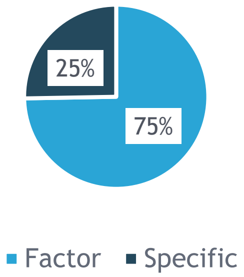

Source: SEI, using data from MSCI, Axioma and Factset. Strategy Name: US All Cap Value. Universe: US All Cap Equities. Risk model comparison as of 31 Dec 2019.

Let's not forget that the portfolio analyzed above, which we name as Portfolio 1, has been optimized using the generic model. As illustrated in the previous section, the optimization rewards a stock's exposure to return factors and penalizes its contribution to risk. We suspect that in Portfolio 1, some exposures to return factors escaped from penalization because these factors were not included in the generic model. Now we re-run the optimization using the customized risk model. The resulting Portfolio 2 has lower active risk, with less risk coming from factors. This highlights the problem stemming from misalignment in factor specification, where the generic model underestimates risk and optimization using it as input produces riskier portfolio.

US All Cap Value 2019-12-31	Portfolio 1	Portfolio 2
Optimized by	Generic Model	Customized Model
Active Risk	1.77%	1.39%
Active Factor Risk	1.53%	1.14
Active Specific Risk	0.89%	0.80%
Active Factor Risk (% of Variance)	75%	67%
Active Specific Risk (% of Variance)	25%	33%

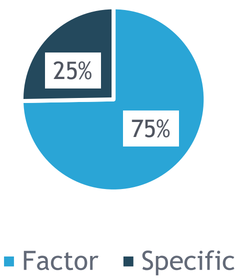

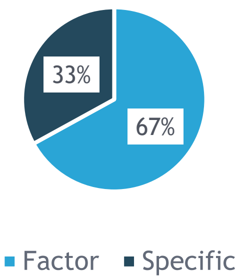

Source: SEI, using data from MSCI, Axioma and Factset. Strategy Name: US All Cap Value. Universe: US All Cap Equities. Risk model comparison as of 31 Dec 2019.

Ex-post risk analysis

While the previous section takes a snapshot of portfolio risk, the below simulation of how a Value strategy would have been constructed that goes as far back as 1996 shows that using customized risk model in optimization creates portfolios with more desirable ex-post characteristics. The first simulation of Value strategy optimized using the generic risk model earns an annualized active return of 2.3% with 4.2% active return volatility, or ex-post active risk over the entire simulation period. The second simulation using the customized risk model containing the return factors managed to earn slightly higher annualized active return of 2.4 but is much less volatile with annualized active risk only at 3.7%. More importantly, we see almost a 1/5 lift in the risk-adjusted return which goes from 0.55 to 0.64, as measured by the Information Ratio, when we optimize using the customized model instead of the generic model.

US All Cap Value 1996 - 2023	Portfolio 1	Portfolio 2
Optimized by	Generic Model	Customized Model
Annualized Active Return	2.3%	2.4%
Annualized Active Risk	4.2%	3.7%
Information Ratio	0.55	0.64

Source: SEI, using data from MSCI, Axioma and Factset. Universe: US All Cap Equities 1996 - 2023. Factor portfolios are rebalanced monthly. Past performance is no guarantee of future results. Updated to 31 Jan 2023.

From the above table we have seen that optimization using the customized risk model generates lower ex-post active risk over the long term than that using the generic risk model. The gap in active risk became larger since 2019. A closer look into the active factor exposure to one of the main target factor, Stable Value, explains the reason. While the two portfolios had similar ex-ante active exposure to Stable Value for most of the simulation period, Portfolio 2 became less exposed to Stable Value from 2019 to 2021. This happened because the customized risk model used by Portfolio 2 observed that the Stable Value factor had gone through a volatile period, updated its risk forecast and hence enabled the optimizer to reduce allocation to stocks with exposure to this factor. The generic model on the other hand, doesn't include this particular Stable Value factor and was unable to make the optimizer reduce its tilt towards the factor. As the Stable Value factor continued to underperform over the next 2 years, the portfolio optimized by the customized model absorbed less loss from the factor thanks to its lower exposure than the one optimized by the generic model.

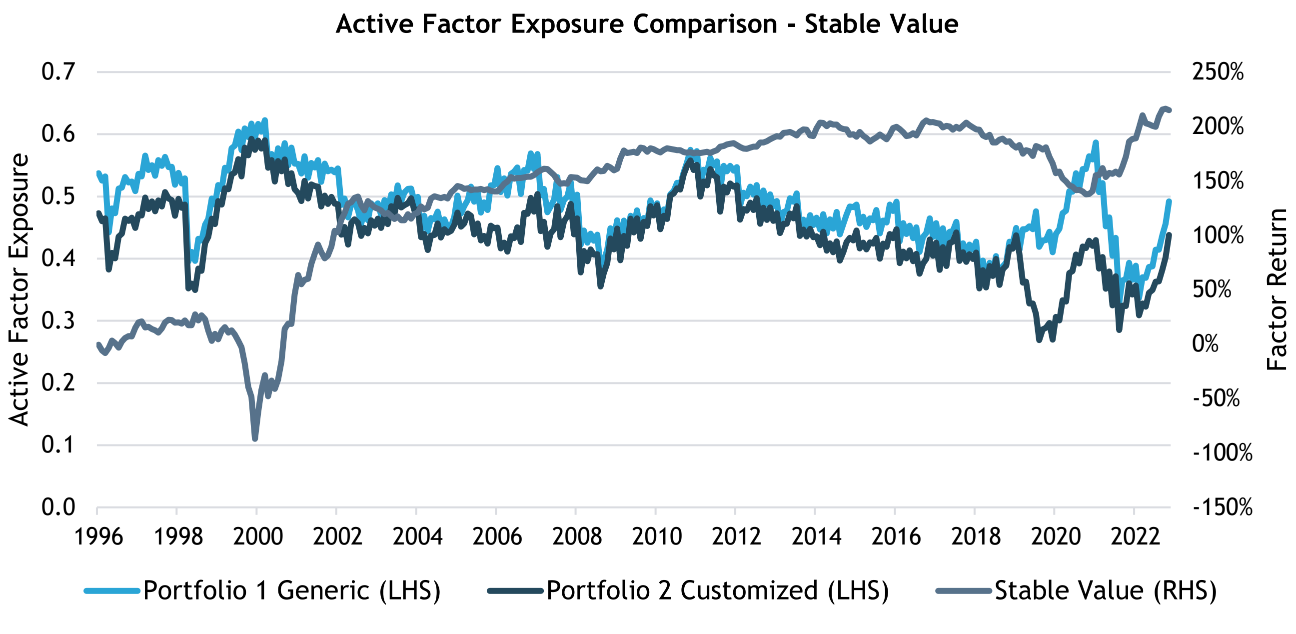

Source: SEI, using data from MSCI, Axioma and Factset. Universe: US All Cap Equities 1996 - 2023. Factor portfolios are rebalanced monthly. Factor returns are calculated as the total return difference between stocks in the top quintile of factor scores and stocks in the bottom quintile. Tax and transaction cost are not included in return calculation. Past performance is no guarantee of future results. Updated to 31 Jan 2023.

Model-risk diversification

We also need to be pragmatic to acknowledge that however careful we or third parties build a fundamental risk model, it may still from time to time fail to capture latent factors that may be transient or not directly connected to current or historical fundamental characteristics of stocks. A solution to this problem is to use a Principal Component Analysis (PCA) model to complement the customized fundamental risk model.

Elton and Gruber (1994) provides a comprehensive discussion of the techniques and characteristics of a PCA model. Unlike the fundamental model which has a pre-specified set of factors, the factors, factor returns, and factor exposures are solved for simultaneously under a PCA model and re-estimated independently for each model update. This feature makes PCA model more adaptive to market conditions.

As well as being more capable of capturing the impact from short-term market trend, PCA models also have the advantage of reducing estimation error. Take the simplest form of PCA model, the CAPM as an example. Each asset's market beta is estimated in the time-series independently. While some beta estimates may be overstated and some may be understated, estimation errors can be diversified away once beta estimates are grouped together to calculate portfolio beta, or risk.

Here we have another simulation example for a Momentum strategy in US All Cap universe between 1996 and 2003. Alongside the customized fundamental risk model, we ran another simulation using a PCA model. Other than observing that the portfolio optimized by the PCA model generated higher risk-adjusted return (Information Ratio), we can see from the ex-ante active risk comparison chart that it has had higher ex-ante active risk than the one optimized by the fundamental model. The gap in ex-ante active risk is bigger at times when the market experienced extreme volatility, such as when the dot com bubble burst in early 2000's, during the Global Financial Crisis in 2008, and since COVID-19 broke out in 2020.

A decomposition of ex-ante active risk suggests that the gap is driven by the higher factor risk contribution in Portfolio 2. In periods of market regime shifts, the PCA model is better at identifying the true driving forces of the market. The optimizer using PCA model penalizes exposures to such forces rather than the return factors. The fundamental model on the other hand thinks that return factors have become more volatile, which the consequence rather than the causes of market volatility. The optimizer using fundamental model then increases penalties to exposures to return factors and results in portfolios with less ex-ante active risk. In out-of-sample test, we compare ex-ante active risk and 2-year rolling realized active return volatility (ex-post active risk) for the two portfolios. The ratio of ex-post active risk over ex-ante risk is much closer to 1 from Portfolio 2 which has been optimized by the PCA model.

US All Cap Momentum 1996 - 2023	Portfolio 1	Portfolio 2
Optimized by	Customized Fundamental	Customized PCA
Annualized Active Return	3.0%	3.6%
Annualized Active Risk	3.0%	3.3%
Information Ratio	1.02	1.07

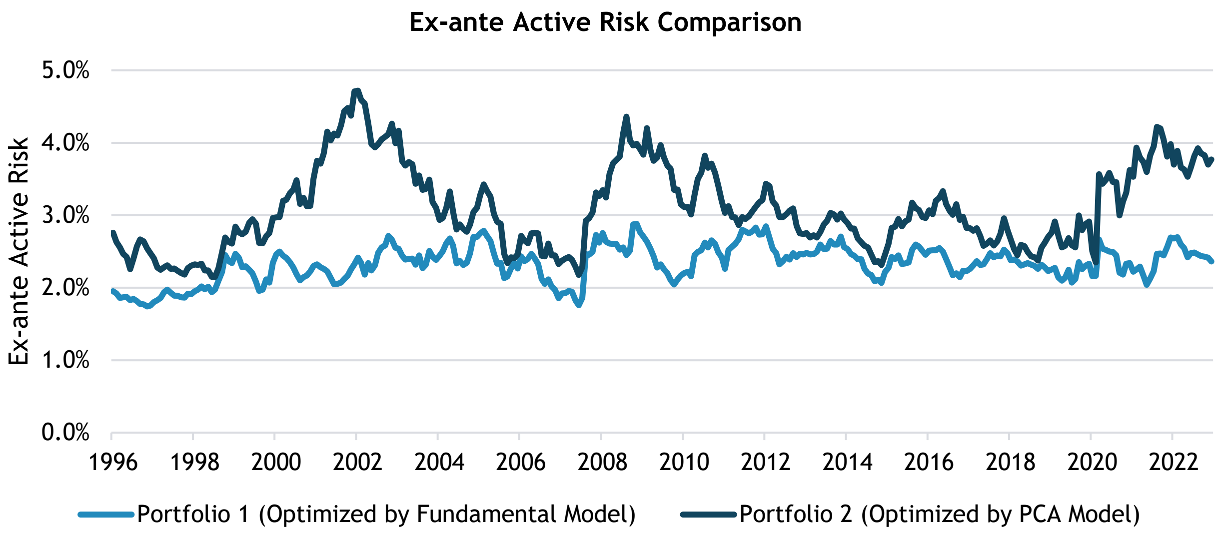

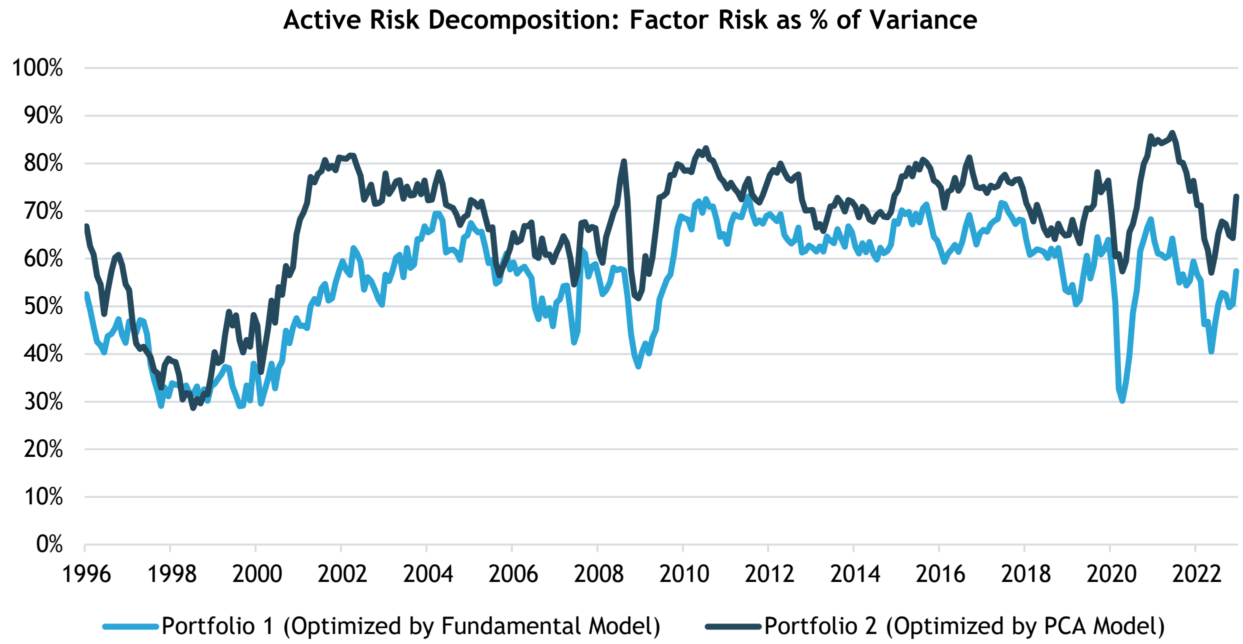

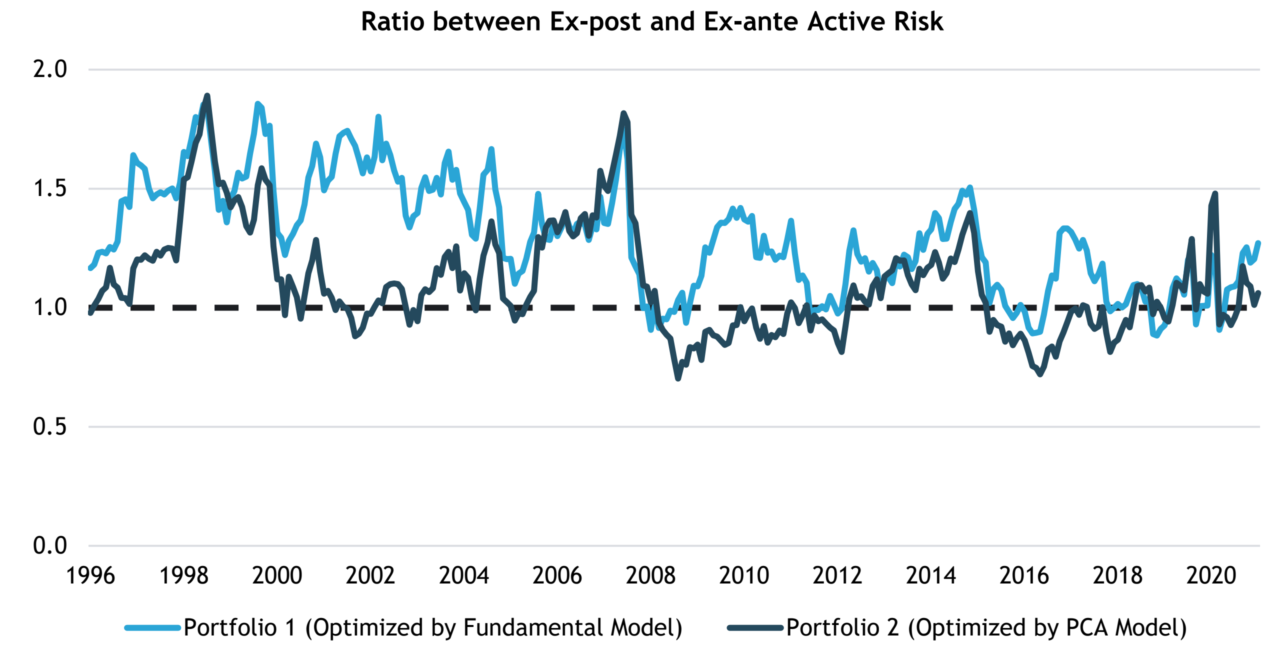

Source: SEI, using data from MSCI, Axioma and Factset. Universe: US All Cap Equities 1996 - 2023. Factor portfolios are rebalanced monthly. Tax and transaction cost are not included in return calculation. Past performance is no guarantee of future results. Updated to 31 Jan 2023.

Let's also not forget that the fundamental risk model already has an unfair advantage over the PCA model in back-testing because factors of fundamental risk models are pre-specified. A fundamental risk model may do a good job in back-testing using specification of risk factors that are known to us today but didn't exist in the past. Take the Internet industry factor as an example. The industry did not exist in 1999, but it exists as a risk factor today. Back-cast to 1999, it improves the performance of the risk model on the basis of hindsight knowledge that internet/dot com was a large risk factor.

Misalignment in horizon

While return factors in alpha generation may experience periods of underperformance, they have been validated over long history across different cycles and chosen for their timeless and universal properties. Risk models, on the other hand, usually estimate volatility and correlations using relatively shorter historical returns data with higher importance assigned to more recent observations. After a factor has experienced periods of elevated volatility, an optimizer using generic risk models penalizes it more harshly than after periods of subdued volatility.

We believe that the choice of look-back period and half-life for the estimation of risk in a customized fundamental risk model should be consistent with the horizon of alpha generation. Furthermore, the look-back period needs to be representative enough for the forward investment horizon. Shall we really extrapolate the benign 2003 - 2006 period for Value into 2007 - 2010? Likewise, is 2019 - 2023 going to represent 2023 - 2026? Risk modeling is only as good as the sample. Having a more recent sample does not mean a better sample.

Take the Low Volatility strategy as an example, where the optimizer is designed not to tilt towards any return factors but only penalizes risk. In this case, factor specification is less of an issue but optimized portfolio has far less turnover and stronger performance when risk is estimated with data over long term history than if it focuses on short term relationship which is likely to be noisier. The below table shows simulated quarterly optimization results coming from 3 different risk models. The first two models are Axioma generic risk models that have a 2-year look-back for volatility and 4-year look-back for correlation estimation. The short term model has a 2-month half-life for volatility estimation and 6-month for correlation estimation. The medium term model's half-life is longer at 6-month and 1-year for volatility and correlation respectively. Our customized long term model which is specifically used on Low Volatility strategies has much longer look-back as it utilizes data of the past 4 years for volatility and 10 years for correlation. Its half-life is 1-year for volatility and 3-year for correlation.

Without imposing any turnover constraint, it is clear that the generic risk model focusing on short term and medium term data results in much higher turnover than that of the simulation using customized model, which benefits from the much more stable risk estimates from the customized model. It's worth noting that replacing 45% or 31% of a factor portfolio every quarter as required by the generic risk models is simply not practical to implement. As for performance, while the models resulted in similar amount of active risk for the optimized portfolios, the customized model that is more consistent with the investment thesis has much stronger active return and risk-adjusted performance as measured by both Information Ratio and Sharpe Ratio. Assuming average transaction cost (implementation shortfall + commission) in US equity markets to be 30 bps per trade, then the portfolios optimized by the generic models wouldn't be profitable at all, while the one optimized by the long-term customized model still beats the benchmark by 0.83% net of transaction cost.

US Large Cap Low Volatility 1996 - 2023	Portfolio 1	Portfolio 2	Portfolio 3
Optimized by	Short-Term Generic	Mid-Term Generic	Long-term Customized
Annualized Active Return	-0.5%	0.4%	1.2%
Annualized Active Return (net)	-1.51%	-0.37%	0.83%
Annualized Active Risk	10.0%	10.4%	10.6%
Turnover (one-way, quarterly)	45%	31%	15%
Information Ratio	-0.05	0.04	0.11
Sharpe Ratio	0.55	0.62	0.66

Source: SEI, using data from MSCI, Axioma and Factset. Universe: US All Cap Equities 1986 - 2023. Factor portfolios are rebalanced quarterly. Returns shown in USD, gross of transaction costs and tax. Past performance is no guarantee of future results. Updated to 31 Mar 2023.

Other than carefully choosing the horizon of history data to use in risk estimation, we also need to consider direct measures that can be applied to improve consistency of optimization over time. Take an unconstrained mean-variance optimization problem as example, the optimal portfolio weight is increasing in alpha exposure but decreasing in risk. If risk estimates are volatile over time, optimal portfolio weights will be volatile even if stocks' alpha exposures are stable. On the other hand, we still need the risk estimates to be dynamic as they convey useful information of the market. To level the playground when trading off between stable factor exposures and dynamic risk estimates, we can also adjust factor exposure by market volatility to make the trade-off consistent across different market volatility regimes.

Misalignment between investment process and performance attribution

A generic multi-factor model estimates factor returns from many periods of cross-sectional regression of stocks' returns on stocks' factor exposures. Each period, such a regression also tells us how much of a stock's return can be attributed to each factor and how much remains unexplained, i.e. as stock specific return. Similar to the problem that arises from misalignment in factor specification between alpha generation and risk modelling, not including the same return factors in performance attribution leads to the result being misleading. Contribution from exposures to return factors identified for alpha generation may be wrongly attributed to risk factors or left in the unreasonably large stock specific returns.

Even if factor specification is aligned, we still have observed consistently big specific returns in portfolio return attribution reports, which confirms the findings of Sanne de Boer (2012). This is mostly due to the misalignment between how stock returns are weighted in alpha generation and ex-post factor return estimation regression.

Here we have two issues. Firstly, we propose that a new weighting scheme called liquidity weighting should be used in alpha generation, as it more accurately represents the investment opportunity set faced by a long-only portfolio manager. Then we demonstrate why empirically it is the better weighting scheme to use in cross-sectional regressions that estimate factor returns.

Introducing the Liquidity Weighting Scheme

We know that a scheme that weights stocks by the square root of market capitalization is better than equal-weighting that overstates the importance of small cap stocks and also better than purely market-cap-weighting that will be dominated by the mega-cap stocks. However, it still doesn't truly reflect the investment opportunity set in portfolio construction. A long-only equity mandate usually faces regulatory constraints on liquidity and diversification, as well as self-imposed constraints on the size of trades and holdings. For most active managers, there's little liquidity concern when buying large cap stocks. The required trading amount is usually fairly small comparing to the average daily volume (ADV) of the stock. But in the small cap space, the liquidity constraint kicks in and the max amount that one can buy becomes linearly decreasing in the stock's ADV or market cap.

Therefore, we find that a more realistic Liquidity Weighting Scheme should be a step function of a stock's market capitalization. Any stock with a benchmark weight larger than a threshold should be considered "large cap" and equally weighted. Any stock below this threshold should be weighted by market cap. The below charts illustrates the differences of these weighting schemes. The first one shows the whole spectrum of market caps of the Russell 1000 index, while the second and third charts zoom into the top 200 largest stocks and the rest of the index respectively. We can clearly see that weighting by market cap gives large stocks higher weights and small stocks smaller weights than equal-weighting. The widely used square-root of market-cap weighting offers a compromise in between. The liquidity weighting scheme assigns equal weights to the largest 100 stocks and scales the rest of the index by their market cap.

Comparing to square-root of market-cap weighting, liquidity weighting reduces the weights of the extremely large and small stocks. This is beneficial from a portfolio construction perspective, otherwise our return factors would either become effectively Mega Cap Tech factors under current market condition or distorted towards the very small cap universe that is not really investable once the assets under management reaches a certain level.

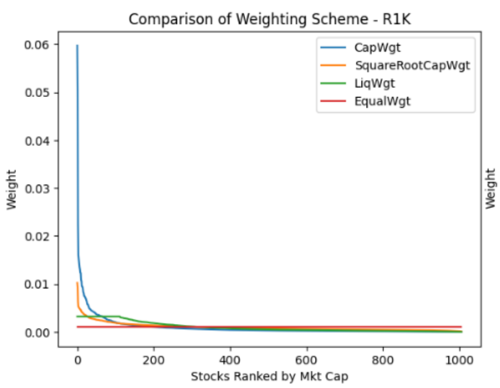

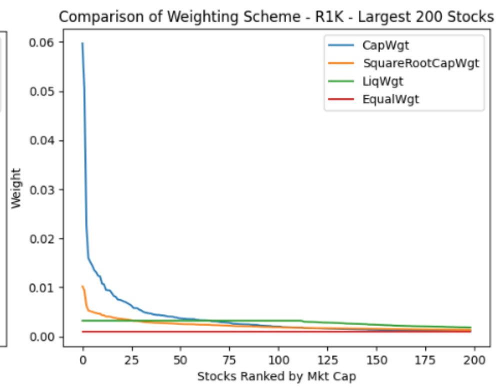

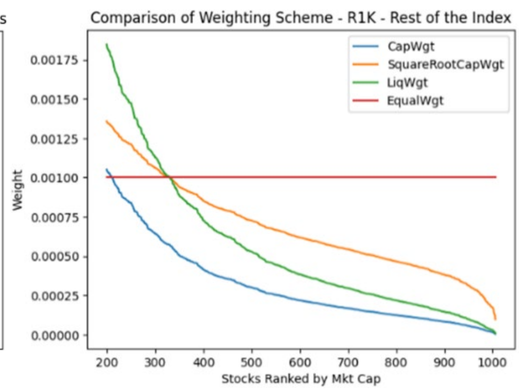

Source: SEI, using data from MSCI, Axioma and Factset. Universe: US Large Cap Equities as of 30 Nov 2022.

Liquidity weighting in Weighted Least Squares regression

When building a fundamental factor model, factor returns need to be estimated each period in a cross-sectional regression of stocks' excess returns on their factor scores. Variances of the residuals of an Ordinary Least Squares (OLS) regression are assumed to be the same across all stocks but they are not in reality, leading to the problem of heteroskedasticity which causes the OLS estimator to be inefficient. Because of this, confidence intervals and hypotheses tests cannot be relied on.

One common solution in theory is to use Weighted Least Squares (WLS) regression, where each observation is weighted by the inverse of its residual variance, i.e., an observation with a larger residual variance has a smaller weight and vice versa. Since residual variances are not known, as mentioned by Grinold and Kahn (1994), practitioners in building factor risk models have been using the square-root of market-cap weights as a proxy for the inverse of the residual variance.

However, liquidity weighting scheme provides an even closer proxy. In the below example, we first estimate residual returns from equal-weighted cross-sectional regression of stocks' excess returns on their factor scores monthly. Then we estimate each stock's residual variance and take the inverse of it as its theoretical weight to be used in a WLS regression.

The chart below shows the average of inverse of residual variance weights from the largest stock to the smallest stock in the Russell 1000 index over time between 1998 and 2022. It's worth noting that while the inverse of residual variance does decrease with market cap of the stock, there isn't much of a difference among the largest 100 stocks. This pattern is in contrast with those displayed in market cap weighting or square-root of market cap weighting schemes but agrees with liquidity weighting scheme which equal weights the mega cap stocks. Not surprisingly, liquidity weights have stronger correlation with inverse of residual variance weights than market cap and square-root of market cap weights do.

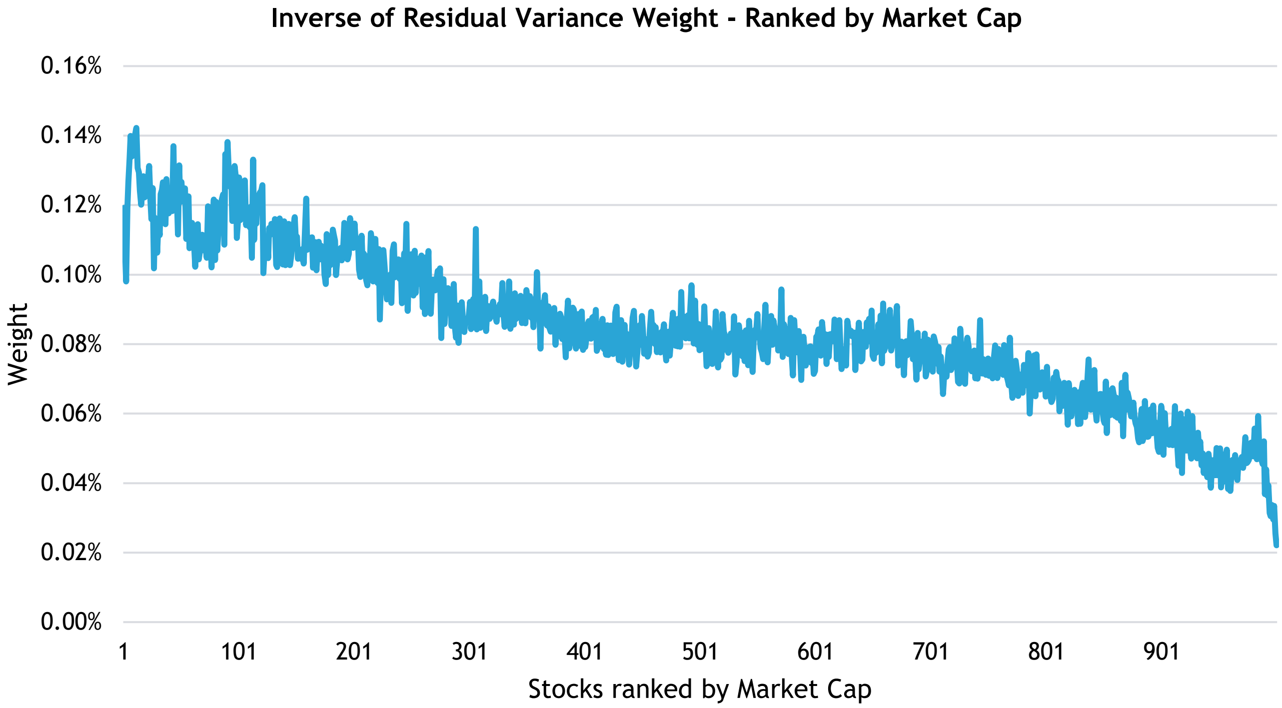

US Large Cap 1998 - 2022	Market Cap Weighted	Square-root of Market Cap Weighted	Liquidity Weighted
Correlation with the Inverse of Residual Variance Ranked by Market Cap	0.53	0.78	0.88

Source: SEI, using data from MSCI, Axioma and Factset. Universe: US Large Cap Equities 1998 - 2022.

Portfolio performance attribution comparison

Here we conduct performance attribution on factor top decile portfolios using different weighting schemes in the regressions that estimate factor returns. The total returns of these factor top decile portfolios are meant to be most driven by factors once we strip out market returns. In most cases, liquidity weighting scheme results in the smallest residual return when measured as a percentage of total return. The chart zooming into the top decile of the Profitability factor shows that residual return from liquidity weighting scheme stayed closer to zero over long period of time than that from the square-root of market-cap weighting scheme.

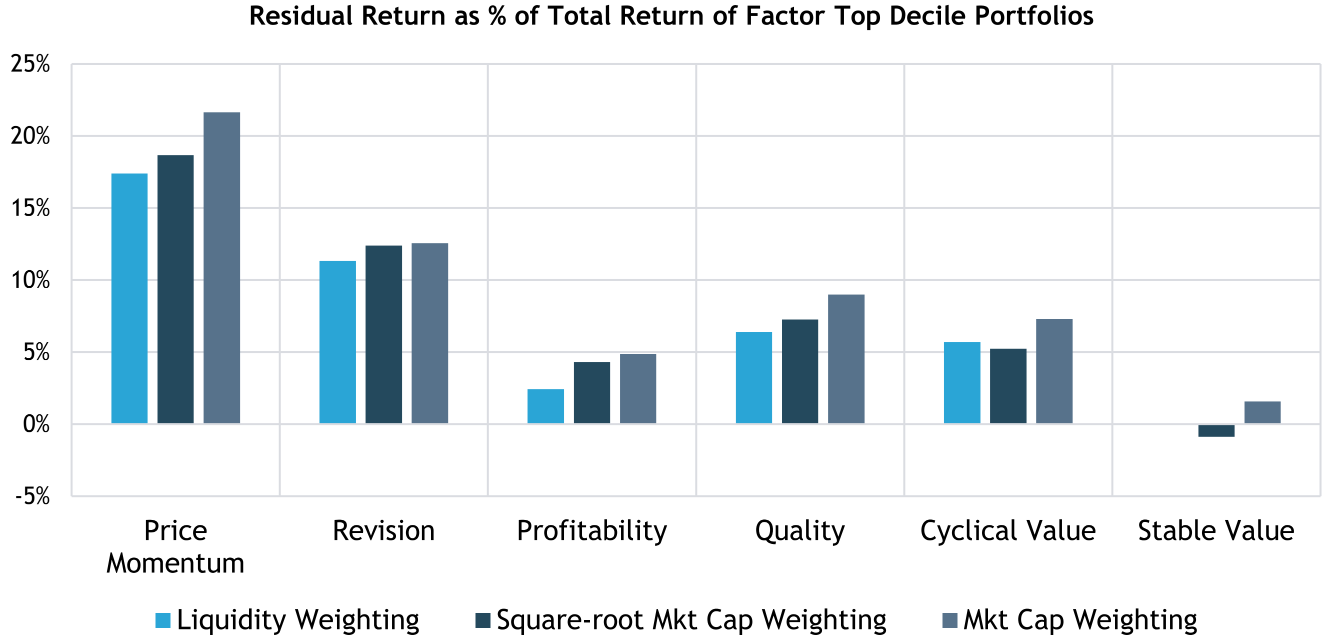

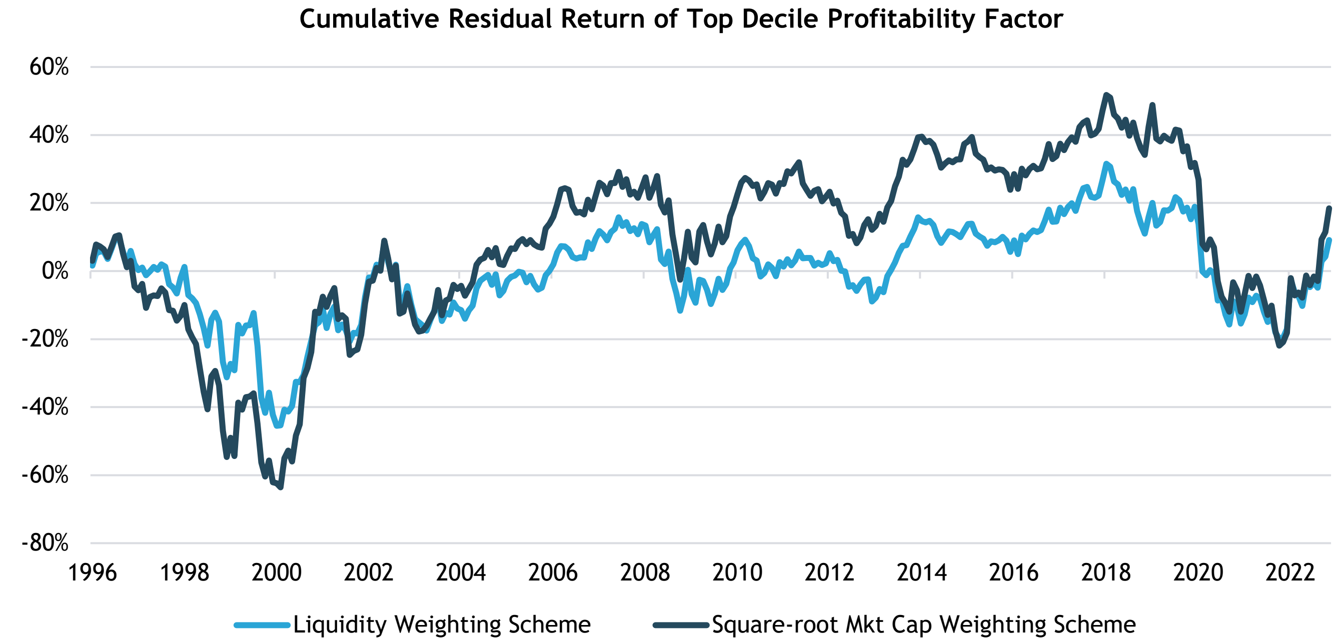

Source: SEI, using data from MSCI, Axioma and Factset. Universe: US Large Cap Equities 1996 - 2022. Factor portfolios are constructed using the top decile of the index and rebalanced monthly. Past performance is no guarantee of future results. Updated to 31 Dec 2022.

Conclusion

While factor investing starts with factor research, we highlight the importance of aligning factor models used in different stages of the investment process to its success. We propose several customized measures to be implemented in portfolio construction and performance attribution processes.

Having demonstrated the factor misalignment problem in optimization, where the optimizer makes unintended bet on return factors that are not included risk model, we propose using a customized fundamental risk model which incorporates both return factors used in alpha generation and other risk factors from generic third-party risk model to address the problem. We then discuss the benefits of adding risk input from a statistical model into the optimization, before showing how aligning the risk-estimation horizon with that of the investment strategy adds value. Individually, each of these customized measures helps improve risk-adjusted performance of the factor strategies under consideration by 5-20% in our simulation. Retaining even a fraction of the combined benefits would be a sizeable overall portfolio efficiency improvement.

Finally, weighting scheme used in the regressions by which we attribute returns to factors also plays an important role in both alpha generation and performance attribution. We argue that a more realistic liquidity weighting scheme should be used to more truthfully represent the investment opportunity set faced by a long-only active mandate.

References

1. Sanne de Boer. "Factor attribution that adds up" Journal of Asset Management volume 13, pages 373–383 (2012)
2. Ceria, S., Saxeena, A., and Stubbs, R. “Factor Alignment Problems and Quantitative Portfolio Management” Journal of Portfolio Management, Winter 2012, 38 (2), 29–43
3. Edwin J. Elton and Martin J. Gruber, "Multi-Index Models Using Simultaneous Estimation of All Parameters", "A Practitioner's Guide to Factor Models", The Research Foundation of the Institute of Chartered Financial Analysts, 1994
4. Richard Grinold and Ronald N. Kahn. "Multiple-Factor Models for Portfolio Risk", "A Practitioner's Guide to Factor Models", The Research Foundation of the Institute of Chartered Financial Analysts, 1994
5. Jyh-Huei Lee and Dan Stefek. “Do Risk Factors Eat Alphas?” The Journal of Portfolio Management, Summer 2008, 34 (4) 12 - 25
6. Anureet Saxena and Robert A. Stubbs. "Augmented Risk Models to Mitigate Factor Alignment Problems" Journal Of Investment Management, Vol. 13, No. 3, (2015), pp. 57–79