JOIM
www.joim.com

JOURNAL OF INVESTMENT MANAGEMENT, Vol. 13, No. 4, (2015), pp. 71–86

EFFICIENTLY COMBINING MULTIPLE SOURCES OF ALPHA

Jose Menchero and Jyh-Huei Lee†*

In this article, we examine the question of efficiently combining multiple sources of alpha. We begin with a comparison of the various methods used by practitioners for constructing portfolios that capture a single alpha signal. These methods are broadly categorized as either: (a) simple factor portfolios, (b) pure factor portfolios, or (c) minimum-volatility factor portfolios. We then derive an equation that shows the optimal alpha weights given the expected returns and covariance matrix of the alpha signals. We provide a discussion on how the required inputs can be estimated in practice, and conclude with an empirical example to illustrate these effects.

1 Introduction

Alpha represents the opportunity for investors to "beat the market." For active managers, the search for alpha is usually regarded as the most important component of the investment process. Nonetheless, poor implementation of even the best-quality alpha signal may result in mediocre risk-adjusted performance.

A crucial question, therefore, concerns how to construct a portfolio that captures a particular alpha signal. Various portfolio construction techniques are used in practice to accomplish this task. One focus of our paper is to compare and contrast these various techniques.

A common approach for capturing an alpha signal is to create a *simple* factor portfolio. This is constructed by directly overweighting stocks with positive alpha, and underweighting those with negative alpha. While such a portfolio certainly captures the alpha factor,[1] it will also have incidental exposure to many other risk factors. These exposures typically add unnecessary risk to the portfolio, without improving the expected performance.

A more risk-aware approach to capturing an alpha signal is to tilt the portfolio toward the alpha

*Jose Menchero is Founder and CEO of Menchero Portfolio Analytics Consulting.
†Jyh-Huei Lee is Director of Quantitative Equity Portfolio Management at TIAA-CREF.

factor, while hedging away exposures to all other risk factors. We refer to this as the *pure* factor portfolio. While this portfolio has the reassuring quality that it has exposure only to the alpha factor, it does not represent the theoretical optimal implementation of the signal.

The most efficient implementation of the alpha signal is given by the *minimum-volatility* factor portfolio.[2] As discussed by Grinold and Kahn (2000), the minimum-volatility factor portfolio has the lowest volatility of all portfolios with unit exposure to the factor. Applied to an alpha factor, therefore, the minimum-volatility factor portfolio has the highest expected information ratio.

An important characteristic of minimum-volatility factor portfolios is that they have non-zero exposures to all risk factors. Many portfolio managers regard these exposures as non-intuitive, and may therefore set tight constraints on these factor exposures as part of the optimization process. We argue against blindly imposing such constraints. One of the aims of our paper is to build the intuition behind these factor exposures, and to demonstrate the practical benefit in volatility reduction that can be achieved by allowing these risk factors to *hedge* the risk of the alpha factor.

In practice, most portfolio managers employ more than a single source of alpha. Another critical question, therefore, is how to optimally combine multiple alpha signals. Qian *et al.* (2007) provided one solution to this problem in terms of information coefficients of the signals and the covariance matrix of these information coefficients. In Qian's approach, the information coefficient is defined as the cross-sectional correlation between the forecasts and the subsequent returns. Our approach is quite different in that we solve for the optimal combination of weights in terms of the expected returns and covariance matrix of minimum-volatility factor portfolios.

Another salient feature of Qian's approach is that the optimal portfolio is constrained to have zero exposure to all risk factors. This contrasts with our approach, where we explicitly allow for such exposures in order to hedge the risk of the alpha factor.

Another method for combining multiple sources of alpha was described by Grinold (2010). In this approach, one constructs efficient portfolios for each alpha signal and then scales them to unit volatility. The optimal weights are then solved for in terms of the correlation matrix and information ratios of these unit-volatility portfolios.

Our approach to signal weighting is similar to Grinold's. The focus of our paper, however, is very different. The main objective of Grinold's article was to provide a framework for incorporating the effects of transaction costs for signals with different levels of information turnover. By contrast, the focus of our paper is how to estimate the required inputs for the optimization process itself. More specifically, we provide explicit detail showing how to estimate the risk and expected returns for each signal within a self-consistent framework.

The remainder of this paper is organized as follows. We begin with a detailed description of the various types of factor portfolios. Next, we provide definitions for "alpha" factors and "risk" factors. These definitions are essential for understanding how risk factors can be used to improve the risk-adjusted performance of a portfolio. More specifically, we show that risk factors can be used to hedge alpha factors, without impacting expected returns. We then present empirical results to compare and contrast the risk and return characteristics of these various factor portfolios. Finally, we present our approach for combining multiple alpha signals, together with an illustrative example. Mathematical details are provided in the technical Appendix.

2 Capturing a single alpha source

A common source of confusion and ambiguity in the realm of factor modeling comes from generic usage of the term "factor." To avoid such confusion, it is important to clearly distinguish between factor *exposures* and factor *returns*. In fundamental factor models, the exposures represent intuitive attributes of individual stocks. These might correspond to industry or country membership, or to some other stock characteristics such as beta, book-to-price ratio, or alpha.

When we speak of factor returns, by contrast, we refer to the returns of portfolios that are *derived from* the factor exposures via some defined algorithm. These factor portfolios are generally long/short, although not always dollar neutral. In this article, we examine three distinct types of factor portfolios. Each is characterized by the same unit exposure to the factor in question. What distinguish them are their exposures to the other factors.

2.1 *Simple factor portfolios*

The first type we consider is the *simple* factor portfolio, which may be regarded as a simple tilt on the factor. This portfolio is constructed by directly taking long positions in stocks with positive exposure to the factor, and short positions in stocks with negative exposure. Mathematically, this is equivalent to performing a univariate regression against the factor in question, with the resulting weights being given by Equation (A.5) of the technical Appendix.

If the factor has some degree of collinearity with other factors—as is typically the case—then the portfolio construction process guarantees that the resulting portfolio will take "inadvertent" tilts on those factors as well. For instance, if positive momentum stocks tend to have high beta, then the momentum simple factor portfolio will also have positive exposure to the beta factor. Similarly, if energy stocks have recently performed well, then the simple momentum factor portfolio will also be overweight in the energy sector. Such unintended bets generally add risk to the portfolio, without enhancing the expected return.

2.2 *Pure factor portfolios*

By contrast, *pure factor* portfolios have unit exposure to the factor in question, with zero exposure to all other factors. Mathematically, pure factor portfolios are formed by multivariate cross-sectional regression, with the resulting weights given by Equation (A.8) of the technical Appendix. Pure factor portfolios are powerful constructs because they provide a means of disentangling the confounding effects of collinearity.

Before proceeding to a concrete example, we first describe the set of factors employed in our study. We focus our analysis on the Barra Global Equity Model (GEM3), which contains the following four types of equity factors: (a) a world factor, representing the regression intercept, to which every stock has unit exposure, (b) multiple country factors, spanning developed and emerging markets, with exposures given by 0 or 1, (c) 34 industry factors, again with exposures given by 0 or 1, and (d) 11 style factors, which are standardized to be mean zero and unit standard deviation.

The estimation universe for the GEM3 Model is the constituents of the MSCI All Country World Investable Market Index (ACWI IMI), a broad market index spanning both emerging and developed markets. Factor returns are estimated for each period by cross-sectional regression of stock returns against the start-of-period factor exposures. Specific returns represent the portion of stock return that cannot be explained by the factors. The GEM3 Model assumes that the specific returns are mutually uncorrelated, thus implying that the model factors fully capture all sources of equity return covariance. A detailed description of the GEM3 Model is provided by Morozov *et al.* (2012).

Note that the GEM3 factor structure contains two exact collinearities. Namely, the sum of column vectors corresponding to the industry factor exposures replicates the exposure to the world factor, with the same identity holding for country factors. To obtain a unique regression solution, we must impose two constraints on the factor returns: Every period, we set the cap-weighted average industry and country factor returns to zero. To solve for the pure factor portfolio weights, we use square root of market capitalization as the regression weights. Further details are provided in the technical Appendix.

As discussed by Menchero (2010), the pure factor portfolio corresponding to the world factor essentially represents the cap-weighted world portfolio. Pure country factor portfolios are 100 percent long the particular country and 100 percent short the world portfolio, with zero net weight in every industry and zero exposure to every style. In other words, pure country factor portfolios capture the performance of the country net of the market, industries, and styles. Similarly, pure industry factor portfolios capture the performance of the industry net of the market, countries, and styles. Pure style factor portfolios have unit exposure to the particular style, and zero exposure to all other countries, industries, and styles.

2.3 Minimum-volatility factor portfolios

The third kind of factor portfolio that we consider is the *minimum-volatility* factor portfolio. This represents the unique portfolio that has the lowest predicted volatility among the set of all portfolios with unit exposure to the factor. Mathematically, this portfolio is formed by mean–variance optimization, with the weights being given by Equation (A.9) of the technical Appendix. Note that the minimum-volatility factor portfolio has non-zero exposure to all other factors. These exposures, however, are designed to reduce the overall risk of the portfolio by hedging the factor in question.

Turning now to a concrete example, we consider the minimum-volatility factor portfolio for the GEM3 momentum factor on December 31, 2012. The asset covariance matrix was obtained from the GEM3 risk model. The universe used to construct the portfolio is the constituents of MSCI ACWI IMI Index. In Table 1, we attribute portfolio risk to the GEM3 factors using the *x-sigma-rho* risk attribution formula, as described

Table 1 Risk attribution analysis on 31-Dec-2012 for the GEM3 momentum minimum-volatility factor portfolio. For illustrative purposes, only a representative subset of the GEM3 factors is exhibited.

Factor	Exposure	Volatility	Correlation	Risk contrib
World	0.04	11.25	0.00	0.00
Banks	0.06	3.66	0.00	0.00
Momentum	1.00	2.91	0.80	2.34
Beta	0.02	4.40	0.00	0.00
Size	0.04	1.11	0.00	0.00
USA	−0.02	3.64	0.00	0.00

by Menchero and Davis (2011). For illustrative purposes, we report only a small subset of the factors. In the *x-sigma-rho* scheme, the risk contribution is given by the product of the exposure (x) to the return source, the volatility (sigma) of the return source, and the correlation (rho) between the return source and the overall portfolio. For the example at hand, the return sources represent the GEM3 pure factor portfolios and the overall portfolio is the minimum-volatility momentum factor portfolio. As shown in Equation (A.10) of the technical Appendix, the minimum-volatility factor portfolio is uncorrelated with all of the pure factor portfolios, except its own. Consequently, a defining characteristic of minimum-volatility factor portfolios is that the pure factors contribute zero to their risk.

Note that the exposure to the momentum factor, by construction, is exactly 1. The other exposures, however, are non-zero. We argue that these exposures are intuitive and serve to reduce portfolio volatility. For instance, Table 1 shows that the minimum-volatility momentum factor portfolio has small positive exposures to both the world factor (0.04) and the beta factor (0.02). On the analysis date, both of these factors were negatively correlated with the pure momentum factor portfolio. As a result, positive exposure to these factors helps hedge the risk of the momentum factor.

The fact that the pure factors contribute zero to the risk of the minimum-volatility factor portfolio does not imply that they are ineffective at reducing portfolio risk. On the contrary, it is precisely because of these factor exposures that the correlation of the pure momentum factor is only 0.80 with the momentum minimum-volatility factor portfolio. Note that the correlation of the pure momentum factor portfolio with itself is exactly 1.0; therefore, the correlation of 0.80 leads to a significant reduction in portfolio volatility.

3 Alpha factors and risk factors

When building equity multi-factor risk models, the primary objective is to identify factors that capture and explain equity return co-movement. A "good" risk factor is therefore one that is very volatile and explains cross-sectional differences in stock returns. For risk model construction purposes, it is largely irrelevant whether the factor has any return premium associated with it. For instance, an industry grouping may represent a good risk factor if stocks within the industry tend to strongly co-move. This does not imply, however, that the industry factor portfolio will have a significant drift (i.e., earn "abnormal" returns) across time.

For portfolio construction purposes, by contrast, it is crucial to identify which factors represent alpha factors, and which do not. In other words, only a small subset of the risk factors may exhibit significant drift across time. *We define an "alpha" factor as one for which the pure factor portfolio has non-zero expected return.* A "good" alpha factor therefore is one that has a strong drift and high information ratio. Conversely, the other risk factors (i.e., those that do not qualify as alpha factors) are defined to have expected returns of zero.

It is worth mentioning that while not all risk factors are expected to be alpha factors, there is strong reason to believe that the converse is not true. That is, if markets are efficient, any source of alpha should have a systematic risk associated with it.

Determining which risk factors represent alpha factors and which do not requires a combination of empirical analysis and subjective insight. One of the criteria that we will employ is that the alpha factor must exhibit *persistence*. Therefore, we adopt the view that to qualify as a potential alpha factor, the realized information ratio should differ *significantly* from zero.

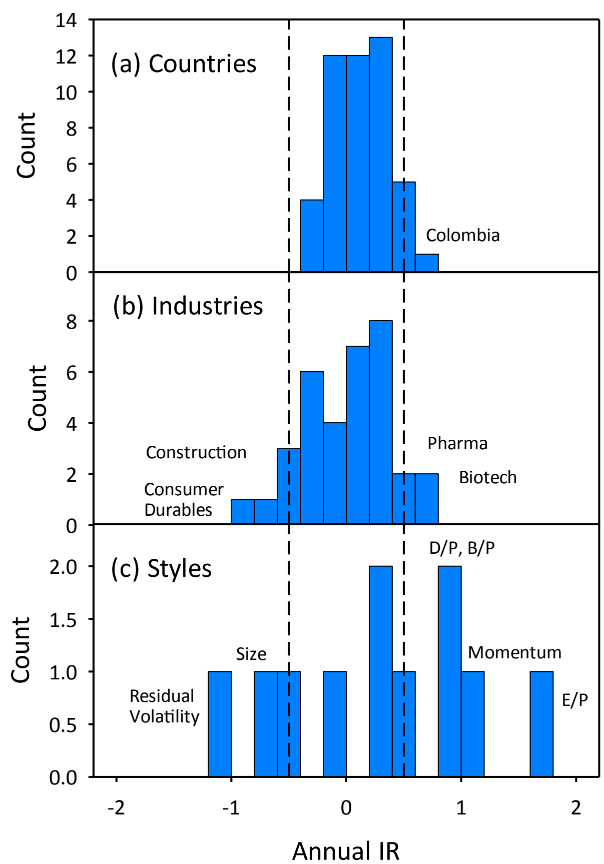

Figure 1 Histogram of information ratios for (a) country factors, (b) industry factors, and (c) style factors. The 95 percent confidence level is indicated by the dashed vertical lines. The sample period runs from January 1997 through December 2012.

In Figure 1, we present histograms of the realized information ratios for the GEM3 pure factor portfolios over a 16-year period spanning January 1997 to December 2012. The factor returns were estimated using monthly cross-sectional regressions with MSCI ACWI IMI as the estimation universe. Even if the *true* information ratio were zero, the *realized* information ratio would never be exactly zero due to sampling error. The confidence interval can be established by assuming normality and a true information ratio of zero; in this case the realized information ratio over the 16-year period should fall between −0.5 and 0.5 with the probability of 95 percent.

In Panel (a) of Figure 1, we plot the histogram of realized information ratios for the 48 GEM3 country factors representing both developed and emerging markets. Only one country (Colombia) fell outside of the confidence interval. This is not much different from what would be expected by chance. In this paper, therefore, we will assume that country factors do not represent alpha factors.

In Panel (b) of Figure 1, we show the realized information ratios for the 34 GEM3 industry factors. In this case, four factors fall outside of the confidence interval. This is only slightly more than the five percent that would be expected by pure chance. Pharmaceuticals and biotechnology were statistically significant on the positive side, whereas as consumer durables and construction were significant on the negative side. For some of these, however, we find that the outcome was dominated by brief spurts of performance. For instance, the best-performing industry factor was biotechnology, which was characterized by a short window of spectacular performance in 2001. In other words, it was not persistent drift that caused biotechnology to fall outside of the confidence interval, but rather a one-time event that is unlikely to repeat itself in the future. In this paper, therefore, we adopt the position that industry factors are not alpha factors.

In Panel (c) of Figure 1, we present the histogram of realized information ratios for the 11 GEM3 style factors. In this case, six of the factors lie outside of the confidence interval—much greater than the five percent that would be expected by chance. We further observe that several of these lie well outside of the confidence interval. For instance, earnings yield and momentum had realized information ratios between 1 and 2. Dividend yield and book-to-price also performed well, with information ratios between 0.5 and 1.0. On the negative side, size and residual volatility had significant information ratios.

These results are also consistent with the academic literature. For instance, Basu (1977) showed that high earnings-to-price stocks have historically outperformed on a risk-adjusted basis. Similarly, Banz (1981) documented the size effect, in which small-cap stocks tend to outperform their large-cap counterparts. The momentum effect, in which recent "winners" outperform "losers," was described by Jegadeesh and Titman (1993). More recently, Ang *et al.* (2009) demonstrated the low-volatility anomaly in which high volatility stocks have underperformed.

We might reasonably conclude that all six represent valid alpha factors. Nonetheless, for expositional simplicity, we reduce the list to just three. In this study, we consider only earnings yield, dividend yield, and book-to-price as *alpha* factors; we assume the other eight style factors have zero expected returns.

4 Empirical comparison

In this section, we compare the risk and return characteristics of the three types of factor portfolios. In Figure 2, we plot the cumulative performance of the simple, pure, and minimum-volatility factor portfolios for the earnings yield

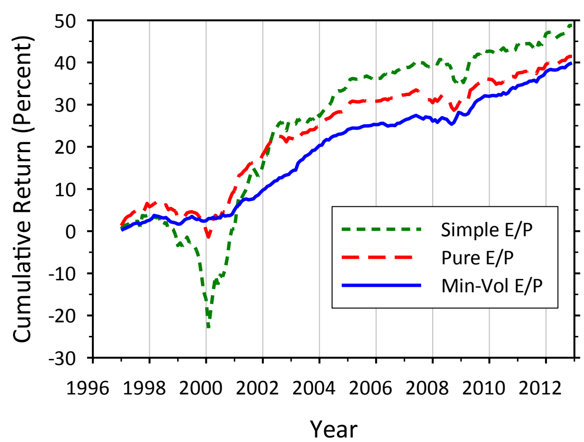

Figure 2 Cumulative performance of simple, pure, and minimum-volatility factor portfolios for the GEM3 earnings yield factor.

factor. All portfolios are constructed using MSCI ACWI IMI as the universe and are rebalanced on a monthly basis. Our first observation is that all three portfolios had similar cumulative performance over the 16-year sample period. The volatilities, by contrast, were dramatically different. Over the sample period, the annualized volatilities of the simple, pure, and minimum-volatility factor portfolios were 4.63 percent, 2.24 percent, and 1.28 percent, respectively. In other words, the minimum-volatility factor portfolio captured the return most efficiently, whereas the simple factor portfolio exhibited much higher risk.

To illustrate the differences between these factor portfolios, it is particularly instructive to examine the Internet Bubble period. We see that the simple factor portfolio exhibited a spectacular drop during 1999 and early 2000, followed by an equally impressive rebound beginning in April 2000. Qualitatively, the pure earnings yield factor portfolio followed a similar pattern, except that the returns were less extreme. Most remarkably, the minimum-volatility factor portfolio "sailed through" the Internet Bubble virtually unperturbed. It is worth stressing that the return contribution from the earnings yield factor itself was identical in all three cases, since each portfolio had the same unit exposure to the factor. The starkly differing behaviors of these portfolios were therefore due to differences in exposures to other factors.

To gain further insight into these results, we report in Table 2 the portfolio exposures to a representative subset of GEM3 factors at the start of April in 2000. We see that the simple factor portfolio had large negative exposures to both beta and momentum. These exposures are intuitive, given that stocks with high earnings yield had strongly underperformed the broad market (i.e., negative momentum) and tended to have

Table 2 Exposures of earnings yield factor portfolios to GEM3 factors, as of 31-Mar-2000. For illustrative purposes, only a subset of GEM3 factors are reported. The simple factor portfolio had positive exposures to "old-economy" industries and negative exposures to "new-economy" industries; for the minimum-volatility factor portfolios, the reverse holds.

Factor	Simple exposure	Pure exposure	Min-vol exposure
Beta	−0.347	0.000	0.089
Momentum	−0.554	0.000	0.014
Earnings Yield	1.000	1.000	1.000
Dividend Yield	0.565	0.000	0.327
Book-to-Price	0.579	0.000	0.183
Food Beverage	0.018	0.000	−0.011
Diversified Financials	0.019	0.000	−0.011
Construction	0.018	0.000	−0.026
Telecommunications	−0.054	0.000	0.024
Internet	−0.031	0.000	0.016
Software	−0.041	0.000	0.018

low beta. Also note that the simple earnings yield factor had large positive exposures to the "old-economy" industries (e.g., food & beverage, diversified financials, and construction), while it had large negative exposures to the "new-economy" industries (e.g., telecommunications, Internet, and software). This example illustrates that the risk of the earnings yield simple factor portfolio was not driven by the earnings yield factor itself, but rather by the "incidental" exposures to other factors.

The factor exposures of the pure earnings yield factor portfolio are very straightforward: By definition, the portfolio has unit exposure to itself and zero exposures to all other factors. As the pure factor portfolio has no incidental exposures, the performance is entirely attributed to the earnings yield factor itself.

The factor exposures of the minimum-volatility factor portfolio are the most interesting to interpret. In April 2000, the pure earnings yield factor returns were positively correlated with the old-economy industries, and negatively correlated with the new-economy industries. As a result, we find that the minimum-volatility factor portfolio had *positive* exposure to the new-economy industries and *negative* exposure to the old-economy industries in April 2000. These served as excellent hedges for the risk of the earnings yield factor, and explain the "smooth sailing" that the portfolio experienced during this otherwise stormy period.

In Figure 3, we plot the cumulative performance of the simple, pure, and minimum-volatility factor portfolios for the dividend yield factor. Again, we used the MSCI ACWI IMI as the universe and rebalanced monthly. Qualitatively, the results are very similar to the earnings yield factor in Figure 2. Namely, the annualized returns of the three portfolios were roughly comparable over the entire sample period, whereas the volatilities were drastically different. Note also the striking similarity between the behavior of the earnings yield and dividend yield simple factor portfolios during the Internet Bubble period. Again, this was mainly due to similar "incidental" exposures to other factors.

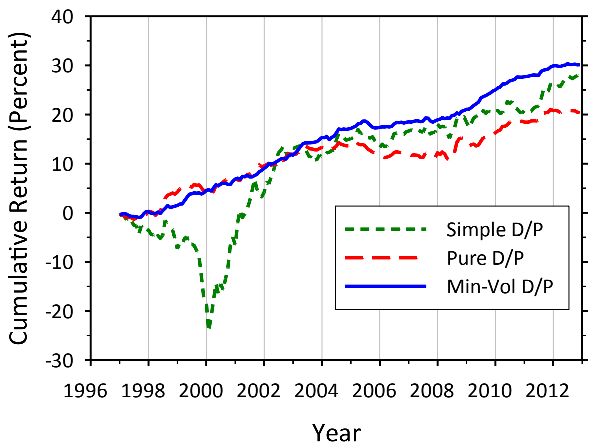

Figure 3 Cumulative performance of simple, pure, and minimum-volatility factor portfolios for the GEM3 dividend yield factor.

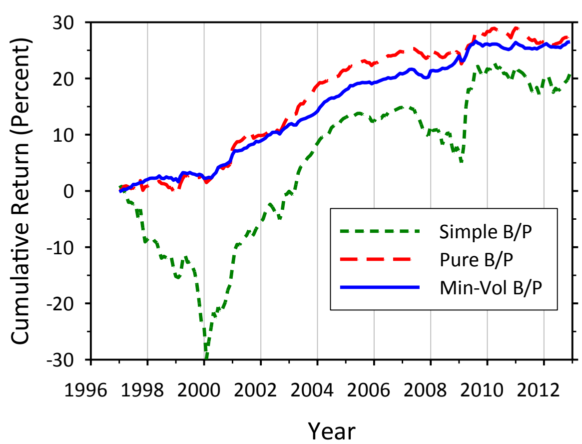

Figure 4 Cumulative performance of simple, pure, and minimum-volatility factor portfolios for the GEM3 book-to-price factor.

In Figure 4, we plot the cumulative performance of the book-to-price factor. Once more, we used the MSCI ACWI IMI universe and rebalanced the portfolios monthly. Qualitatively, the results are again very similar to those for the earnings yield and dividend yield factors. That is, the net returns over the sample period were quite similar, whereas the volatilities were starkly different.

In Table 3, we summarize results for the return, risk, and information ratio of the three types of factor portfolios for earnings yield, dividend yield, and book-to-price factors. As previously noted, the annualized returns of the simple, pure, and minimum-volatility factor portfolio were comparable. For earnings yield, the simple factor portfolio slightly outperformed the other two, whereas the minimum-volatility factor portfolio performed best for dividend yield and the pure factor portfolio had the highest returns for the book-to-price factor. Nonetheless, in each case, the return difference between the top-performing and bottom-performing factor portfolio was rather small (typically 50–60 bps).

While the annualized returns were quite similar, the volatilities were dramatically different. In each case, the volatility of the pure factor portfolio was *less than half* the volatility of the corresponding simple factor portfolio. Similarly, the volatilities of the minimum-volatility factor portfolios were roughly 40 percent lower

Table 3 Summary return, risk and information ratios for earnings yield, dividend yield, and book-to-price factors. For each style factor, we report results for simple, pure, and minimum-volatility factor portfolios. The sample period runs from January 1997 through December 2012.

	(Earnings yield)			(Dividend yield)			(Book-to-price)		
	Simple	Pure	Min-vol	Simple	Pure	Min-vol	Simple	Pure	Min-vol
Return	3.05	2.59	2.49	1.71	1.27	1.88	1.30	1.74	1.66
Risk	4.63	2.24	1.28	4.47	1.55	0.98	4.66	1.81	1.09
IR	0.67	1.25	2.11	0.39	0.67	1.74	0.31	0.96	1.50

than the volatilities of their pure factor portfolio counterparts.

Also reported in Table 3 are the information ratios[3] for the three different types of factor portfolios. Not surprisingly, we find that the minimum-volatility factor portfolios had the highest information ratios over the sample period. This was primarily due to their lower volatilities. Conversely, the lowest information ratios occurred for the simple factor portfolios, owing to their much higher volatilities.

5 Combining multiple sources of alpha

Next, we turn our attention to the case in which the alpha factor is composed of several distinct alpha signals. For this study, we take the alpha factor to be a weighted combination of the earnings yield, dividend yield, and book-to-price signals. Our task is to construct the maximum information ratio portfolio combining these three alpha sources.

In Equation (A.21) of the technical Appendix, we derive the intuitive result that the optimal portfolio is simply a weighted combination of the minimum-volatility portfolios for each alpha signal taken separately. We therefore treat each minimum-volatility factor portfolio as a separate "asset" and use mean–variance optimization to solve for the optimal weights. This, in turn, requires an asset covariance matrix and a set of expected returns.

Our first task is to construct the 3×3 asset covariance matrix. As discussed in the technical Appendix, our approach is to compute the covariance matrix directly using the GEM3 model. This is a straightforward exercise, since we know the holdings of each minimum-volatility portfolio.

Our next task is to estimate the expected returns for each minimum-volatility factor portfolio. Observe that the factor exposures of the minimum-volatility factor portfolio are known. For example, these are shown in Table 2 for the earnings yield factor. Note that only the alpha factors will contribute to the expected return of any portfolio. That is, since the other risk factors, by definition, have zero expected return, they do not contribute to the expected portfolio return.

To illustrate our technique, consider the following simple example. Suppose that the expected returns of the pure factor portfolios were known. For the sake of the argument, say that the expected returns were 300 bps for earnings yield, 100 bps for dividend yield, and 200 bps for book-to-price. From Table 2, we saw that the exposure of the earnings yield minimum-volatility factor portfolio was exactly 1.0 to earnings yield, 0.327 for dividend yield, and 0.183 for book-to-price. The expected return of the earnings yield minimum-volatility factor portfolio therefore would have been 370 bps. This is attributed as 300 bps from the earnings yield factor itself, plus 33 bps from dividend yield (0.327×100), and another 37 bps from book-to-price (0.183×200).

The final remaining task is to obtain the expected return of each pure alpha factor. To accomplish this, we follow a two-step process. First, we estimate the information ratio of each pure factor portfolio. We use the first five years of data to make our first estimate in January 2002. Every month, we add a new observation to our expanding window and update the information ratio estimate. By construction, such information ratios will evolve slowly across time. The second step is to scale the information ratio of the pure alpha factor by its GEM3 predicted volatility to obtain the expected return.

We do not make any claim that the procedure described above represents the "only" way or even the "best" way to compute the expected returns and covariance matrix of the minimum-volatility factor portfolios. Rather, we believe that our approach provides a reasonable and transparent method useful for illustrative purposes. Sophisticated asset managers will certainly employ prudent judgment and their investment insights to refine these methods according to their own views.

6 Empirical results

With the covariance matrix and expected returns of the three minimum-volatility factor portfolios determined, we can now easily compute the optimal set of portfolio weights. These weights, plotted in Figure 5, are quite stable. The main exception was during the "Quant Meltdown" of August 2007, which saw a sudden decline in the weight of the earnings yield portfolio. Over the course of the month, the volatility of the earnings yield factor nearly tripled, whereas the volatilities of dividend yield and book-to-price increased more modestly. From a mean–variance perspective, therefore, earnings yield temporarily became relatively less attractive. As the volatility of the earnings yield factor receded in the months following the Quant Meltdown, we see that the weight of the earnings yield factor increased accordingly.

The weights in Figure 5 are also quite intuitive when one considers the *ex ante* information ratios of the minimum-volatility factor portfolios. These are plotted in Figure 6 for the three alpha signals, as well as for the composite alpha factor. We see that earnings yield typically had the highest information ratio among the three signals, so it is not surprising that it should receive the largest weight in the optimization. Similarly, the dividend yield factor typically had the lowest information ratio and received the smallest weight. Figure 6 also illustrates the benefits of diversifying across multiple alpha sources. That is, the *ex ante* information ratio for the composite factor was significantly greater than the stand-alone information ratios of the individual signals.

In Table 4, we report the risk attribution for the optimal portfolio on analysis date 31-Dec-2012. As before, we show only a subset of representative factors. We see that only the three alpha factors contribute to portfolio risk, whereas the other risk factors contribute zero. This is consistent with the property that for an unconstrained optimal portfolio, the expected return contribution

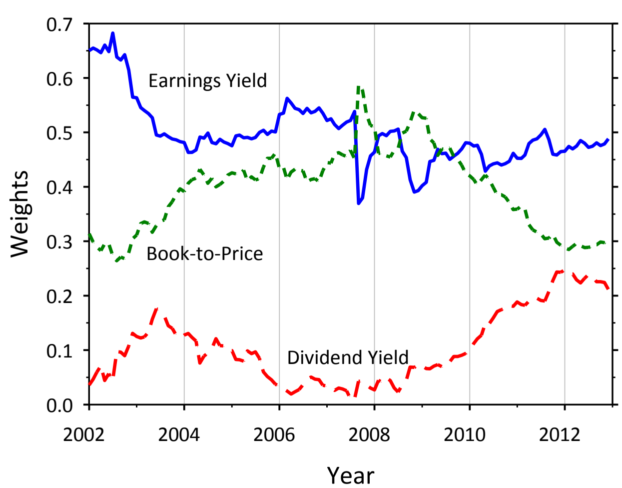

Figure 5 Weights of minimum-volatility factor portfolios versus time.

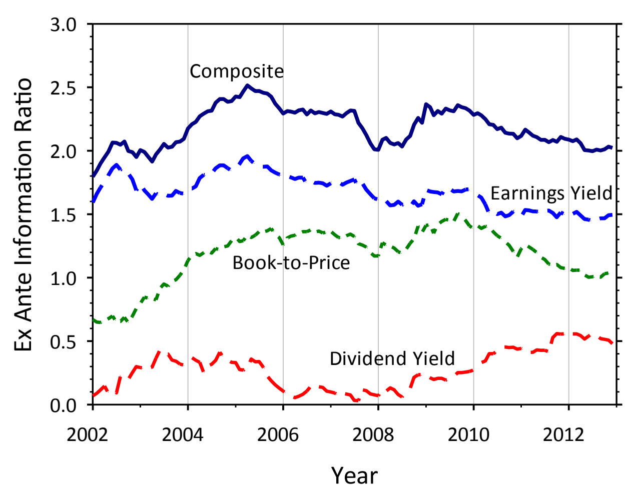

Figure 6 *Ex ante* information ratio of minimum-volatility portfolios for individual signals and composite alpha versus time.

Table 4 Risk attribution of optimal minimum-volatility factor portfolio, for analysis date 31-Dec-2012.

Factor	Exposure	Volatility	Correlation	Risk contrib
Earnings Yield	0.559	1.01	0.60	0.34
Book-to-Price	0.338	1.16	0.40	0.16
Dividend Yield	0.317	0.82	0.17	0.05
World	−0.004	11.25	0.00	0.00
Banks	0.005	3.66	0.00	0.00
Momentum	0.054	2.91	0.00	0.00
USA	0.010	3.64	0.00	0.00

is proportional to the risk contribution for every component of the portfolio. Since the other risk factors have zero expected return, it follows that they must contribute zero to the risk of the unconstrained optimal portfolio.

In Figure 7, we plot the cumulative performance of the optimal portfolio, represented by the solid blue line. We see that the portfolio earned a consistently positive return, with the exception of the period from mid-2007 until early 2009, during which time it was flat. Over the full sample period, the annualized return was 2.25 percent, with a volatility of only 87 bps, leading to an impressive information ratio of 2.78. Of course, this analysis does not include the effect of trading costs. Neither does it take into account real-world investment constraints such as long-only portfolios with monthly turnover limits. These constraints would serve to further reduce the information ratio that could be achieved in practice.

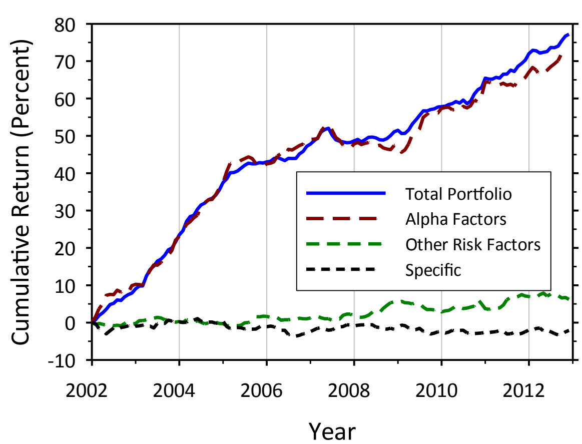

Figure 7 Return attribution of optimal minimum-volatility factor portfolio. The annualized return was 2.25 percent, with a volatility of 87 bps, leading to an information ratio of 2.78.

In Figure 7, we also plot the return contribution from three distinct sources: (a) the three alpha factors, (b) all other risk factors, and (c) the stock-specific component. We see that virtually all of the return came from the alpha factors, consistent with our portfolio design objective. The return contribution of the other risk factors was essentially flat over the entire sample period, reflecting the notion that these factors do not exhibit drift. It is also interesting to note that from July 2007 until February 2009, the alpha factors detracted nearly 600 bps from performance. Over this same time period, the other risk factors contributed more than 500 bps, thus largely mitigating the drawdown.

While the other risk factors made little impact on the cumulative returns of the portfolio, they were quite effective at reducing overall volatility levels. This is illustrated in Figure 8, which represents a scatterplot of monthly return contributions from alpha factors and all other risk factors. The realized correlation was −0.52, indicating that the

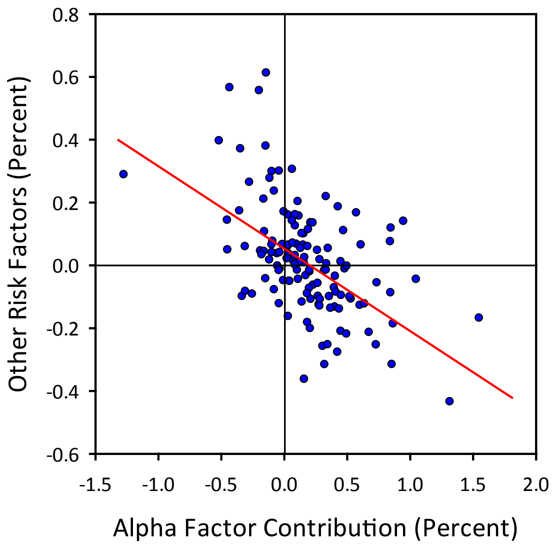

Figure 8 Scatterplot of monthly return contributions from alpha factors and from other risk factors, from January 2002 through December 2012. The realized correlation was −0.52. The regression fit through the points is also indicated.

other risk factors were highly effective at hedging the risk of the alpha factors. Under normal assumptions, the standard error of the correlation estimate is approximately 0.087. Thus, the realized correlation was also highly statistically significant.

7 Summary

We have discussed several approaches for capturing the return premium of an alpha factor. Simple factor portfolios tilt on the factor by overweighting positive alpha stocks and underweighting those with negative alpha. Such portfolios typically contain "incidental" exposures that may add considerable risk without enhancing the return profile. Another approach is to construct a pure factor portfolio by neutralizing exposures to all other factors. The most efficient implementation is given by the minimum-volatility factor portfolio, which exploits correlations between risk factors and alpha factors to reduce overall volatility levels.

We also presented a methodology for combining multiple sources of alpha. We showed that the optimal portfolio is a weighted combination of the minimum-volatility factor portfolios for each alpha signal. We provided a detailed example that illustrates our technique as well as builds the intuition behind these portfolios.

Appendix A

A.1 Simple factor portfolios

Simple factor portfolios are formed by univariate cross-sectional regression,

$$r_n = f_w + X_{nk} f_k^S + e_n, \qquad \text{(A.1)}$$

where r_n is the return of stock n, f_w is the world factor return (the regression intercept), X_{nk} is the stock exposure to factor k, f_k^S is the simple factor return, and e_n is the residual return. We assume that the exposures are standardized to be mean zero,

$$\sum_{n=1}^{N} w_n X_{nk} = 0, \qquad \text{(A.2)}$$

where w_n is the regression weight of stock n and N is the total number of stocks. We use regression weights that are proportional to the square root of market capitalization. The factor exposures are standardized to have unit standard deviation,

$$\sum_{n=1}^{N} w_n X_{nk}^2 = 1. \qquad \text{(A.3)}$$

Under this specification, the regression coefficient for the world factor is given by

$$f_w = \sum_{n=1}^{N} w_n r_n, \qquad \text{(A.4)}$$

which represents the regression-weighted mean return of the universe. The return of the simple factor portfolio is given by

$$f_k^S = \sum_{n=1}^{N} (w_n X_{nk}) r_n. \tag{A.5}$$

In other words, the weight of each stock in the simple factor portfolio is simply the product of the regression weight and the factor exposure. Equation (A.3) immediately shows that the simple factor portfolio has unit exposure to the underlying factor.

A.2 *Pure factor portfolios*

Pure factor portfolios are formed by multivariate cross-sectional regression,

$$r_n = \sum_{k=1}^{K} X_{nk} f_k^P + u_n, \tag{A.6}$$

where f_k^P is the return of the pure factor, u_n is the stock-specific return, and K is the total number of factors. The GEM3 factor structure contains two exact collinearities. Thus, we impose two constraints in order to obtain a unique regression solution: (a) the cap-weighted industry factor returns sum to zero, and (b) the cap-weighted country factor returns sum to zero. These constraints are embodied in the restriction matrix R, with dimensionality $K \times (K-2)$. Following the approach of Menchero (2010), we define a $K \times N$ matrix

$$\Omega^P = R(R'X'WXR)^{-1}R'X'W, \tag{A.7}$$

where W is the regression weighting matrix whose diagonal elements are given by w_n, and whose off-diagonal elements are zero. The pure factor returns are given by

$$f_k^P = \sum_{n=1}^{N} \Omega_{kn}^P r_n. \tag{A.8}$$

That is, the k^{th} row of matrix Ω^P gives the stock weights in the pure factor portfolios. Each pure factor style portfolio has unit exposure to the underlying factor, and zero exposure to all other factors.

A.3 *Minimum-volatility factor portfolios*

Minimum-volatility factor portfolios are formed by mean–variance optimization.

$$\Omega_k^{MV} = \frac{V^{-1}X_k}{X_k' V^{-1} X_k}, \tag{A.9}$$

where V is the $N \times N$ asset covariance matrix, and X_k is the k^{th} column vector of factor exposure matrix X. The denominator in Equation (A.9) ensures that the resulting portfolio has unit exposure to the underlying factor. We obtain the asset covariance matrix using the GEM3 risk model. Minimum-volatility factor portfolios have the interesting property that they are uncorrelated with all pure factor portfolios, except their own underlying factor. This can be easily seen as follows:

$$\mathrm{cov}(\Omega_m^P, \Omega_k^{MV}) = \frac{(\Omega_m^P)' V (V^{-1} X_k)}{X_k' V^{-1} X_k}. \tag{A.10}$$

If $k \neq m$, Equation (A.10) shows that the covariance is zero, since pure factor portfolio m has zero exposure to factor k. According to the *x-sigma-rho* formula, described by Menchero and Davis (2011), the risk contribution is proportional to the correlation with the portfolio. Consequently, pure factor portfolios contribute zero to the risk of the minimum-volatility factor portfolio. For the case $k = m$, of course, the contribution is non-zero.

A.4 *Alpha factors and risk factors*

We define an *alpha factor* to be any factor whose pure factor portfolio has non-zero expected return,

$$E[f_k^{P(\alpha)}] = \alpha_k^P, \tag{A.11}$$

where α_k^P denotes the expected return of the pure alpha factor portfolio. The other risk factors are assumed to have an expected return of identically zero,

$$E[f_k^{P(\sigma)}] = 0. \tag{A.12}$$

Segmenting factors into alpha factors (α) and the other risk factors (σ), Equation (A.6) can be rewritten as

$$r_n = \sum_{k=1}^{K_\alpha} X_{nk}^{(\alpha)} f_k^{P(\alpha)} + \sum_{k=1}^{K_\sigma} X_{nk}^{(\sigma)} f_k^{P(\sigma)} + u_n, \tag{A.13}$$

where K_α is the total number of alpha factors and K_σ is the total number of other risk factors. The optimizer can exploit covariances between the alpha factors and the other risk factors to reduce portfolio volatility. Furthermore, since the other risk factors have zero expected return, the volatility reduction comes without incurring a price in expected portfolio performance.

A.5 Combining multiple alpha signals

We write the composite alpha signal as a linear combination of the individual alpha signals,

$$\alpha_n = \sum_{k=1}^{K_\alpha} \upsilon_k X_{nk}^{(\alpha)}. \tag{A.14}$$

Our task is to solve for the coefficients υ_k. The optimal portfolio is given by the standard expression,

$$h = \frac{V^{-1}\alpha}{\lambda}, \tag{A.15}$$

where λ is the risk-aversion parameter and α is the $N \times 1$ vector whose elements are given by α_n. Note that since h has the lowest volatility for a given level of alpha, it represents the maximum information ratio portfolio. Substituting Equation (A.14) into Equation (A.15), we obtain

$$h = \frac{1}{\lambda} V^{-1} \left[\sum_{k=1}^{K_\alpha} \upsilon_k X_k^{(\alpha)} \right], \tag{A.16}$$

which can be rewritten as

$$h = \frac{1}{\lambda} \left[\sum_{k=1}^{K_\alpha} \upsilon_k V^{-1} X_k^{(\alpha)} \right]. \tag{A.17}$$

The minimum-volatility portfolio for alpha signal k can be expressed as an $N \times 1$ vector

$$h_k = \frac{V^{-1} X_k^{(\alpha)}}{s_k}, \tag{A.18}$$

where s_k is a scalar determined by Equation (A.9),

$$s_k = X_k^{\prime(\alpha)} V^{-1} X_k^{(\alpha)}. \tag{A.19}$$

Now substituting Equation (A.18) into Equation (A.17), we find

$$h = \frac{1}{\lambda} \left[\sum_{k=1}^{K_\alpha} \upsilon_k s_k h_k \right]. \tag{A.20}$$

This states that the optimal portfolio is a linear combination of minimum-volatility factor portfolios for the individual alpha signals. Therefore, we can write the optimal portfolio as

$$h = \sum_{k=1}^{K_\alpha} w_k h_k, \tag{A.21}$$

where w_k is the weight of minimum-volatility factor portfolio k. Comparing Equation (A.21) with Equation (A.20), we can relate the portfolio weights to the alpha weights

$$\upsilon_k = \frac{\lambda w_k}{s_k}. \tag{A.22}$$

The portfolio weights w_k can be solved via mean–variance optimization. The "assets" in this case represent the minimum-volatility factor portfolios. Consequently, we have reduced the optimization problem from N dimensions to the much smaller dimensionality K_α of the alpha factors.

To apply mean–variance optimization, we require the asset covariance matrix and the asset expected returns. We compute the asset covariance matrix directly

$$G_{kl} = h_k' V h_l, \qquad \text{(A.23)}$$

where V is the $N \times N$ asset covariance matrix from the GEM3 risk model.

To estimate the asset expected returns, first note that the exposure X_{km}^{MV} of minimum-volatility factor portfolio k to factor m is known. Since only alpha factors contribute to the expected return of any portfolio, we can write

$$\alpha_k^{MV} = \sum_{m=1}^{K_\alpha} X_{km}^{MV} \alpha_m^P. \qquad \text{(A.24)}$$

A reasonable way to obtain the expected return of the pure alpha factor is to estimate the information ratio of the pure alpha factor portfolio over a trailing window and then to multiply this information ratio by the predicted volatility of the pure factor portfolio. The optimal set of weights is then given by the standard solution,

$$w = \frac{1}{\lambda} G^{-1} \alpha^{MV}. \qquad \text{(A.25)}$$

Reverting to summation notation, this can be expressed as

$$w_k = \frac{1}{\lambda} \sum_{j=1}^{K_\alpha} G_{kj}^{-1} \alpha_j^{MV}. \qquad \text{(A.26)}$$

Plugging Equation (A.26) into Equation (A.22) allows us to solve for the weights v_k

$$v_k = \frac{1}{s_k} \sum_{j=1}^{K_\alpha} G_{kj}^{-1} \alpha_j^{MV}. \qquad \text{(A.27)}$$

This expresses the optimal set of alpha weights in terms of quantities that can be easily computed or estimated. Also note that the final result is independent of the risk-aversion parameter, as should be expected.

Notes

1 In this article, we use the terms "alpha signal" and "alpha factor" interchangeably.

2 Grinold and Kahn (2000) also refer to this as the "characteristic" factor portfolio.

3 To avoid undue influence from periods of high volatility, we compute the information ratios from monthly z-scores. This explains why the information ratios may deviate slightly from the conventional information ratio, defined as annualized return divided by annualized risk. The conventional information ratios are easily computed from Table 3.

References

Ang, A., Hodrick, R., Xing, Y., and Zhang, X. (2009). "High Idiosyncratic Volatility and Low Returns: International and Further U.S. Evidence," *Journal of Financial Economics* **91**, 1–23.

Banz, R. (1981). "The Relationship Between Return and Market Value of Common Stock," *Journal of Financial Economics* **9**, 3–18.

Basu, S. (1977). "Investment Performance of Common Stocks in Relation to Their Price-Earnings Ratio: A Test of the Efficient Market Hypothesis," *Journal of Finance* **32** (June), 663–682.

Grinold, R. (2010). "Signal Weighting," *Journal of Portfolio Management* **36**(4) (Summer), 24–34.

Grinold, R. and Kahn, R. (2000). *Active Portfolio Management* (New York: McGraw-Hill).

Jegadeesh, N. and Titman, S. (1993). "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance* **48**, 65–92.

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

Menchero, J. and Davis, B. (2011). "Risk Contribution is Exposure Times Volatility Times Correlation," *Journal of Portfolio Management* **37**(2) (Winter), 97–106.

Morozov, A., Wang, J., Borda, L., and Menchero, J. (2012). "The Barra Global Equity Model Empirical Notes," *MSCI Model Insight* (January).

Qian, E., Hua, R., and Sorensen, E. (2007). *Quantitative Equity Portfolio Management* (Boca Raton: Chapman & Hall).

Keywords: Portfolio construction; alpha estimation; mean-variance optimization