JOIM
www.joim.com

Journal Of Investment Management, Vol. 15, No. 2, (2017), pp. 92–107

RETHINKING THE FUNDAMENTAL LAW OF ACTIVE MANAGEMENT

*Jose Menchero**

The fundamental law of active management provides a powerful framework for analyzing portfolio diversification and risk-adjusted returns. It states that the information ratio of an unconstrained optimal portfolio is given by the product of the information coefficient (a measure of skill) and the square root of breadth, where breadth is the number of "independent" bets. A basic limitation of previous formulations of the fundamental law is that it was not possible to determine portfolio breadth for realistic portfolios under a general covariance structure. In this paper, we present a new formulation of the fundamental law of active management. We derive a new measure of skill, denoted the Signal Quality, and obtain an exact closed-form expression for the square root of breadth, which we denote as the Diversification Coefficient. Our formulation is easily applied to real-world portfolios described by general covariance matrices. We conclude with a discussion of the transfer coefficient, which measures the drop in portfolio efficiency due to investment constraints.

The primary objective of disciplined investing is to achieve the highest possible risk-adjusted performance. For active managers, this is measured by the *Information Ratio* (*IR*), which represents the portfolio outperformance (relative to a benchmark) divided by the risk taken to achieve the outperformance. Asset owners also make use of the Information Ratio, both as a key determinant in manager selection, and for deciding the optimal asset allocation to maximize overall risk-adjusted performance.

*This work was completed at Menchero Portfolio Analytics Consulting. Jose currently serves as Head of Portfolio Analytics Research at Bloomberg.

In a pioneering work, Grinold (1989) introduced the fundamental law of active management, in which he decomposed the portfolio *IR* into a product of two terms,

$$IR = IC\sqrt{BR}. \tag{1}$$

The first term, which he called the *Information Coefficient* (*IC*), represents the skill of the portfolio manager (i.e., their ability to accurately forecast asset returns). The second term is given by the square root of *Breadth*, where Breadth represents the number of "independent" bets placed by the portfolio manager. Grinold and Kahn (2000) later codified this formulation in their celebrated book, *Active Portfolio Management.*

The main lesson of the fundamental law is clear: In order to achieve high risk-adjusted performance, the portfolio manager should play well (high IC) and place as many independent bets as possible (i.e., large Breadth). The fundamental law thus represents a powerful framework for analyzing investment problems. For instance, suppose a large-cap equity manager considers extending a successful strategy into the small-cap space. In theory, expanding the investment universe always leads to an increase in the expected, or *ex ante*, Information Ratio. The *magnitude* of the increase, however, depends on the extent to which the new bets are truly "independent" of the existing bets; clearly, if the two bets are highly correlated, any benefit may be minimal. Assuming that Breadth can be reliably estimated, the fundamental law can help answer whether the expected gain in risk-adjusted performance is worth the extra costs associated with extending the strategy to a broader universe. A similar concept was explored by Goldsticker (2013), who used the fundamental law to analyze the tradeoffs between high skill within a narrow universe, versus lower skill across a wider universe.

It is important to bear in mind that the Information Ratio analyzed by Grinold and Kahn represents the maximum possible Information Ratio (*ex ante*), given a set of alpha forecasts. This implicitly assumes that the portfolio is constructed using mean–variance optimization, free of investment constraints. In the real world, however, investment constraints invariably lead the manager to hold a portfolio that differs from the theoretical optimal portfolio.

Clarke *et al.* (2002) provided an important generalization to the fundamental law by introducing the notion of the *Transfer Coefficient* (TC), which quantifies the drop in portfolio efficiency due to investment constraints. The fundamental law, stated in its generalized form, is given by

$$IR = TC \cdot IC\sqrt{BR}, \qquad (2)$$

where IR is the Information Ratio of the portfolio held by the manager, IC is the manager's Information Coefficient, and BR is the Breadth. Note that $TC = 1$ for the unconstrained optimal portfolio, whereas $TC < 1$ for all other portfolios.

The Transfer Coefficient represents an important tool for evaluating the impact of investment constraints. For instance, unconstrained optimal portfolios often contain negative weights in securities that cannot be shorted under the long-only constraint. The Transfer Coefficient enables the manager to quantify the impact of this constraint on portfolio efficiency. More specifically, by comparing Transfer Coefficients across different sets of constraints, the manager can evaluate the impact of lifting a particular constraint.

While the fundamental law provides a powerful theoretical framework for analyzing investment problems, to be effectively applied in practice, skill (IC) and Breadth must be reliably estimated. Unfortunately, under the Grinold–Kahn formulation, certain unrealistic assumptions make it difficult to estimate these quantities in practice. For instance, Grinold and Kahn assume that the manager has equal forecasting skill across every asset, whereas in reality the manager may be more skillful in certain market segments than others. Similarly, whereas Breadth is simply equal to the number of assets under the assumption of zero correlation, it is not clear how to compute Breadth under more realistic scenarios.

In this paper, we present a new formulation of the fundamental law. Specifically, our formulation consists of two main contributions. First, by starting with the basic definition of IR, and recognizing that skill represents the manager's ability to forecast risk-adjusted returns, we derive a new measure of skill, denoted the *Signal Quality*. This new measure of skill allows the fundamental law to be easily derived in exact form without relying on the set of simplifying assumptions and approximations used by other formulations. As a result, our new formulation is both simpler and more general than previous formulations. Second, our analysis allows us to derive an exact closed-form expression for Breadth, which we reinterpret and relabel as the *Diversification Coefficient*. This resolves a long-outstanding practical issue in the application of the fundamental law, namely, how to compute the Breadth of a real portfolio. Together, these two contributions allow the fundamental law to be easily applied to real-world portfolios.

The remainder of this paper is organized as follows. First, we provide a brief review of previous formulations of the fundamental law. We then present our new formulation of the fundamental law. We also provide several illustrative examples showing how it can be applied in practice. We conclude with a discussion of the Transfer Coefficient. Finally, to preserve the readability and flow of the main text, mathematical details have been relegated to several technical appendices.

1 Previous formulations of the fundamental law

In this section, we briefly review previous formulations of the fundamental law. We begin with the formulation of Grinold and Kahn (2000), and trace the development until the most recent work, by Ding and Martin (2015).

Grinold and Kahn. The starting point in Grinold and Kahn is to segment asset returns into a benchmark component and a residual component,

$$r_{nt} = \beta_{nt} R_t^B + e_{nt}, \tag{3}$$

where r_{nt} is the return of asset n over period t, β_{nt} is the beta of the asset relative to the benchmark, R_t^B is the return of the benchmark over the period, and e_{nt} is the residual return. Note that residual returns, by construction, are uncorrelated with the benchmark.

A central result of Modern Portfolio Theory is that if the benchmark is mean–variance efficient, then the expected return of any asset is given by the beta of the asset multiplied by the expected return of the benchmark. This would imply that the residual returns are mean zero. Active managers, of course, do not believe that the benchmark is efficient. From their viewpoint, the residuals have non-zero expected returns, denoted *alpha*, i.e.,

$$\alpha_{nt} \equiv E[e_{nt}]. \tag{4}$$

Philosophically, the difference between passive and active investors is that the former believe that alphas are equal to zero, whereas the latter believe that alphas are non-zero and can be exploited to outperform the benchmark.

In the Grinold–Kahn formulation, given by Equation (1), the IC for a particular stock n is computed by the time-series correlation

$$IC_n = \text{corr}(\alpha_{nt}, e_{nt}), \tag{5}$$

where α_{nt} is the start-of-period forecast for the residual return, and e_{nt} is the realized residual return over the period.

Grinold and Kahn assume that all stocks have the same IC. Although it is possible to relax this rather unrealistic assumption, one then loses the simplicity of Equation (1). In practice, managers will often average the IC across stocks to obtain a composite portfolio value

$$IC = \frac{1}{N}\sum_{n=1}^{N} IC_n. \tag{6}$$

Another important element of the Grinold–Kahn formulation is to assume that the stock alphas adhere to a certain structure. More specifically, the alphas are assumed to be of the form

$$\alpha_n = \sigma_n \cdot IC \cdot z_n, \tag{7}$$

where σ_n is the volatility of stock n, and z_n is a z-score for the stock, standardized to have mean zero and unit variance. Equation (7) represents the famous scaling rule described by Grinold (1994), known as "alpha equals volatility times IC times score."

One shortcoming of the Grinold–Kahn formulation is that unless the assets are strictly uncorrelated, it is unclear how to compute Breadth. Many practitioners incorrectly assume that the Breadth represents the number of stocks in the portfolio, which would only be true if the residual returns were mutually uncorrelated. In practice, making this rather unrealistic assumption typically results in overly optimistic Information Ratios.

It is also important to stress that the alpha scaling rule, given by Equation (7), is not universally applicable. The form of this scaling rule results from estimating alphas by a time-series regression of asset returns against the z-scores. Nevertheless, some practitioners blindly apply the rule, even when inappropriate. Specifically, if alphas are not estimated by time-series regression, there is no basis for applying the alpha scaling rule. For instance, a portfolio manager may bucket stocks into five groupings: strong buy, buy, hold, sell, and strong sell. The manager may assign alphas of two percent to the strong buys, one percent to the buys, zero for holds, and so on. Although such an alpha model may be simplistic, it nevertheless represents a valid set of alphas since it faithfully reflects the views of the portfolio manager.[1] Hence, rescaling the alphas in this case would not be warranted. Another example in which alpha rescaling is not warranted is the estimation technique described by Menchero and Lee (2015), who used a multi-factor cross-sectional approach to estimate stock alphas.

Buckle. Another significant advance in the development of the fundamental law was due to Buckle (2004), who derived a closed-form solution for portfolio Breadth. Buckle's derivation rested upon several key assumptions, including: (1) return forecasts are unbiased, (2) there are no long-run abnormal returns, (3) return forecasts and their errors are independent, (4) return forecasts are normally distributed, (5) all assets have the same IC, and (6) the IC is small. Subject to these assumptions, Buckle shows that the Breadth is given by

$$BR = \sum_{i=1}^{N} \sum_{j=1}^{N} \rho_{ij} P_{ij}^{-1}, \tag{8}$$

where ρ_{ij} is the correlation between return forecasts i and j, and P_{ij}^{-l} denotes the element in row i and column j of the inverse asset *correlation* matrix. Buckle then goes on to illustrate the formula for several idealized scenarios in which the correlation ρ_{ij} is pre-specified.

While Buckle's formula for Breadth represents an important theoretical result, it suffers from two basic shortcomings. First, it rests upon a number of assumptions, which may be violated in practice to varying degree. Second, whereas P_{ij}^{-1} may be easily derived from the asset covariance matrix, it is not obvious how to estimate the correlation ρ_{ij} for a general set of forecasts. This limits the usefulness of the Buckle formula in practice.

Qian, Hua, and Sorensen (QHS). Another important formulation of the fundamental law is due to Qian *et al.* (2007). In the QHS formulation, the IC for period t is defined as the *cross-sectional* correlation between the risk-adjusted forecasts and the risk-adjusted realizations

$$IC_t = \text{corr}\left(\frac{\alpha_{nt}}{\sigma_{nt}}, \frac{e_{nt}}{\sigma_{nt}}\right), \tag{9}$$

where σ_{nt} represents the volatility of stock n at time t. Note that *ex post*, the IC will be negative for periods in which the signal "didn't work." *Ex ante*, however, the *IC* must be positive to justify active management. In practice, the forecast *IC* is often computed by assuming that the future will be similar to the past, and simply averaging over a suitably long back-testing window

$$IC = \frac{1}{T}\sum_{t=1}^{T} IC_t. \tag{10}$$

The QHS formulation rests upon several assumptions, including: (1) the portfolio has zero factor risk, (2) the asset-level *IRs* are cross-sectionally mean zero, and (3) the risk-adjusted residual returns are mean zero. Subject to these assumptions, Qian *et al.* derive a formula for the portfolio *IR*,

$$IR = \frac{IC\sqrt{N}}{\sqrt{1 - IC^2}}. \tag{11}$$

This result looks tantalizingly similar to the Grinold–Kahn form, except for the appearance of $\sqrt{1 - IC^2}$ in the denominator. It should be noted that the absence of this term in Grinold and Kahn results from an approximation that assumes the *IC* is small (as is typical for a universe of stocks), in which case the denominator in Equation (11) may be safely ignored.

An attractive feature of the QHS formulation is that many practitioners prefer to think of the *IC* as a cross-sectional measure across a universe of stocks, as opposed to a time-series measure for an individual stock as in the Grinold–Kahn formulation. Despite this, it should be noted that the validity of the QHS formulation ultimately rests upon several assumptions that may be violated in practice. For instance, the assumption that asset-level *IRs* are cross-sectionally mean zero may be a good approximation for a universe of stocks, but it breaks down in other applications, such as a fund-of-funds investment problem in which the "assets" constitute individual hedge funds, each with a positive *IR*. Another significant limitation is the assumption of zero factor risk; in reality, virtually every alpha signal will have some degree of systematic factor risk associated with it.

Ding and Martin. Another significant contribution to the development of the fundamental law is due to Ding and Martin (2015), who relax the assumption of zero factor risk. They introduced a one-factor model in which the "factor returns" take the form of Information Coefficients. The model is estimated by performing cross-sectional regressions every period, leading to a time series of Information Coefficients, IC_t. They then derive a formula for the portfolio *IR*,

$$IR = \frac{IC\sqrt{N}}{\sqrt{1 - IC^2 + N\sigma_{IC}^2}}, \tag{12}$$

where IC and σ_{IC}^2 are the mean and variance of IC_t, respectively. In the absence of factor risk, $\sigma_{IC} = 0$, and Equation (12) reduces to Equation (11) as a special limiting case. Another interesting limit occurs when the number of assets approaches infinity. In this case,

$$IR_{\max} = \frac{IC}{\sigma_{IC}}, \tag{13}$$

which represents the maximum attainable *IR*. In other words, the Ding–Martin formulation shows that factor risk imposes an upper bound on the *IR*. This crucial insight is not evident in either the Grinold–Kahn or the QHS formulations.

The main drawback of the Ding–Martin formulation is that their analytic result was derived using a highly idealized one-factor model. Although the model can be extended to multiple factors, such an extension adds considerable complexity.

2 New formulation of the fundamental law

In this section, we present a new formulation of the fundamental law. Our only assumption is that the $N \times N$ asset covariance matrix $\boldsymbol{\Omega}$ and the

$N \times 1$ expected return vector $\boldsymbol{\alpha}$ are known. We first derive an exact expression for the *IR* of an unconstrained optimal portfolio. We then decompose the *IR* into a product of skill and Breadth, offering a new interpretation for each of these measures.

To find the optimal portfolio, we apply mean–variance optimization. Our utility function contains a reward for portfolio alpha, with a penalty for portfolio variance,

$$U = \mathbf{h}'\boldsymbol{\alpha} - \lambda \mathbf{h}'\boldsymbol{\Omega}\mathbf{h}, \tag{14}$$

where λ is the risk-aversion parameter. Let Y denote the unconstrained optimal portfolio, with the risk-aversion parameter selected to produce a portfolio with volatility σ_Y. The $N \times 1$ holdings vector of portfolio Y is given by

$$\mathbf{h}_Y = \frac{\sigma_Y \boldsymbol{\Omega}^{-1}\boldsymbol{\alpha}}{\sqrt{\boldsymbol{\alpha}'\boldsymbol{\Omega}^{-1}\boldsymbol{\alpha}}}. \tag{15}$$

The reader may easily verify that $\mathbf{h}_Y'\boldsymbol{\Omega}\mathbf{h}_Y = \sigma_Y^2$, as required. It can be shown that portfolio Y has a beta of zero relative to the benchmark. This is an intuitive result, since any component aligned with the benchmark adds risk to the portfolio while leaving the portfolio alpha unchanged.

It is worth pointing out that the solution for the optimal portfolio involves the inverse of the asset covariance matrix. If this is computed naively (e.g., using the sample covariance matrix), then the matrix is not invertible whenever the number of assets exceeds the number of time periods. To overcome this difficulty, asset covariance matrices are typically estimated using multi-factor risk models. Such models are well behaved and invertible so long as the number of time periods exceeds the number of factors, as is typically the case.

The expected return of portfolio Y is the inner product of the holdings vector with the alpha vector, i.e., $\boldsymbol{\alpha}'\mathbf{h}_Y$, thus giving

$$E[R_Y] = \sigma_Y \sqrt{\boldsymbol{\alpha}'\boldsymbol{\Omega}^{-1}\boldsymbol{\alpha}}. \tag{16}$$

Since the *IR* is the ratio of expected return to volatility, the *IR* of portfolio Y follows immediately from Equation (16),

$$IR_Y = \sqrt{\boldsymbol{\alpha}'\boldsymbol{\Omega}^{-1}\boldsymbol{\alpha}}. \tag{17}$$

Equation (17) represents the exact expression for the *IR* of an unconstrained optimal portfolio. This well-established result is given by Equation (5A.6) in Grinold and Kahn (2000), and serves as the launch point for our new formulation of the fundamental law.

Next, we follow the basic structure of Grinold and Kahn and write the portfolio *IR* as a product of two terms representing skill and Breadth. However, since we offer a new interpretation for these quantities, we also adopt a different terminology. Specifically, we write

$$IR_Y = Q \cdot D, \tag{18}$$

where Q denotes the *Signal Quality* and D is the *Diversification Coefficient*. The Signal Quality specifies the strength of the risk-adjusted alphas and is used to measure the skill of the portfolio manager. The Diversification Coefficient is akin to the square root of Breadth and represents the gain in *IR* that can be achieved through optimal allocation of the risk budget. Note that for N uncorrelated assets, the Diversification Coefficient must satisfy $D = \sqrt{N}$, consistent with other formulations of the fundamental law.

The main job of the portfolio manager is to identify assets with superior risk-adjusted returns. Notably, portfolio managers are not in the business of forecasting asset correlations; this task is typically outsourced to a third-party risk provider. Therefore, *any measure of manager skill should be independent of asset correlations*. This implies that if we can solve for manager skill using *one* set of asset correlations, we have found manager skill for *any* set of asset correlations.

Fortunately, it is quite easy to solve for skill in the special case of uncorrelated assets. In this case, the asset covariance matrix $\boldsymbol{\Omega}$ is diagonal, and Equation (17) reduces to

$$IR_Y^{Diag} = \sqrt{\sum_{n=1}^{N} (\alpha_n/\sigma_n)^2}. \tag{19}$$

We now use Equation (18) to solve for the skill, given that $D = \sqrt{N}$ for uncorrelated assets,

$$Q = \sqrt{\frac{1}{N}\sum_{n=1}^{N} (\alpha_n/\sigma_n)^2}. \tag{20}$$

Hence, we reinterpret skill as the root-mean-square (RMS) of the *IRs* of the underlying assets, i.e., $Q = \text{RMS}(\alpha_n/\sigma_n)$. Note that this definition of skill is easily applied to important use-cases such as fund-of-funds investing, in which all assets have positive *IR*.

Next, we compute the Diversification Coefficient for the general case of correlated assets. This is simply the ratio of IR_Y to the Signal Quality Q. Hence, the Diversification Coefficient is given by

$$D = \frac{\sqrt{\boldsymbol{\alpha}'\boldsymbol{\Omega}^{-1}\boldsymbol{\alpha}}}{\text{RMS}(\alpha_n/\sigma_n)}. \tag{21}$$

Note that the Diversification Coefficient possesses three essential properties. First, D is *scale invariant*, meaning that if we scale the alphas by any multiplicative constant, the Diversification Coefficient remains unchanged. Second, $D \geq 1$, meaning that diversification always provides a benefit.[2] Finally, it is easy to verify that $D = \sqrt{N}$ for the case of uncorrelated assets.

It is also interesting to note that Equation (21) can be expressed in terms of the Rayleigh Quotient, defined as

$$R(\mathbf{M}, \mathbf{x}) = \frac{\mathbf{x}'\mathbf{M}\mathbf{x}}{\mathbf{x}'\mathbf{x}}, \tag{22}$$

where $\mathbf{M}$ denotes an $N \times N$ symmetric matrix and $\mathbf{x}$ represents an $N \times 1$ vector. To ascertain the relationship between Breadth and the Rayleigh Quotient, we begin by writing the asset covariance matrix $\boldsymbol{\Omega}$ as a product of volatilities and correlations,

$$\boldsymbol{\Omega} = \boldsymbol{\sigma}\mathbf{P}\boldsymbol{\sigma}, \tag{23}$$

where $\boldsymbol{\sigma}$ is an $N \times N$ diagonal matrix of volatilities σ_n, and $\mathbf{P}$ denotes the $N \times N$ asset correlation matrix. The inverse of the asset covariance matrix may thus be written as

$$\boldsymbol{\Omega}^{-1} = \boldsymbol{\sigma}^{-1}\mathbf{P}^{-1}\sigma^{-1}. \tag{24}$$

Next, we define an $N \times 1$ vector of risk-adjusted alphas

$$\mathbf{a} = \boldsymbol{\sigma}^{-1}\boldsymbol{\alpha}. \tag{25}$$

Using these relationships, Equation (21) is easily manipulated into the following form

$$D^2 = N\frac{\mathbf{a}'\mathbf{P}^{-1}\mathbf{a}}{\mathbf{a}'\mathbf{a}}. \tag{26}$$

Equation (26) states that the squared Diversification Coefficient (i.e., the Breadth) represents the number of assets N multiplied by the Rayleigh Quotient involving the risk-adjusted alphas and the inverse asset correlation matrix. The Rayleigh Quotient has several interesting properties, including that it is maximized when $\mathbf{a}$ corresponds to the eigenvector of $\mathbf{P}^{-1}$ with the largest eigenvalue. Further explorations of the Rayleigh Quotient, however, are beyond the scope of this paper.

2.1 *Example* 1: *Optimal two-asset portfolio*

In this example, we apply our formulation of the fundamental law to analyze a fund-of-funds manager who uses mean–variance optimization to allocate between two hedge funds A and B. Five inputs are required to fully specify the optimization problem: (a) the expected returns α_A and α_B of the two funds, (b) the volatilities σ_A and σ_B of the funds, and (c) the correlation ρ between the funds. Technical details are provided in Appendix A.

For simplicity, we assume that each fund has a volatility of 10 percent. We further assume that Fund A has an expected return of 11.3 percent, versus 8.5 percent for Fund B. These values were conveniently selected to produce a Signal Quality of exactly 1.

We first provide some intuition on the portfolio composition. In Figure 1, we plot the optimal fund weights as a function of the correlation between the funds. The weights are standardized so that they sum to unity. Note that for large negative correlation, the fund weights are essentially equal. This is because the two funds become near-perfect hedges for one another, thus virtually eliminating portfolio risk while still capturing the positive return premium of each fund. As the correlation increases, the optimal portfolio shifts more weight into Fund A, which has superior risk-adjusted performance. Eventually, if the correlation between the funds becomes sufficiently large, the optimal portfolio takes a short position in Fund B. This serves as an effective hedge for Fund A, while still allowing the portfolio to capture the return difference between the funds.

In Appendix A, we derive exact expressions for the expected return, risk, and *IR* of the optimal two-asset portfolio. These quantities are plotted in Figure 2 as a function of fund correlation. For large negative correlations, the expected return is essentially constant, since the portfolio splits the weights nearly equally across the two funds. The risk, however, declines to zero, since Fund B becomes a near-perfect hedge for Fund A as the correlation approaches -1. As a result, the *IR* increases rapidly for strong negative correlations. By contrast, as the correlations become strongly positive, both the expected return and the risk rise dramatically due to increasing leverage. Nevertheless, expected return increases more rapidly than risk, causing the *IR* to increase for correlations above roughly 0.90.

Next, we study how the Diversification Coefficient of the portfolio depends on the correlation between the funds. Since $Q = 1$ in this example, it follows that the Diversification Coefficient is equal to the portfolio *IR*. Hence, we see from Figure 2 that the Diversification Coefficient

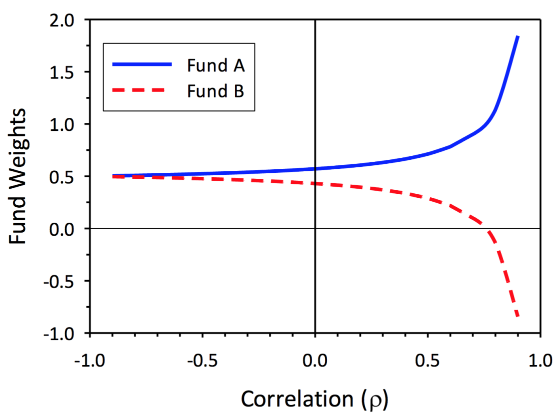

Figure 1 Optimal fund weights versus fund correlation.

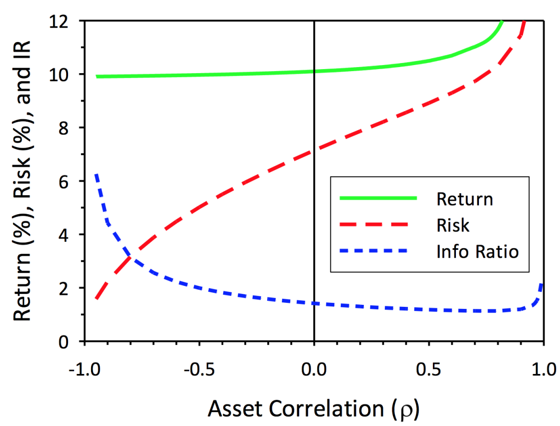

Figure 2 Return, risk, and Information Ratio (*IR*) of optimal portfolio versus fund correlation.

reaches a minimum for a correlation of roughly 0.90. If the two funds are uncorrelated, the Diversification Coefficient is exactly equal to $\sqrt{2}$, as required for two assets. Note that for negative correlations, the Diversification Coefficient is always greater than $\sqrt{2}$. For this reason, we prefer not to think of the Diversification Coefficient in terms of the number of "independent" bets, since this language would imply that it is possible to form more than two independent bets from only two assets. By contrast, it is quite intuitive that the Diversification Coefficient would exceed $\sqrt{2}$ when the correlations are negative, as this constitutes a "dream scenario" in which the manager has identified two negatively correlated funds, each with positive expected returns.

This example illustrates the ease with which our formulation of the fundamental law can be applied in practice. It is worth stressing, once again, that in previous formulations it was not clear how to apply the fundamental law for general sets of asset correlations, or in situations in which all assets had positive expected returns, such as a fund-of-funds investment problem.

2.2 *Example* 2: *One-factor model*

In this example, we study the effect of factor risk on the Diversification Coefficient. We consider an optimal portfolio containing N assets. We assume that both risk and expected returns are driven by a single underlying factor, with volatility σ_F. To make the model analytically tractable, we assume that all stocks have the same specific risk σ_S. In Appendix B, we provide technical details on how to compute the Signal Quality and Diversification Coefficient for this problem.

In Figure 3, we plot the Diversification Coefficient versus $\sqrt{N}$ for several values of the ratio σ_F/σ_S. If the factor risk is identically zero, all assets are mutually uncorrelated, in which case D increases linearly with $\sqrt{N}$, as required.

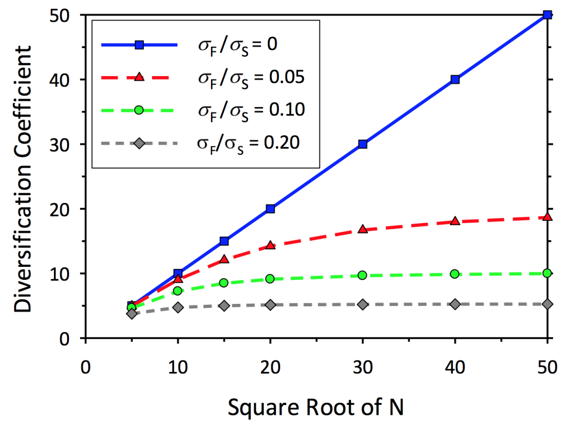

Figure 3 Diversification Coefficient versus the square root of the number of assets. The various lines correspond to different levels of factor risk relative to specific risk.

However, even modest levels of factor risk can have a major impact on the Diversification Coefficient. For instance, as the factor volatility increases from zero to merely 5 percent of the specific risk,[3] the Diversification Coefficient for a universe of 2,500 assets drops from 50 to below 20. This implies that neglecting the role of factor risk would lead to an overly optimistic *IR* by more than a factor of two. Finally, it is worth noting that as the factor volatility goes to infinity, the Diversification Coefficient goes to 1, as discussed in Appendix B. In this case, factor risk dominates all else and there is no benefit to diversifying across stocks.

2.3 *Example* 3: *Multi-factor model of US equity market*

In this example, we study the Diversification Coefficient for a variety of factors across different segments of the US equity market. To construct the optimal portfolios, we use the Axioma US equity risk model, which is estimated by cross-sectional regression against a set of industry dummies and 11 style factors. For analysis purposes, we treat the 11 style factors contained within the Axioma model as our "alpha factors" and construct unconstrained optimal portfolios with unit exposure to alpha.

In Table 1, we list the 11 style factors in the risk model, together with the annualized *ex ante* factor volatilities as of April 17, 2015. We also report the Diversification Coefficient, computed using Equation (21), for three investment universes: the S&P 500, the Russell 1000, and the Russell 3000. Note that for the S&P 500, several factors (e.g., growth) have Diversification Coefficient that exceeds the $\sqrt{N}$ limit for uncorrelated assets. This is certainly possible, since the Diversification Coefficient depends in part on the correlations between the factors. For larger universes, such as the Russell 3000, none of the factors exceed the $\sqrt{N}$ limit.

Another point worth highlighting in Table 1 is that higher factor volatility tends to result in lower Diversification Coefficient. This is broadly consistent with the findings in Example 2. Note that there is no reason why the Diversification Coefficient would *always* be lower for high-volatility factors, since in a multifactor setting the Diversification Coefficient also depends on the correlations between the factors. Finally, we note from Table 1 that larger universes tend to have larger Diversification Coefficient. This result is intuitive, since there is more opportunity to diversify residual risk in a larger universe. The size factor, however, represents an apparent anomaly: The Diversification Coefficient for size actually *decreases* as the investment universe is expanded.

In Table 2, we explore this apparent anomaly for the size factor, and compare it with a more typical factor, such as value. For each universe, we build unconstrained optimal portfolios for value and size. By construction, these portfolios have

Table 1 Annualized factor volatility and Diversification Coefficient for style factors in the Axioma US equity risk model, for analysis date April 17, 2015. The Diversification Coefficient is reported for three different universes: (a) S&P 500, (b) Russell 1000, and (c) Russell 3000.

		Diversification Coefficient		
Factor	Volatility	S&P 500	RU 1000	RU 3000
Dividend yield	0.77%	21.39	23.58	28.51
FX sensitivity	1.02%	23.66	29.20	37.17
Growth	0.75%	26.72	32.42	41.57
Leverage	1.15%	20.96	23.13	30.30
Liquidity	1.87%	19.61	22.99	24.91
Market sensitivity	3.47%	14.82	17.79	24.55
Momentum	2.71%	13.97	16.33	18.79
ROE	0.94%	23.92	29.84	36.80
Size	5.24%	15.98	12.99	9.47
Value	1.41%	18.53	21.98	27.33
Volatility	3.36%	17.65	21.68	22.74
SQRT(N)	N/A	22.36	31.62	54.77

Table 2 Analysis of value and size factors for analysis date April 17, 2015.

Factor	Stock index	Expected return	Portfolio volatility	Information Ratio	Signal Quality	Diversif. Coeff.
Value	S&P 500	1.00	1.84	0.54	0.029	18.53
Value	RU 1000	1.00	1.38	0.72	0.033	21.98
Value	RU 3000	1.00	1.18	0.85	0.031	27.33
Size	S&P 500	1.00	4.15	0.24	0.015	15.98
Size	RU 1000	1.00	3.62	0.28	0.021	12.99
Size	RU 3000	1.00	3.31	0.30	0.032	9.47

unit exposure to the factor, and therefore identical expected returns, regardless of the universe. Note, however, that in all cases the volatility decreases as the investment universe is expanded. This must always be the case, since the *IR* necessarily increases as we expand the size of the universe. This result is most easily seen by recognizing that the optimal portfolio for the S&P 500 universe can be exactly replicated using the Russell 3000 universe, under the constraint that the optimal portfolio holds zero weight in any stock outside the S&P 500. Such constraints always serve to reduce the *ex ante IR.*

In Table 2, we see that the Signal Quality for the value factor is nearly identical across the different investment universes. This reflects the fact that the distribution of value exposures (i.e., the "alphas") is very similar across the three different universes. The increasing *IR*, coupled with a nearly constant Signal Quality, leads to an increase in the Diversification Coefficient for the value factor as the investment universe expands.

The situation for the size factor, however, is quite different. In this case, we see that the Signal Quality increases dramatically as we move from the S&P 500 to the Russell 3000. This is an artifact of the distribution of the size exposure. For the S&P 500, the size exposures are clustered on the right-hand side of the distribution. The Russell 3000, by contrast, includes the left tail of the size distribution, which contains many small-cap stocks with large negative size exposures. Hence, although the *IR* increases as we expand the size of the investment universe, the Signal Quality increases at an even faster rate, leading to a reduction in the Diversification Coefficient.

3 The transfer coefficient

Up to now, we have assumed that the optimal portfolio Y is constructed free of constraints. In the real world, however, portfolio managers are faced with investment constraints such as turnover limits and/or long-only restrictions. These constraints cause the portfolio manager to hold a different portfolio P, which necessarily has a lower *ex ante IR*. The Transfer Coefficient quantifies the drop in portfolio efficiency due to these constraints, i.e., $IR_P = TC \cdot IR_Y$.

In Appendix C, we derive an exact expression for the Transfer Coefficient using a risk-budget approach. We show that the TC is simply the predicted correlation between portfolio P and portfolio Y,

$$TC = \text{corr}(R_P, R_Y). \qquad (27)$$

In practice, the correlation is computed using the same asset covariance matrix $\mathbf{\Omega}$ that was used to construct the optimal portfolio.

Clarke *et al.* (2006) derived a different expression for the Transfer Coefficient,

$$TC = \frac{\boldsymbol{\alpha}'\mathbf{h}_P}{\sqrt{\boldsymbol{\alpha}'\boldsymbol{\Omega}^{-1}\boldsymbol{\alpha}}\sqrt{\mathbf{h}_P'\boldsymbol{\Omega}\mathbf{h}_P}}, \quad (28)$$

where $\mathbf{h}_P$ is the portfolio holdings vector. Although this expression looks very different from Equation (27), in Appendix C we show that these two expressions are in fact equivalent. Finally, putting everything together, we arrive at the following exact expression for the Information Ratio of portfolio P,

$$IR_P = TC \cdot Q \cdot D. \quad (29)$$

Equation (29) represents the fundamental law in its exact form and full generality.

4 Summary

We have introduced a new formulation of the fundamental law of active management. By recognizing that manager skill is independent of asset correlations, we derived an exact expression for skill, which we denote the Signal Quality. We showed that the Signal Quality is given by the root-mean-square (RMS) Information Ratio of the underlying set of assets. We also derived an exact closed-form expression for the Diversification Coefficient, which is analogous to the square root of Breadth in the Grinold–Kahn formulation. In addition, we showed that the Transfer Coefficient is exactly given by the return correlation between the actual portfolio P and the unconstrained optimal portfolio Y. We included several examples to illustrate how the fundamental law of active management is applied in practice.

Appendix A: Optimal two-asset portfolio

To solve for the optimal portfolio of two assets A and B, five parameters are required: the expected returns α_A, and α_B of the assets, the volatilities σ_A and σ_B of the assets, and the correlation ρ between the assets. The Information Ratio of Asset A is therefore $IR_A = \alpha_A/\sigma_A$, with a corresponding expression for Asset B.

The Signal Quality Q is the root-mean-square *IR* of the two assets,

$$Q = \sqrt{\frac{IR_A^2 + IR_B^2}{2}}. \quad (A.1)$$

The asset covariance matrix is given by

$$\Omega = \begin{bmatrix} \sigma_A^2 & \rho\sigma_A\sigma_B \\ \rho\sigma_A\sigma_B & \sigma_B^2 \end{bmatrix}. \quad (A.2)$$

The optimal portfolio weights are proportional to $\mathbf{h} = \boldsymbol{\Omega}^{-1}\boldsymbol{\alpha}$, where $\mathbf{h}$ is a 2×1 vector of asset weights, and $\boldsymbol{\alpha}$ is the 2×1 vector of asset expected returns. Using the standard solution for the inverse of a 2×2 matrix, we obtain the optimal weights h_A and h_B,

$$\begin{bmatrix} h_A \\ h_B \end{bmatrix} = \frac{1}{1-\rho^2}\begin{bmatrix} (\alpha_A/\sigma_A^2) - (\rho\alpha_B/\sigma_A\sigma_B) \\ (\alpha_B/\sigma_B^2) - (\rho\alpha_A/\sigma_A\sigma_B) \end{bmatrix}. \quad (A.3)$$

Note that these weights do not automatically satisfy the full-investment constraint. In most instances, the weights can be simply rescaled so that the resulting portfolio weights sum to unity. Sometimes, however, the optimal portfolio is net short, in which case the optimal weights can never sum to a positive number.

The expected return of optimal portfolio Y is given by

$$E[R_Y] = \frac{IR_A^2}{1-\rho^2}(1 + q^2 - 2\rho q), \quad (A.4)$$

where q is defined as $q \equiv IR_B/IR_A$. The risk of portfolio Y is given by

$$\sigma_Y = \frac{IR_A}{\sqrt{1-\rho^2}}(1 + q^2 - 2\rho q)^{1/2}, \quad (A.5)$$

which leads to a portfolio Information Ratio of

$$IR_Y = \frac{IR_A}{\sqrt{1-\rho^2}}(1+q^2-2\rho q)^{1/2}. \quad \text{(A.6)}$$

Finally, the Diversification Coefficient D is given by

$$D = \sqrt{2}\left[\frac{(1+q^2-2\rho q)}{(1-\rho^2)(1+q^2)}\right]^{1/2}. \quad \text{(A.7)}$$

Note that when $\rho = 0$, the Diversification Coefficient becomes $D = \sqrt{2}$, as required.

Appendix B: One-factor model

We consider a simple model to illustrate the interplay between specific risk (diversifiable) and factor risk (non-diversifiable) in portfolio optimization. This model was first used by Lee and Stefek (2008) to study the effect of factor misalignment in portfolio construction. It assumes that stock returns are driven by a single risk factor with volatility σ_F, while further assuming that all stocks have the same specific risk, denoted σ_S. We consider a universe of N stocks, with covariance matrix $\mathbf{\Omega}$ given by

$$\mathbf{\Omega} = \mathbf{XFX}' + \Delta, \quad \text{(B.1)}$$

where $\mathbf{X}$ is the factor exposure matrix, $\mathbf{F}$ is the factor covariance matrix, and Δ is the $N \times N$ diagonal matrix of specific variances. Since our model includes only a single factor, $\mathbf{X}$ is an $N \times 1$ column vector, and $\mathbf{F}$ is given by the factor variance σ_F^2.

In order to find the optimal portfolio, we require the inverse of the asset covariance matrix. The solution is found in Grinold and Kahn (2000),

$$\mathbf{\Omega}^{-1} = \Delta^{-1} - \Delta^{-1}\mathbf{X}(\mathbf{X}'\Delta^{-1}\mathbf{X} + \mathbf{F}^{-1})^{-1} \times \mathbf{X}'\Delta^{-1}. \quad \text{(B.2)}$$

The inverse of the diagonal specific variance matrix is given by

$$\Delta^{-1} = \frac{\mathbf{I}}{\sigma_S^2}, \quad \text{(B.3)}$$

where $\mathbf{I}$ is the $N \times N$ identity matrix. We assume that the risk factor $\mathbf{X}$ has mean zero and unit standard deviation, which implies

$$\mathbf{X}'\mathbf{X} = N. \quad \text{(B.4)}$$

Substituting these relations into Equation (B2), we obtain

$$\mathbf{\Omega}^{-1} = \frac{\mathbf{I}}{\sigma_S^2} - \frac{1}{\sigma_S^2}\left(\frac{\sigma_F^2}{N\sigma_F^2 + \sigma_S^2}\right)\mathbf{XX}'. \quad \text{(B.5)}$$

We assume that our alpha signal is perfectly aligned with the risk factor, i.e.,

$$\boldsymbol{\alpha} = a\mathbf{X}, \quad \text{(B.6)}$$

where a represents the standard deviation of the alpha signal.

The holdings vector $\mathbf{h}_Y$ of the unconstrained optimal portfolio Y, up to an arbitrary multiplicative constant, is given by

$$\mathbf{h}_Y = \mathbf{\Omega}^{-1}\boldsymbol{\alpha}. \quad \text{(B.7)}$$

Substituting Equation (B5) and Equation (B6) into Equation (B7), and utilizing Equation (B4), we obtain

$$\mathbf{h}_Y = \frac{a\mathbf{X}}{N\sigma_F^2 + \sigma_S^2}. \quad \text{(B.8)}$$

The expected return of portfolio Y is computed as $\mathbf{h}_Y'\boldsymbol{\alpha}$, which gives

$$E[R_Y] = \frac{a^2 N}{N\sigma_F^2 + \sigma_S^2}. \quad \text{(B.9)}$$

The variance of portfolio Y is given by

$$\sigma_Y^2 = \mathbf{h}_Y'\mathbf{\Omega}\mathbf{h}_Y, \quad \text{(B.10)}$$

which reduces after a few lines of algebra to

$$\sigma_Y^2 = \frac{a^2 N}{N\sigma_F^2 + \sigma_S^2}. \quad \text{(B.11)}$$

The portfolio Information Ratio is given by

$$IR_Y = \frac{E[R_Y]}{\sigma_Y}, \quad \text{(B.12)}$$