*"Forthcoming in the Journal of Performance Measurement"*

# Measuring Investment Skill Using The Effective Information Coefficient

*This paper proposes a new portfolio performance metric called the Effective Information Coefficient. The EIC offers a key advantage relative to traditional performance metrics such as the information ratio: rather than get one observation of portfolio return for each period of time to evaluate, we get a one observation per time period for each asset in the investor's permissible universe of investments. As such, sample sizes are much larger and we are able to obtain statistically significant performance evaluations much more quickly.*

## Introduction

This paper proposes a new portfolio performance metric called the *Effective Information Coefficient*. The return performance of an investment manager is often summarized by the *information ratio,* the coefficient of variation of the active returns.

Similarly, the predictive power of asset return forecasts made by investment managers is summarized in a correlation statistic called the *information coefficient*. When portfolio managers form portfolios of actual securities, some of this predictive power can be lost due to risk limitations, portfolio constraints and trading costs. The extent to which the predictive power of the manager is translated into portfolio performance is usually summarized in a statistic called the *transfer coefficient*. We also know that most investors seek to have portfolios which are consistent with Markowitz's concept of mean-variance efficiency, wherein portfolios are formed so that *expected returns and risks are in balance at the margin*. We will combine these three concepts into a new measure called the "Effective Information Coefficient" as a way of evaluating investment skill. The EIC offers a key advantage relative to traditional performance metrics: rather than get one observation of portfolio return for each period of time to evaluate, we get a one observation per time period for each asset in the investor's permissible universe of investments. As such, sample sizes are much larger and we are able to obtain statistically significant performance evaluations much more quickly.

## Previous Research on Skill

There are numerous performance metrics used as proxies for investment manager skill such as realized alpha, and information ratio. In practice, we rarely obtain statistically significant values for these measures because you need a long time series of active return data over which conditions are stable. Unfortunately, real-world conditions rarely are stable, making this form of evaluation problematic. It would be helpful to have a measure that uses more information so we can get statistically meaningful results over a shorter time window.

Another important aspect to consider is that active managers occasionally experience very bad return outcomes for a period of time. It would be valuable to investors to be able to discriminate a meaningful decline in a manager's skill level from large, but random, negative outcomes.

There is an enormous literature in finance regarding whether investment managers collectively exhibit skill. The answer to that question has important implications for the issue of market efficiency, and the theory of asset pricing. Most of this research is based on the concept of "performance persistence". It assumes that those managers who perform consistently well must be skillful. Examples of this research include Brown and Goetzmann (1995), Elton, Gruber and Blake (1996), and Stewart (1998). There is also an extensive related literature such as Brown, Goetzmann, Ibbotson and Ross (1992), Carpenter and Lynch (1999), and Carhart, Lynch and Musto (2002) that debates whether such persistence effects are artifacts of survivorship bias in the data used for empirical studies

The issue as to whether or not managers collectively exhibit skill is of limited consequence in this paper. The task before us is the evaluation of single managers. For this purpose, there is a great deal of literature that centers on using traditional return based performance statistics as proxies for manager skill. The seminal paper is Kritzman (1986), introduced specific statistical analyses of past returns as a metric of investment manager skill. Other interesting papers include Marcus (1990) which incorporates the issue of selection bias, and Lee and Rahman (1991) which tries to distinguish between security selection and market timing skills among mutual fund managers. Bailey (1996) introduces a graphical approach to skill detection.

As previously noted, the limiting conditions on use of time series performance statistics as measures of manager skill are substantial. We must always have a sufficiently large sample of return observations while also meeting the statistical criteria for stationarity (stability of conditions). To the extent that the real world is constantly evolving, there is a natural tension between these two needs that makes it generally impossible to obtain statistically significant results on the performance records of individual managers when using typical return observation frequencies (e.g. monthly). One simplistic fix to this problem is to use high frequency observations such as daily returns, but using daily returns for skill evaluation is problematic on numerous fronts. The conceptual and statistical difficulties are detailed in diBartolomeo (2003) and diBartolomeo (2007).

Some researchers have tried to detect manager skill, or changes in the level thereof, using statistical process control methods. Philips, Stein and Yaschin (2003) use CUSUM methods to directly evaluate active manager performance. Bolster, diBartolomeo and Warrick (2006) use CUSUM as a method for detecting regime change so as to isolate the most relevant portion of a manager's track record for evaluation.

**The Breakdown Problem**

Let us consider an actual example of an institutional equity manager. Using a commercially available risk assessment system this manager managed his portfolio so as to keep the ex-ante risk forecast of tracking error (standard deviation of benchmark-relative return) below 3% per year. During a particular year, the manager's fund underperformed its benchmark index by 6.3%. Upon experiencing this event, the manager considered two possible rationales. The first is that he had been very unlucky and had randomly experienced a more than two standard deviation negative event. The other possible rationale was the risk assessment model was at fault, and was grossly underestimating the active risk of the portfolio.

However, when monthly returns were examined, a rather different picture emerged. The average value of the month-end ex-ante risk expectation was 2.74%, while the realized standard deviation of the twelve monthly returns during the year in question was 2.80% annualized. The risk model was almost exactly on target. Active performance was as consistent as it was expected to be. Unfortunately for our manager, it was consistently bad, with a mean monthly return of negative .54% per month during the sample year. What the manager had neglected is that the standard deviation of anything is a measure of dispersion around the mean, not around zero. For active returns, the dispersion around the mean and the dispersion around zero should be expected to be equivalent only for index funds. The common confusion around active return volatility and its implications for skill assessment are described in Huber (2001).

**The Information Ratio as a Proxy for Skill**

The most commonly used proxy for investment manager skill is the *information ratio*. In Grinold (1989), it is defined as the coefficient of variation of the manager's active returns.

IR = alpha / tracking error

The paper goes on to derive information ratio as the product of the *information coefficient* and the *breadth* of an active management strategy. Grinold refers to this relationship as the "Fundamental Law of Active Management".

$IR = IC * Breadth^{.5}$

Where

IC = correlation of your return forecasts and outcomes

Breadth = number of independent "bets" taken per unit time

If we know how good we are at forecasting returns (prediction skill) and how many bets we act on, we can forecast how good our active performance should be for any given risk level. However, the Fundamental Law makes big assumptions**.** One assumption is there are no constraints at all on portfolio construction, so positions can be long or short and of any size. A second is that transaction costs are zero, so bets in one time period are independent of bets in other periods. A third implicit assumption is that research resources are limitless so our forecasting effectiveness (IC) is constant as we increase the number of investment bets we chose to investigate.

Most crucially, the Fundamental Law requires that we measure only independent bets in our estimation of breadth. For example, if we choose to invest in twenty different stocks for twenty different reasons we can consider this set of actions as twenty different bets. However, if we choose to invest in twenty different stocks because they share a common trait we find preferable (e.g. a generous dividend, or a low P/E ratio), this is not twenty bets, just one very big bet! Once we've tilted the odds in our favor through positive return forecasting capability, we want to take lots of bets to maximize the information ratio. Unfortunately for investors, managers are rarely willing to disclose sufficient details of their investment process to make accurate estimation of breadth possible from the "outside in".

**Enter the Transfer Coefficient**

Many practitioners are uncomfortable with the use of information ratios as a measure of skill because the assumptions of no limitations on position sizes, zero trading costs and the availability of unlimited short positions are unrealistic for most investment portfolios. Clarke, de Silva and Thorley (2002) tries resolve this issue by introducing a scaling factor into the calculation of the information ratio that they call the *transfer coefficient*. We can think of the transfer coefficient as a scalar less than one which describes how much of the potential economic value added from our investment strategy actually contributes to actual performance. It points out the extent to which our potential value is lost due to the interference of constraints on position size and portfolio turnover.

$$IR = IC * TC * Breadth^{.5}$$

TC = the efficiency of your portfolio construction ($TC < 1$)

Imagine a manager with a diverse team of analysts that are great at forecasting monthly stock returns on a large universe of stocks, but whose portfolio is allowed to have only 1% per year turnover. The existence of good monthly forecasts, diverse reasons for actions (independence of bets) and a large universe imply high IC and high breadth. However, if we can never act on the forecasts because of the turnover constraint, the transfer coefficient can be zero or even negative. If we can't short a stock that we correctly believe is going down, or take a big position in a stock that we correctly believe is going up, the transfer coefficient will decline. The more binding constraints we have on our portfolio construction, the more return we fail to capture when our forecasts are good. For bad forecasters, a low transfer coefficient is good. You hurt yourself less when you constrain your level of activity. In some sense it is disingenuous for asset managers to simultaneously tout their forecasting skills, while simultaneously advocating layers of tight constraints on portfolio construction.

For situations where the information coefficient can be measured (i.e. a quantitative manager analyzing their own performance) another relationship emerges:

$$EIC = IC * TC$$

So for asset managers, measuring EIC and IC can provide an approach for the estimation of the transfer coefficient.

**Limitations of the Information Ratio**

While investment managers (especially hedge funds) often evidence their skills via realized information ratios, this measure really doesn't correspond to investor utility except in extreme cases. Consider a manager with an alpha of 1 basis point and a tracking error of zero. The information ratio is infinite but the economic value added for the investor is very, very small and inconsequential. A substantial investigation of this issue appears in deGroot, and Plantinga (2001).

Another problem with using the information ratio as a proxy for manager skill is that the statistical significance of differences across managers is difficult to calculate. For example if Manager A has an information ratio of .5 for the past sixty months, and manager B has an information ratio of .6 for the past sixty months, can we actually say those two values are materially different, and hence Manager B performed better than Manager A? Although algebraically complex, a method for this calculation is available by a slight modification of methods in Jobson and Korkie (1981).

Another limitation of the Fundamental Law is that it assumes that information coefficients (IC) are constant over time. This implies that the predictive skill level of a manager is a constant. Most practitioners assess the information coefficient through a series of cross-sectional analyses. To the extent that each cross-section represents a particular time period, information coefficients can vary. Qian and Hua (2004) define strategy risk as the standard deviation of the manager's IC over time, which leads to corresponding variations in excess returns. They define "forecast true active risk" as a combination of both "risk model predicted tracking error" (random return variation due to things outside the manager's control) and the return variation arising from strategy risk.

$$\text{Forecast Active Risk} = \text{std(IC)} * \text{Breadth}^{1/2} * \text{Forecast Tracking Error}$$

**The Effective Information Coefficient (EIC)**

Successful active management involves forecasting what returns different assets will earn in the future (the information coefficient), and forming portfolios that will efficiently use the valid information contained in the forecast to generate returns (the transfer coefficient). Typically, an investment manager will have a large universe of assets from which to choose. This implies that we can judge the statistical significance of our information coefficient (one observation of our forecast quality per asset per period) far more quickly than we can our information ratio (one observation of portfolio performance per period).

Our proposal is to extend the concept of the information coefficient to include the quality of portfolio construction, normally characterized by the transfer coefficient. We call this new measure, the *effective information coefficient*. This measure retains the cross-sectional nature of the information coefficient so statistical significance can be judged quickly, while also capturing the impact of portfolio constraints and limitations.

The basis of the effective information coefficient is the concept of portfolio optimality as first put forward by Markowitz (1952). In mathematical terms, optimality means that the position sizes within our portfolios balance the marginal returns, risks and costs. The requirement of this "balance at the margin" comes from the Kuhn-Tucker conditions which describe how we can find the maximum or minimum of a smooth algebraic function. *Every portfolio manager must believe that the portfolio they hold is optimal for their investors. If they didn't they would hold a different portfolio*. If we describe investor goals as maximizing risk adjusted returns, we know that the marginal risks associated with every active position must be exactly offset by the expected active returns. We can infer the manager's expectations of returns from the marginal risks they choose to accept in their portfolios. For every portfolio, there exists a set of alpha (active return) expectations that would make the portfolio optimal. We call these the *implied alphas*. Sharpe (1974) provides the basics of estimating implied returns, while Fisher (1975) demonstrates the linkage between analyst forecasts and portfolio changes.

We will define the *effective information coefficient* as the cross-sectional correlation between the implied alphas from portfolio security positions at each moment in time, and the residual returns realized by those individual securities in the subsequent period. We can also pool these values over time for a longer term estimate of the EIC.

$$EIC_t = \text{Correlation}\ (\text{Implied alphas}_{t-1}, \text{Realized alphas}_t)$$

The role that active weights play in the Clarke, et. al. (2002) procedure are impounded into our formulation of implied alphas. As such, we are able to avoid certain simplifying assumptions as described in in Suntharam, Khilnani and Demoiseau (2007).

To sum up the idea, we will use the *effective information coefficient* as the measure of investment manager skill. If our forecasting skill is good (high IC) and our portfolio construction skill is good (high TC) then effective information coefficient will be high. If either information coefficient or the transfer coefficient is low, then the effective information coefficient will be low. As this measurement involves every active position during each time period, the sample is large and statistical significance is obtained quickly. To the extent that the effective information coefficient is simply a form of correlation coefficient, the standard error calculation needed to calculate statistical significance is well known.

**Subtleties and Caveats for Use of the Effective Information Coefficient**

There are some subtleties and potential pitfalls in using the effective information coefficient. Most of these issues are analytical but potential users of the EIC technique may have operational concerns as well.

In order to estimate implied alphas, we must first estimate the marginal risks of portfolio positions. To the extent that different investment organizations hold different views of the marginal risks of positions they will obtain different estimates of implied alphas. In practice, however, there is a high degree of concordance among investment managers about portfolio risk. This is demonstrated by the fact that nearly every major asset manager uses a risk assessment model provided by one of just a few commercial vendors. Managers see their "value added" in superior return forecasting. As long as everyone roughly agrees on the covariance among securities, then we can reliably infer manager "alpha" forecasts from the portfolio they choose to hold.

In addition, studies such as Best and Grauer (1991), Broadie (1993) and Chopra and Ziemba (1993) show that estimation errors in risk have a relatively small impact on portfolio optimality as compared to errors in return estimation. A related instance of implying returns from covariance estimates (that are *assumed* to be accurate) can be found in the well-known Black-Litterman model (1991) for asset allocation. While implied alphas can be biased through estimation errors in the risk model, such usage imposes no greater risk than conventional management that is using the same risk model (i.e. you are no worse off than just about everybody else).

Estimating implied alphas directly also requires us to know the manager's level of aggressiveness (risk tolerance). If we don't know this, we can't estimate the magnitude of implied alphas but we can still estimate the implied rank value of the implied alphas from the marginal risks. Our first alternative is to estimate the effective information coefficient as a rank correlation measure such as the Spearman Rho or Kendall's Tau. This may mask the influence of transaction costs in defining optimality if trading costs are heterogeneous across securities. A second approach would be to "map" the implied alpha rank values into an estimated cross-sectional distribution for returns. de Silva, Sapra and Thorley (2001) and Lilo, Mantegna, Bouchard and Potters (2002) provide methods for estimating the cross-sectional distribution of returns. Finally, we can try to infer the manager's risk tolerance from the observed level of portfolio risk itself. Wilcox (2003) argues that rational investors maximize the long term growth of their discretionary wealth (the portion of wealth they can afford to lose). If we are willing to define an investor's "worst case scenario" as a particular probability of catastrophic loss (e.g. a three standard deviation event), then we can directly estimate risk tolerance from the magnitude of portfolio risk undertaken.

Another concern about the use of implied alphas is how they can be biased by constraints on portfolio position size. Most obviously, most portfolio managers are prohibited from taking short positions. This issue is particularly acute because we are implying benchmark relative returns rather than absolute returns. Without the ability to short positions, the distribution of implied alphas will lack the large magnitude negative values that would be implied by short positions. As such, the distribution may exhibit positive skew. Similar truncation of the upper tail of the implied alphas distribution can occur from a maximum weight bounds on position sizes in portfolios. To determine if this problem is material to a given portfolio we can check the distribution of implied portfolio returns to see if it has the expected properties. The distribution of implied returns should be roughly symmetric about the mean, skew should be close to zero and the expected alpha on the benchmark index portfolio should be zero. If the observed properties of the implied alphas distribution are not satisfactory, we can adjust the implied alphas on *only those securities whose portfolio position is constrained by a weight bound*. A simple adjustment rule consistent with Grinold (1994) is:

Adjusted Implied Alpha (i) = Implied Alpha (i) + (x * Specific Risk (i))

The logic of this process is that the potential for security (i) to underperform or outperform the benchmark index is proportional to the security's specific risk. For those securities whose implied alpha is constrained by a portfolio weight bound (e.g. long only), we make an additive adjustment to the implied alpha by selecting a single value x for all bounded securities in the portfolio. The value of x is chosen to minimize the extent to which the distribution of implied alphas is different from expectations.

From an operational perspective, the entity doing the analysis must have access to the portfolio positions on a periodic basis, have at least rough estimates of trading costs for different securities in the portfolio, and have a detailed analytical model of how each security position contributes to the risk of the portfolio. The routine process of monthly statements from a custodian or portfolio accounting system fulfills the first need. As previously noted, commercially available analytical models of risk are widely used by asset managers, consultants and custody banks in their reporting of risk levels. All that is required for the EIC analysis is that the reports include "marginal contributions to tracking variance" which are a standard output of the widely used systems. The EIC analysis is relatively insensitive to trading costs, except for very illiquid securities so it is of lesser importance in most cases. .In addition, as previously mentioned, we can also modify the analytical procedure to reduce the need for trading cost information.

**Using EIC to Test Risk Model Effectiveness**

For active managers to generate excess returns in a given time period, there must be cross-sectional dispersion in the individual asset returns. If all assets had the same return during a particular period, no active returns would be available to any portfolio, as every portfolio and benchmark would also have the same return. Even if the magnitude of the common return was different in different periods, the realized active risk would also be zero since every portfolio and every benchmark would have the same return in each period. As such, a manager's expected active return is a function of their EIC (are they skillful?), their risk tolerance (are they willing to take bets?) and the opportunity set afforded them as measured by the cross-sectional dispersion of asset returns. The empirical relationship between cross-sectional dispersion of asset returns and manager active returns has been confirmed in Ankrim and Ding (2002).

So we can look at returns as:

$P_t - B_t = \text{Expected Alpha}_t + e_t$

$P_t$ = portfolio return during period t

$B_t$ = benchmark return during period t
$e_t$ = residual returns due to luck

If risk model is predicting accurately, the annualized value of the time series standard deviation of the $e_t$, should be consistent with the risk model forecast tracking error.

**Conclusions**
Most traditional measures of investment performance, such as information ratios, have weak statistical power because they require long time series of stationary conditions to come to statistically significant conclusions.   Our new measure, the *effective information coefficient*, is able to take advantage of a far large sample of data, allowing for rapid statistical significance and also incorporates important information about a manager's portfolio construction efficiency as well as proficiency in forecasting asset returns.   The EIC offers the additional benefit that in can be used without knowledge of the manager's security level return expectations, making it a practical investigative tool for investors who employ external asset managers.

**References**

Stewart, Scott D. "Is Consistency Of Performance A Good Measure Of Manager Skill?," Journal of Portfolio Management, 1998, v24(3,Spring), 22-32.

Brown, Stephen J. and William N. Goetzmann. "Performance Persistence," Journal of Finance, 1995, v50(2), 679-698.

Elton, Edwin J., Martin J. Gruber and Christopher R. Blake. "The Persistence Of Risk-Adjusted Mutual Fund Performance," Journal of Business, 1996, v69(2,Apr), 133-157.

Brown, Stephen J., William Goetzmann, Roger G. Ibbotson and Stephen A. Ross. "Survivorship Bias In Performance Studies," Review of Financial Studies, 1992, v5(4), 553-580.

Carpenter, Jennifer N. and Anthony W. Lynch. "Survivorship Bias And Attrition Effects In Measures Of Performance Persistence," Journal of Financial Economics, 1999, v54(3,Dec), 337-374.

Carhart, M. M., J. N. Carpenter, A. W. Lynch and D. K. Musto. "Mutual Fund Survivorship," Review of Financial Studies, 2002, v15(5), 1439-1463.

Kritzman, Mark. "How To Detect Skill In Management Performance," Journal of Portfolio Management, 1986, v12(2), 16-20.

Lee, Cheng F. and Shafiqur Rahman. "New Evidence On Timing And Security Selection Skill Of Mutual Fund Managers," Journal of Portfolio Management, 1991, v17(2), 80-83.

Marcus, Alan J. "The Magellan Fund And Market Efficiency," Journal of Portfolio Management, 1990, v17(1),85-88.

diBartolomeo, Dan. “Just Because We Can Doesn’t Mean We Should: Why Daily Attribution is Not Better”, Journal of Performance Measurement, Spring 2003

diBartolomeo, Dan. “Fat Tails, Tall Tales, Puppy Dog Tails”, Professional Investor, Autumn 2007.

Bailey, Jeffery V. "Evaluating Investment Skill With A VAM Graph," Journal of Investing, 1996, v5(2,Summer), 64-71.

Philips, Thomas K., David Stein and Emmanuel Yashchin, “Using Statistical Process Control to Monitor Active Managers”, *Journal of Portfolio Management*, 2003

Bolster, Paul, Dan diBartolomeo and Sandy Warrick, “Forecasting Relative Performance of Active Managers”, Northfield Working Paper, 2006

Huber, Gerard. “Tracking Error and Active Management”, Northfield Conference Proceedings, 2001, http://www.northinfo.com/documents/164.pdf

Grinold, Richard C. "The Fundamental Law Of Active Management," Journal of Portfolio Management, 1989, v15(3), 30-37.

Clarke, Roger, Harindra de Silva and Steven Thorley. "Portfolio Constraints And The Fundamental Law Of Active Management," Financial Analyst Journal, 2002, v58(5,Sep/Oct), 48-66.

deGroot, Sebastien and Auke Plantinga, “Risk Adjusted Performance Measures and Implied Risk Attitudes”, *Journal of Performance Measurement*, Winter 2001/2002

Jobson, J. D. and Bob M. Korkie. "Performance Hypothesis Testing With The Sharpe And Treynor Measures," Journal of Finance, 1981, v36(4), 889-908.

Qian, Edward and Ronald Hua. “Active Risk and the Information Ratio”, Journal of Investment Management, Third Quarter 2004.

Markowitz, Harry. "Portfolio Selection," Journal of Finance, 1952, v7(1), 77-91.

Sharpe, William F. "Imputing Expected Security Returns From Portfolio Composition," Journal of Financial and Quantitative Analysis, 1974, v9(3), 462-472.

Fisher, Lawrence. "Analysts' Input And Portfolio Changes," Financial Analyst Journal, 1975, v31(3), 73-85.

Best, Michael J. and Robert R. Grauer. "On The Sensitivity Of Mean-Variance-Efficient Portfolios To Changes In Asset Means: Some Analytical And Computational Results," Review of Financial Studies, 1991, v4(2), 315-342.

Chopra, Vijay K. and William T. Ziemba. "The Effect Of Errors In Means, Variances, And Covariances On Optimal Portfolio Choice," Journal of Portfolio Management, 1993, v19(2), 6-12.

Broadie, Mark. “Computing Efficient Frontiers with Estimated Parameters”, Annals of Operations Research, 1993.

Black, Fischer and Robert Litterman. "Asset Allocation: Combining Investor Views With Market Equilibrium," Journal of Fixed Income, 1991, v1(2), 7-18.

de Silva, Harindra, Steven Sapra and Steven Thorley. "Return Dispersion And Active Management," Financial Analyst Journal, 2001, v57(5,Sep/Oct), 29-42.

Lilo, Fabrizio, Rosario Mantegna, Jean-Philippe Bouchard and Marc Potters. “Introducing Variety in Risk Management”, WILMOTT December 2002.

Wilcox, Jarrod. "Harry Markowitz and the Discretionary Wealth Hypothesis", Journal of Portfolio Management, Spring 2003.

Grinold, Richard C. "Alpha Is Volatility Times IC Times Score,"Journal of Portfolio Management, 1994, v20(4), 9-16.

Suntharam, Ganesh, Vasant Khilnani and Eric Demoiseau. "Measuring and Targeting Efficiency to Optimise the Use of Turnover", Perpetual Working Paper, 2007, http://www.northinfo.com/documents/292.pdf

Ankrim, Ernest M. and Zhuanxin Ding. "Cross-Sectional Volatility And Return Dispersion," Financial Analyst Journal, 2002, v58(5,Sep/Oct), 67-73.

De Silva, Harindra, Steven Sapra and Steven Thorley. "Return Dispersion And Active Management," Financial Analyst Journal, 2001, v57(5,Sep/Oct), 29-42.