Alpha is Volatility Times IC Times Score

Real alphas don't get eaten.

Richard C. Grinold

RICHARD C. GRINOLD is a managing director of the advanced strategies and research group at Wells Fargo Nikko Investment Advisors in San Francisco (CA 94105).

Alpha is the key to investment success. All the information that distinguishes an investor's judgment from the consensus is distilled into alphas, and then translated into portfolios in the hope of producing superior performance and contented clients.

The explicit use of alpha is one of the many ways to distinguish between a traditional portfolio manager and a systematic portfolio manager.[1] Dealing with alphas is not easy, however, because the alpha is often the endpoint of an informed but ad hoc analysis. My experience indicates that many systematic managers have difficulty translating their quantitative and qualitative insights into expected returns.

In this article we hope to provide assistance and insight for the systematic investor. We will do this by stripping alpha into its component parts. Those parts are: 1) the *volatility* of the return that we are forecasting, 2) an *information coefficient* (IC) that measures the manager's overall skill at forecasting, and 3) a *score* that is a standardized measure of how strongly we feel about a particular stock at a particular time.

This result should be looked on as 90% insight and 10% science. The rule in the title can be derived in a rigorous way, but that rigor gives an aura of false precision that is not useful when one turns to implementation.

We will explain the rule in the next section and provide several examples. In particular, we consider the phenomenon known as "alpha eating." Alpha eating occurs when the portfolios emerging from an optimization appear to be ignoring the alphas. We argue that alpha eating occurs because the alphas are not consistent with the prescription: Alpha is Volatility Times IC Times Score.

A technical appendix is available from the author that derives the main result and provides technical backup for several of the examples.

THE BUILDING BLOCKS OF ALPHA

A technical definition of alpha involves the notion of residual return and the choice of either a forward-looking or backward-looking perspective. The term alpha means one thing if we look back to the world of statistics, realizations, and performance, and quite another if we look forward to the world of probabilities, expectations, and hope.

When an investor's performance is managed relative to a benchmark like the S&P 500 or the Frank Russell 3000, it is convenient to split the excess return of each stock (and portfolio) into a portion that is perfectly correlated with the benchmark's return and a second part that is uncorrelated with the benchmark's return.[2] The uncorrelated part is called the residual.[3]

Looking Back

If we are looking at past returns and evaluating performance, the average realized residual return is frequently called alpha,[4] although historical alpha or realized alpha would be more precise.

Looking Forward

If we are looking to the future, alpha is the *expected* residual return. It is this forward-looking version of alpha that will concern us.

The great benefit of defining the alpha in terms of expected residual return is that it provides a convenient starting point for active management. The baseline value for each stock's alpha is zero. If all the stocks have zero alphas, then the manager is inclined to hold the benchmark portfolio. To the extent that the managers assign positive or negative alphas to certain stocks, the manager's portfolio will tend to differ from the benchmark.

The portfolio's divergence from the benchmark is determined by the alphas' divergence from zero.[5] Alpha is what differentiates one portfolio from the herd; it is thus the driving force in portfolio management.

Understanding alpha is vital for the systematic manager. We hope to promote that understanding by stripping alpha down to its essential components: the residual volatility of the stock, the IC, and a score. We consider these in turn.

Volatility

The residual volatility is easily calculated from a model of asset and portfolio risk. Exhibit 1 shows the total volatility and residual volatility for the stocks in the MMI. The numbers are BARRA forecasts of annual volatility as of the end of July 1993. The stocks are ordered by increasing residual volatility.

The residual volatility for these stocks averages about 70% of the total volatility. Stocks with higher total volatility tend (with the major exception of the oil companies, Exxon and Chevron) to have higher residual volatility.

IC

The IC or information coefficient is a measure of forecasting skill. The IC is the correlation between the manager's forecast of residual return and the residual return. A bit of humility is useful in selecting an IC. A reasonable IC for an outstanding (top 5%) manager forecasting the returns on 500 stocks is about 0.06. If

EXHIBIT 1
Total and Residual Volatility

MMI Company	Total Volatility	Residual Volatility
MMM	23.13%	13.41%
GE	25.77%	14.42%
ATT	24.40%	15.89%
P&G	25.32%	16.29%
Dow Chem	25.98%	16.93%
DuPont	25.15%	17.29%
Coca-Cola	26.89%	18.92%
J&J	29.36%	18.97%
Disney	29.77%	19.17%
Kodak	26.98%	19.20%
Intl Paper	22.68%	19.83%
Philip M	31.64%	20.17%
Merck	31.52%	20.43%
Chevron	24.36%	20.44%
McDonald's	28.92%	20.54%
Exxon	24.95%	21.13%
Sears	33.90%	22.33%
Amex	33.31%	23.26%
GM	33.16%	23.46%
IBM	38.63%	30.32%

EXHIBIT 2
Simple Scoring Rule

Category	Score	Number in Category
Very Positive	2	32
Positive	1	128
No View	0	180
Negative	−1	128
Very Negative	−2	32

the manager is good (top quartile), 0.04 is a reasonable number.[6]

Score

The score is a *standardized measure* that shows how strongly you feel about a particular stock at a particular time. As a standardized measure, the average score is anticipated to be zero and the standard deviation of the score to be one. This means that the average and standard deviation of the scores for a particular stock over many periods of time should be close to zero and one, respectively.

In the same fashion, it means that the average and standard deviations of the scores over many stocks at one particular time should be close to zero and one as well. Scores are controlled for location, because the average is zero, and controlled for scale, because the standard deviation is one.

As an example, consider a simple scoring scheme that allocates stocks to five categories: very positive, positive, no view, negative, and very negative. We will associate numerical scores of 2, 1, 0, −1, and −2 to these five categories. We can even go farther than usual and ration the number of stocks that can be assigned to a category. If we are going to score 500 stocks, Exhibit 2 could be used to define the scores.

This simple assignment rule produces scores with an average of zero and a standard deviation of one. It is easy to imagine more sophisticated versions of the same idea.

Don't confuse the IC with the score. The score tells how strongly you feel about a particular stock at a particular time. The IC tells whether your feelings are in any way linked to the subsequent return. The scores change from period to period and from stock to stock. The IC is constant across stocks and across time.[7]

With these pieces introduced we can write out the formula:

$$\text{Alpha} = \text{Volatility} \times \text{IC} \times \text{Score} \qquad (1)$$

Exhibit 3 demonstrates an example for the MMI. We use an IC level of 0.09 and a random number generator to sample the scores from a standard normal distribution. We use these alphas again in the alpha-eating example (Number 5) below.

HOW TO BUILD ALPHAS

We apply the Volatility × IC × Score rule in several situations. We hope that these specific examples will indicate how one can use the same rule in other cases.

Example Number 1: A Tip

Consider that most ad hoc of all situations, the stock tip. Let's say the stock in question is Philip Morris, with 20.17% residual volatility. To change the subjective stock tip into a forecast of residual return, we need the IC and the score. For the IC you look to the track record of the source: If the source is great set IC = 0.1, if the source is good IC = 0.05, and if the source is a waste of time, then IC = 0. For the score we can give a run-of-the-mill tip (very positive) a 1.00 and a very, very positive tip a 2.00.

EXHIBIT 3
Alphas from Residual Volatility and Scores

MMI Company	Residual Volatility (%)	Score	Alpha (%)
MMM	13.41	0.3507	0.42
GE	14.42	0.7746	1.01
ATT	15.89	0.7094	1.01
P&G	16.29	−2.3165	−3.40
Dow Chem	16.93	−0.769	−1.17
DuPont	17.29	1.5844	2.47
Coca-Cola	18.92	−0.4825	−0.82
J&J	18.97	−1.7717	−3.02
Disney	19.17	0.3579	0.62
Kodak	19.20	−0.0587	−0.10
Intl Paper	19.83	−0.0278	−0.05
Philip M	20.17	−0.8938	−1.62
Merck	20.43	0.7404	1.36
Chevron	20.44	−0.2455	−0.45
McDonald's	20.54	−0.4458	−0.82
Exxon	21.13	−0.0043	−0.01
Sears	22.33	0.8477	1.70
Amex	23.26	0.3482	0.73
GM	23.46	1.9755	4.17
IBM	30.32	−0.6732	−1.84

EXHIBIT 4
Alphas from a Tip

IC	Very Positive Score = 1	Very Very Positive Score = 2
Great: IC = 0.10	2.02%	4.03%
Good: IC = 0.05	1.01%	2.02%
No Info: IC = 0.00	0.00%	0.00%

Exhibit 4 shows the spectrum of possibilities and the ability to transform some unstructured qualitative information into a more useful quantitative form.

Example Number 2: Buy and Sell Recommendations

A more structured example works with a buy and sell list. In this case we give a score of +1.00 to the buys and a score of −1.00 to the sells. If we apply this to the MMI stocks with a random choice of buy and sell and an IC of 0.09, we get the alphas shown in Exhibit 5.

The rule gives higher alphas to the more volatile stocks. If we ignore the rule and give an alpha of +1.00% to the buy stocks and an alpha of −1.00% to the sell stocks, we would be prime candidates for the alpha-eating phenomenon described in Example Number 5.

EXHIBIT 5
The Buy and Sell List

MMI Company	Residual Volatility (%)	View	Score	Alpha (%)
MMM	13.41	Sell	−1.00	−1.21
GE	14.42	Sell	−1.00	−1.30
ATT	15.89	Buy	1.00	1.43
P&G	16.29	Sell	−1.00	−1.47
Dow Chem	16.93	Buy	1.00	1.52
DuPont	17.29	Buy	1.00	1.56
Coca-Cola	18.92	Sell	−1.00	−1.70
J&J	18.97	Buy	1.00	1.71
Disney	19.17	Sell	−1.00	−1.73
Kodak	19.20	Buy	1.00	1.73
Intl Paper	19.83	Sell	−1.00	−1.78
Philip M	20.17	Buy	1.00	1.82
Merck	20.43	Sell	−1.00	−1.84
Chevron	20.44	Buy	1.00	1.84
McDonald's	20.54	Buy	1.00	1.85
Exxon	21.13	Sell	−1.00	−1.90
Sears	22.33	Sell	−1.00	−2.01
Amex	23.26	Sell	−1.00	−2.09
GM	23.46	Buy	1.00	2.11
IBM	30.32	Buy	1.00	2.73

Example Number 3: A Numerical Forecast for Each Stock

In this example we have a raw forecast that attributes a numerical value for each stock. This might be the output of a DDM, some valuation ratio such as consensus earnings forecast divided by price, or a price momentum indicator.

Raw forecasts present a problem. To produce standardized scores, we need to know the standard deviation of the raw forecast *for each stock*. We are not worried about the volatility of the stock's residual return; there are reasonable and robust ways to determine the residual return volatilities. We are concerned about the standard deviation of the raw forecasts and any linkage it may have with the volatility of the stock's residual return. In most cases we will not have enough information to determine the standard deviation of each stock's raw forecast. We must make an informed assumption.

At one extreme, we could say that the raw forecast has the same standard deviation for each stock. This is the "no problem" scenario. If we assume there is "no problem," we would turn the raw forecasts into scores by subtracting the cross-sectional average and dividing by the cross-sectional standard deviation.

At the other extreme, we could say that the standard deviation of each stock's raw forecast is proportional to the residual volatility of that stock's return. In the context of our MMI example, that would mean that the standard deviation of the raw MMM forecast is 44% (= 13.41/30.32) of the standard deviation of the raw forecast for IBM. Then we could turn the raw forecasts into scores in two steps: 1) first divide each stock's raw forecast by that stock's residual volatility, and then 2) standardize by subtracting the cross-sectional average and divide by the cross-sectional standard deviation.

What is the right assumption? Does it make any difference? It will matter if the stocks have a wide range of residual volatilities. They do. Even with the twenty MMI stocks, we see that the largest residual volatility (IBM) is more than twice the smallest residual volatility (MMM).

The investigation should be conceptual and include some statistical backup. We should first examine the process that generates the numbers. If the inputs to that process are (to some extent) surrogates for volatility (such as dispersion in analysts' forecasts, estimated change in analysts' forecasts, beta, etc.), then it is very likely that the outputs will be related to volatility.

As a supporting statistical test, you can rank the stocks by residual volatility and sort them into groups of some reasonable number (say, forty), and then compare the cross-sectional standard deviation of the raw forecasts within each group to the average residual return volatility of the stocks in that group. If this statistical information supports your conceptual investigation, then you have the insight needed to adjust the raw forecasts.

As an example, we considered raw forecasts for the 100 stocks in the OEX that were generated by a dividend discount model. We suspect that there is a strong relationship between these raw forecasts and residual volatility because the DDM has two inputs (estimated growth in earnings per share and required return) that are strongly linked to volatility. Exhibit 6 shows the stock's residual return volatility on the horizontal axis and the model's raw forecast on the vertical axis.

To check the relationship between the residual volatility and the raw forecasts we correlate the *absolute value* of the raw forecasts with volatility.[9] A small amount of analysis indicates that the correlation would be zero if the raw forecasts were independent of the stock's residual volatility and in the neighborhood of 0.35 if the volatility of the raw forecasts were proportional to the residual volatility. In this case the correlation is 0.39, which points strongly toward the conclusion that the raw forecasts are proportional to the residual volatility.

Thus, to transform these DDM outputs into alpha we follow a (nearly circular) route by: 1) dividing the raw forecast for each stock by that stock's residual standard deviation; 2) turning the resulting numbers into scores by subtracting off the cross-sectional average and dividing by the cross-sectional standard deviation, and 3) transforming into alphas by multiplying the scores by the IC and the residual volatility.[10]

The particular raw forecasts for the OEX stocks shown in Exhibit 6 have a cross-sectional average of zero and a cross-sectional standard deviation of 2.42%. The net effect of all the adjustments is to reduce that cross-sectional standard deviation to 1.25%.

EXHIBIT 6
FORECAST AND RESIDUAL VOLATILITY FOR THE OEX STOCKS

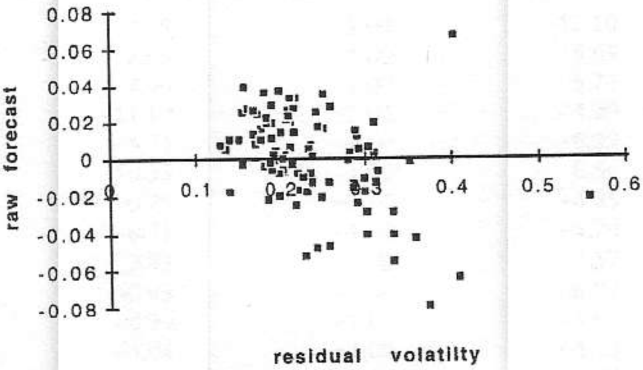

Example Number 4: Multiple Raw Forecasts

This example concerns two or more raw forecasts for several stocks. In that case, there is the same challenge described above of transforming the raw forecasts into scores, and the additional challenge of combining those scores to build an overall alpha.

Let's call the two raw forecasts g and h; these could be the broker's buy-sell recommendations of Example Number 3 and the DDM outputs from Example Number 4. We can assign an IC for each of the raw forecasts; call them IC(g) and IC(h). Let Score(g, n) be the score for stock n from the raw forecasts g, and Score(h, n) be the score for stock n from the raw forecast h.

The easiest way is to build up alphas for each raw forecast separately and then *add* (not average) them. We have

$$\text{Alpha}(n) = \text{Volatility}(n)[\text{IC}(g)\text{Score}(g, n) + \text{IC}(h)\text{Score}(h, n)] \quad (2)$$

This is the best thing to do if the process that generates raw forecast g is uncorrelated with the process that generates raw forecast h. If we believe that the processes that generate g and h are correlated (for example, both are based on some valuation measure such as yield or E/P), we can make a correction to the ICs.

Let $\rho(g, h)$ be the correlation of the process generating the raw F and G forecasts (same correlation for each stock), and let $\text{IC}(g)^*$ and $\text{IC}(h)^*$ be the effective ICs as shown below:

EXHIBIT 7
Alpha Consumption

MMI Company	Residual Volatility (%)	Real Case Alpha (%)	Real Case Positions	Edible Case Alpha (%)	Edible Case Positions
MMM	13.41	0.42	4.38	1.00	13.10
GE	14.42	1.01	2.65	1.00	5.69
ATT	15.89	1.01	4.40	1.00	5.73
P&G	16.29	−3.40	−11.19	−1.00	−4.29
Dow Chem	16.93	−1.17	−5.74	−1.00	−6.39
DuPont	17.29	2.47	10.38	1.00	8.56
Coca-Cola	18.92	−0.82	−0.79	−1.00	−4.85
J&J	18.97	−3.02	−6.71	−1.00	−4.79
Disney	19.17	0.62	2.93	1.00	5.57
Kodak	19.20	−0.10	−0.93	−1.00	−6.37
Intl Paper	19.83	−0.05	−0.99	−1.00	−4.94
Philip M	20.17	−1.62	−4.04	−1.00	−5.13
Merck	20.43	1.36	5.31	1.00	4.97
Chevron	20.44	−0.45	−4.25	−1.00	−3.88
McDonald's	20.54	−0.82	−1.76	−1.00	−4.74
Exxon	21.13	−0.01	0.81	−1.00	−2.59
Sears	22.33	1.70	2.38	1.00	2.57
Amex	23.26	0.73	−0.71	1.00	0.93
GM	23.46	4.17	7.45	1.00	3.01
IBM	30.32	−1.84	−3.57	−1.00	−2.18

$$IC(g)^* = [IC(g) - \rho(g, h)IC(h)]/[1 - \rho(g, h)^2] \quad (3)$$

$$IC(h)^* = [IC(h) - \rho(g, h)IC(g)]/[1 - \rho(g, h)^2] \quad (4)$$

Then use Equation (2) with $IC(g)^*$ replacing $IC(g)$ and $IC(h)^*$ replacing $IC(h)$.

For example, if $IC(g) = 0.07$, $IC(h) = 0.04$, and $\rho(g, h) = 0.35$, we would get $IC(g)^* = 0.0638$ and $IC(h)^* = 0.0176$. The stronger of the two raw forecasts, g, is hardly changed, but the weaker raw forecast h has its effective IC cut in half.

Example Number 5: Real Alphas Don't Get Eaten

Many have heard fables of alphas being eaten by optimizers. The optimization is said to "eat" the alphas when the portfolios that emerge from the optimizer don't emphasize the stocks with the most positive alphas and fail to de-emphasize the stocks with the most negative alphas. The alpha message fails to get through to the portfolio. As in all fables there is a moral: Alphas generated in an offhand manner can be eaten.

The cause of alpha consumption is not the optimization; it is the alphas. If there is a thoughtful and consistent approach to generating the alphas, you can avoid alpha consumption. To illustrate this problem we use the alphas generated in Exhibit 3 for the MMI stocks. We will call these the "real alphas." The real alphas, shown in Exhibit 7, are generated by the rule Volatility × IC × Score.

A second set of "edible" alphas are generated in a shortsighted way. Suppose we get our information as buy and sell signals. We will associate an edible alpha of +1.00% for the buy stocks and an edible alpha of −1.00% for the sell stocks. We have linked the edible to the real alphas by assigning a buy (+1.00%) to each stock with a positive "real alpha" and a sell (−1.00%) to each stock with a negative real alpha.[11] The edible alphas are also shown in Exhibit 7.

To test the effect of an optimizer on these real and edible alphas, we use the respective sets of alphas to generate optimal portfolios. The only requirements on the optimal portfolios are full investment in the MMI assets and a 4% active risk (tracking error) compared to the MMI. The optimal *active* positions using the real and edible alphas are included in Exhibit 7.[12] The MMI stocks in Exhibit 7 are sorted according to residual risk.

The effects of the alpha eating are apparent in this simple example. When the real alphas are used, the four largest active positions belong to the stocks with the four largest alphas: P&G, DuPont, J&J, and GM. When the edible alphas are used, all four of the largest active positions are among the stocks with lower than average volatility. In fact, for the edible optimal portfolio seven of the eight stocks with active positions greater than 5% (or less than −5%) are among the lower ten in residual volatility.[13]

It is easy to see why. With the edible alphas, we have two groups of stocks (the +1.00% and the −1.00%). An optimizer will not be able to distinguish within the groups except on the basis of volatility. If we ignore the impact of correlations, we see that the optimizer will quite naturally emphasize the lower residual risk assets.

SIDE EFFECTS

Alpha eating can be an unforeseen part of a systematic portfolio construction process if you have not been careful about how the alphas are constructed. The process can lead to a portfolio that will have an unintended bias toward lower residual volatility stocks.

It turns out that for most of the past twenty years in the U.S. a low-volatility bias would have had a beneficial effect on the portfolio. The same is true in Australia and the U.K. too. See the reported results on the low-volatility strategies in Grinold and Kahn [1992].

In other words: 1) you did not have alphas that were proportional to volatility; as a consequence, 2) your alpha was eaten, 3) your portfolio was biased toward lower-volatility stocks, and 4) you benefited from that low-volatility bias. You may not be so fortunate in the future. Any emphasis on lower-volatility stocks in the portfolio should be an explicit part of the portfolio strategy and not a fortunate side effect.

CONCLUSIONS

The residual volatility, the information coefficient, IC, and a standardized score are the three components of alpha. In several examples, we indicate how this insight can be used in structured and unstructured situations to analyze raw forecasts and to turn them into suitable forecasts of residual return. At a minimum, investors who absorb this lesson should be able to produce alphas of the right order of magnitude.

Systematic investors should use this information and examine their alpha-generating process to ensure that they are, in fact, making forecasts of residual return. In particular, they should make sure that the absolute value of the stock's alpha is (on average) roughly proportional to the stock's residual volatility. When alphas have little correlation with volatility, they stand a reasonable chance of being eaten by an optimizer.

ENDNOTES

Arjun Divecha insisted that I stay focused on the title line. Mark Engerman and Ronald Kahn made several useful suggestions. Andrew Rudd suggested Example Number 1.

[1]Traditional managers deal with alphas implicitly. It is possible, with the aid of a few useful and somewhat suspect assumptions, to infer the traditional manager's alphas by looking at the manager's portfolios; i.e., look at the portfolios, assume a portfolio construction process, and thus "reverse engineer" the alphas that would lead to that portfolio. This same procedure is used as a reality check by systematic managers.

[2]Excess return is the rate of return from the investment less the rate of return on a risk-free (T-bills) investment over the same time period.

[3]If r_n is the excess return on stock n, and r_B is the excess return on the benchmark, then the residual return is $\theta_n = r_n - \beta_n r_B$, where $\beta_n = \text{Covar}[r_n, r_B]/\text{Var}[r_B]$.

[4]Accomplished using a forecast of beta or using an ex post calculation of beta with a regression.

[5]If α is the vector of asset alphas, V is the covariance matrix of asset returns, and h_B are the holdings in the benchmark portfolio, then the holdings in the manager's portfolio are (ignoring all constraints) $h_P = h_B + \tau V^{-1}\alpha$, where τ is a measure of risk tolerance.

[6]A rough guideline for determining the required IC comes from Grinold [1989]. If you have N stocks, then a truly outstanding manager who has an information ratio of IR = 1.33 (corresponding to a t-stat of 3 over five years) will need an IC (for each stock!) given approximately by $IC \approx \{IR\}/(\#\text{ of Stocks})^{1/2} = 1.33/(500)^{1/2} = 0.06$. Top quartile might have (let's be generous) an information ratio of IR = 0.90 (t-stat of 2 over five years); thus the IC of $0.04 = 0.9/(500)^{1/2}$. These numbers are *rough guidelines*. The guideline can tell us that for 500 stocks and a quality manager ICs of 0.3 or 0.001 are out of range. The rough guideline will not help us tell if 0.03 or 0.04 is a better choice.

[7]One could imagine cases where the IC varies across groups of stocks. For example, if a company follows the Frank Russell 1000 very closely with its own analysts, and follows the Frank Russell 2000 with street and technical research, it might be justified in giving a smaller IC to the FR2000 stocks.

[8]Recalling the rough guideline $IC \approx \{IR\}/(\#\text{ of Stocks})^{1/2}$, a passable manager IR = 0.44 (t-stat of 1 over five years) with a universe of only twenty stocks will require an IC of about $0.44/(20)^{1/2} = 0.098$. This is high, but realism has been sacrificed in order to present examples with twenty familiar names.

[9]Write the raw forecasts as $f(n) = \psi\omega(n)z(n)$ where: ψ is a constant, $\omega(n)$ is the residual volatility of stock n, and z(n) is normally distributed with mean zero and standard deviation one. Assume that both z(n) and $z(n)^2$ are independent of $\omega(n)$. Then the correlation of the absolute values of the raw forecasts and the stock's volatility is

$$\rho = m/[1 + k^2(1 - m^2)]^{0.5}$$

where $m = 0.79788 = (2/\pi)^{0.5}$, and k is the ratio of the average residual volatility to the cross-sectional standard deviation of the residual volatilities. For the OEX at the end of July 1993, we have k = 23.42/6.81 = 3.4369, and thus $\rho = 0.34682$.

[10]In this example we assume that the raw forecasts are uncorrelated across the stocks. This is not a bad first approximation. As you

become more sophisticated in analysis, you might uncover some correlation in the raw forecasts. For example, we may tend to like all the stocks in one sector at one particular time or dislike all the stocks in the same sector at other times. In that case we should have a correlation between the raw forecasts.

[11]There is a variant of this story that we might have used. Suppose managers insist they can forecast the direction of the residual return but not the amount. In that case the scores (not the alphas!) are +1.00 or −1.00. The alphas will be either IC × residual volatility or −IC × residual volatility.

[12]The active position is the portfolio holding in the stock less the MMI holding in the same stock.

[13]You might be worried that the active positions are too large. Indeed some of them, like P&G in the real case, imply net short positions. You can just cut all the active positions in half. This reduces the active risk to 2.00%. It has no effect on the relative sizes of the active positions.

REFERENCES

Grinold, Richard C. "The Fundamental Law of Active Management." *Journal of Portfolio Management*, Spring 1989, pp. 30-37.

Grinold, Richard C., and Ronald N. Kahn. "Information Analysis." *Journal of Portfolio Management*, Spring 1992, pp. 14-21.