

FORECASTING INCENTIVES FOR AN AGGREGATED MUTUAL FUND

by

Carl Lundgren
VALMARPRO FORECASTING, INC.
2121 Columbia Pike, Suite 301
Arlington, VA 22204-4423

(703) 521-1037

Version Ib

July 1995

NOTICE OF COPYRIGHT

© Copyright 1997, Carl Allin Lundgren, All rights reserved.

The right to download and store or output this paper is granted to individual users for their personal use only. Governmental or business entities may internally distribute copies of this paper, provided that such internal distribution and use is for the sole purpose of conducting or contemplating the doing of business with Valmarpro Forecasting, Inc. or its affiliated companies. Any other reproduction, by any means - mechanical or electronic - without the express written permission of Valmarpro Forecasting, Inc. is strictly prohibited.

NOTICE OF PATENT ISSUED

U.S. Patent 5,608,620.

This paper refers to a method of economic incentives, involving plural-forecaster payment systems, upon which the author and inventor has been issued a patent. (U.S. Patent 5,608,620) The patent on that invention only restricts actual use of the described invention; it does not restrict in any way the verbal or written discussion, description, or criticism of that invention. Any actual use of the patented invention without the express written permission of Valmarpro Patents, Inc. is strictly prohibited.

The purpose of a mutual fund is to allow an unsophisticated investor to purchase a diversified portfolio of investments with a minimum of transaction costs, including the opportunity costs of information gathering and analysis. This purpose might be served well or badly by any given mutual fund. Even for a given level of risk, the returns on different mutual funds can vary quite substantially from each other. Empirically, this argues that the unsophisticated investor could be made better off if he bought a diversified portfolio of mutual funds, rather than just one or two mutual funds. This implies that an aggregated mutual fund--a mutual fund that holds a diversified portfolio of other mutual funds--will improve the risk-return trade-off for the typical investor.

Two types of diversification are possible. First, there is diversification of portfolios for all investors having the same risk preference. Second, there can be diversification across investor risk classes. That is, the typical portfolios which might be held by investors with different risk preferences will not have returns which are perfectly correlated. An aggregated mutual fund that serves more than one risk preference class can take advantage of this fact to reduce risks for the same level of returns for each risk preference class. Both kinds of diversification can make investors better off.

To the extent that an aggregated mutual fund attempts to diversify risks across investor risk classes, there is a need to predict the expected returns and variance of returns of the

aggregated portfolio, and perhaps also of the various mutual funds contained within the aggregated portfolio. Such predictions are necessary, at a minimum, in order to divide fairly the aggregate portfolio's risk and return among the various risk preference classes. The predictions are also useful to determine whether the total risk of the aggregated portfolio corresponds well with the aggregate risk preference of the various investor classes.

In order to elicit these predictions, this paper suggests application of a forecasting incentive scheme described in Lundgren (1994).¹ The forecasting incentive scheme is based on paying forecasters in accordance with their value marginal product (VMP) and has desirable incentive properties. This scheme can be applied to the aggregated mutual funds problem by asking forecasters to provide predictions for expected future returns and expected variances (or second moments) of returns for the aggregated mutual fund as a whole, and perhaps to the various mutual funds which make up that aggregated portfolio. Application of the VMP incentive scheme requires a) choice of a criterion value for judging the accuracy of forecaster predictions, b) choice of a loss function to use in the VMP formula, and c) choice of a method for aggregating the predictions of different forecasters.

There are at least three distinct functions which an aggregated mutual fund must perform. The first is to ascertain the risk preferences or utility functions of individual investors

and to allocate the risks and returns of the aggregated portfolio among investors according to their indicated risk preferences. The second is to allocate the assets of the aggregated mutual fund among various mutual funds and other assets (if any). The third is to obtain unbiased forecasts of the expected returns and risks of the aggregated portfolio (and perhaps of subportfolios as well). The first function can be performed by formulae, given the successful performance of the second and third functions.

The second function is similar to the functions performed by any portfolio manager. This paper is not directly concerned with the incentives given to the portfolio manager(s) of the aggregated mutual fund, which might or might not be made similar to those facing other types of portfolio managers. It is the effective performance of the third function to which this paper is primarily addressed.

The actual managers of a portfolio presumably do not have a proper incentive to provide the public with unbiased forecasts of the risks and returns. To the extent that their interest consists in selling more of their product, they will have an incentive to overstate the expected returns and understate the expected risk. Effective performance of the third function requires an independent assessment of the risks and returns of the aggregated mutual fund.

Conflict of interest considerations require that the individuals who perform the third function must be independent of the individuals who perform the second function. Likewise, the

incentives and compensation of these independent forecasters must be completely independent of the incentives and compensation provided to the individuals who select the assets of the aggregated mutual fund. Therefore, the forecasters and their incentives must be completely independent of the portfolio managers of the aggregated mutual fund.

This paper is organized as follows: Section I discusses the market inefficiency which implies the need for an aggregated mutual fund. Section II analyzes empirical evidence that the market efficiency appears to exist. Section III defines an aggregated mutual fund, and presents general formulae for splitting the returns of an aggregated portfolio among various investor classes. Section IV describes the VMP forecasting incentive method. Section V suggests a criterion value for use in the VMP formula. Section VI derives a loss function for use in the VMP formula. Section VII analyzes a possible method of aggregating forecaster predictions. Section VIII concludes.

I. What Market Inefficiency?

It is the presumed purpose of a mutual fund to diversify an investor's portfolio. If an aggregated mutual fund can improve performance for the typical investor, then this implies that the current market inventory of 3,000+ mutual funds are not doing a fully adequate job of diversifying investors' portfolios. This implies the existence of a market inefficiency, which can be exploited for presumed profit, and ultimately drive the market

toward greater efficiency and improvements in welfare for the typical investor. In essence, the claimed market inefficiency is that mutual funds, as they currently exist, are inadequately diversified.

This inadequate diversification likely comes from two sources: transaction costs and unsophisticated investors. Most investors are unsophisticated in that they lack knowledge of the relevant economic/financial theories governing optimal investment portfolios. The small minority of investors who do have such knowledge face significant information and computational costs of evaluating the risk-return trade-offs of different mutual funds. Mutual funds themselves are unlikely to become fully diversified, because of the economies to specializing in the analysis and trading of particular asset classes, and because of disagreements among fund managers about the most profitable investment strategies.

Ideally, a mutual fund investor should be able to engage in one-stop shopping. An investor should be able to choose from an uncomplicated menu of funds and choose one or two fund(s) which most closely match his risk preferences. A mutual fund should already be diversified. An investor should not need to worry about purchasing a diversified portfolio of mutual funds. Since diversification by the individual investor imposes extra transaction costs,² inadequate diversification by the mutual funds themselves implies an inefficiency in the mutual funds market.

This market inefficiency, if it exists, is different in kind from other inefficiencies which have been alleged with respect to the mutual funds market. For example, Jensen (1968) and Ippolito (1989) explore the question of whether mutual funds, on average, outperform or underperform the market relative to some absolute standard (i.e., relative to some diversified portfolio of primary assets, such as the S&P 500). Jensen (1968) finds evidence that no mutual fund outperforms the market, and that the average mutual fund underperforms the market. Ippolito (1989), on the other hand, finds that the mutual fund industry as a whole (not necessarily particular funds) appears to be efficient relative to an absolute standard.

Grinblatt and Titman (1992), Hendricks, Patel, and Zeckhauser (1993), and Goetzmann and Ibbotson (1994) analyze the question of whether some mutual funds perform relatively better than others. These three recent studies analyze monthly returns for different data sets and time periods and basically come to the same affirmative conclusion. Not only do some funds perform better than others ex post, these authors find that differences in mutual fund performance can be predicted ex ante based solely on measures of past performance. Hendricks, et al (1993) claims to show that selection of funds in accordance with a particular formula can improve expected returns by about 6% relative to the typical fund, and by about 3-4% relative to an absolute market standard.

Whether and to what extent it may be proper to select

particular funds based on expected performance is a topic which goes way beyond the scope of this paper. To the extent that it is possible to distinguish, ex ante, the likely performances of different mutual funds, this becomes an additional service which an aggregated mutual fund would be expected to perform on behalf of its clients. For purposes of this paper, it is sufficient to assume that diversification alone, with or without the ability to pick fund "winners" and avoid fund "losers", is an adequate justification for aggregating mutual funds.

II. Empirical Evidence of Inadequate Diversification

Casual inspection of tables of mutual funds returns indicates that different mutual funds have widely varying returns, even if we restrict our attention to funds which fall within the same "class." Such casual inspection tells us that there may be some gain to be had by diversifying one's assets into more than one such fund. Unfortunately, casual inspection cannot tell us how important this gain might be.

To answer this question, I analyze the annual returns of 313 stock mutual funds for the ten-year period, 1983-1992. This sample comes from CDA/Wiesenberger's Mutual Funds Panorama, 1993. All funds with 10 years of data are included, except for stock/bond funds, bond funds, and sector funds which are excluded.

Excess returns for each fund are calculated by considering only the return relative to the risk-free interest rate.³ The

observed average excess return and sample variance of excess returns was then calculated. It is clear from the data that ex post returns and returns variance do not perfectly correlate. However, if we assume investors are rational, we expect that average returns, expected ex ante, should be higher for funds which have higher variances in returns. To calculate the benefits of diversification, it is inappropriate to make calculations based on ex post returns without making some adjustment to determine what investors may have expected ex ante.

In order to make this adjustment, I maintain the assumption that each investor knew ex ante that his chosen fund would display the variance in returns that was observed ex post for his chosen class of fund over the ten-year period. The expected return for each fund class was then calculated according to a regression (without intercept term) of observed average excess returns on all fund classes as a linear function of observed sample standard deviations.

Based on the calculated expected returns and the observed sample variance of returns, each fund was assumed to be most attractive to an investor with a particular level of risk aversion. This level of risk aversion was determined by asking what level of risk aversion would be required for an investor to want to invest precisely 100% of his assets into a fund of that type, if the only alternative is to invest in the risk-free asset.⁴

Investors are assumed to have a quadratic utility function

of the normalized form, $U(W)=W-CW^2$. Each investor with wealth W may purchase k units of the risky asset and $W-k$ units of the riskless asset. The safe asset has excess return of zero; the risky asset has expected excess return of r , with a variance of σ^2 . The investor's expected utility is:

$$E(U) = W + kr - CW^2 - 2CWkr - Ck^2r^2 - Ck^2\sigma^2 \quad (1)$$

Solving for the first-order condition in k , we find that $k=\Phi Z$, where $Z=r/(r^2+\sigma^2)$ and $\Phi=(1-2CW)/(2C)$. Z is the expected return divided by the second moment. Φ is the absolute risk tolerance (inverse of absolute risk aversion) of the investor's utility function. If we set $W=1$ and ask what value of C in the utility function would yield $k=1$, we find that $C=Z/(2+2Z)$ and $\Phi=1/Z$. The implied relative risk aversion is Z .

We compute an aggregated portfolio by giving equal weight to each of the 313 funds. The average risk tolerance corresponding to these funds is 0.452 (relative risk aversion of 2.21). The 313 funds are divided into 10 classes of 15-43 funds each.⁵ The utility of each of the ten investor types is computed under three scenarios: 1) each investor type holds only one fund, 2) each investor type holds an aggregated fund consisting of the 15-43 funds within his investor class, 3) each investor type holds an aggregated portfolio of all 313 funds.

The computed utility of each investor type is converted into a dollar value certainty equivalent. In computing the utility of investors under scenarios (2) and (3), it was assumed that the aggregated mutual fund allowed investors to form any combination

(including negative combinations) of the safe asset with the risky aggregated portfolio. Because the aggregated portfolio reduces risk for any given level of expected return, it is optimal for the average investor to increase his exposure to the risky assets.

The results for each of the ten investor classes are shown in Tables 1(a) and 1(b). The average investor experiences a certainty-equivalent gain of 1.19% in annual returns from an aggregated mutual fund. This 1.19% increase is composed of a 0.93% increase from aggregation within each investor class, and an additional 0.26% increase from aggregating together all ten investor classes. Compared with the safe asset, the certainty equivalent value of mutual funds rises from 3.20% to 4.39% in annual excess returns--a 37% gain in value for the risky assets.

III. Careful Definition of an Aggregated Mutual Fund

An Aggregated Mutual Fund is an aggregate of mutual funds in which the risks and returns for each participant are not necessarily the same or proportional for each participant and do not necessarily correspond to the risks and returns of the whole portfolio or of any subportfolio within the aggregate mutual fund. That is, unlike the typical mutual fund or common stock, the aggregate mutual fund allows for disproportionate risk sharing among participants. With respect to such a fund, we can state the following informal theorem:

Theorem 1: In the absence of any tax disadvantages or

transaction costs of investing in an aggregated mutual fund, an aggregated mutual fund can always do at least as well (and sometimes better) for all participants as would separate portfolios for each participant.

We may supply the following verbal "proof" of this theorem: Suppose that without the aggregated mutual fund, each participant would hold the best individually rational portfolio of assets. The aggregated mutual fund can hold the same portfolio of assets for all participants, and thus can always do at least as well for each participant. In addition, the aggregated mutual fund may be able to formulate contracts to trade risks and returns among participants so that one, some, or all participants are made better off.

The question of whether an aggregated mutual fund really can do better is essentially an empirical question, which the previous section attempted to shed some light on, by showing that current mutual funds appear to be inadequately diversified. This section focuses on the theoretical question of how risks and returns of an aggregated fund might best be divided up among investor classes with different risk preferences.

The distribution problem of an aggregated mutual fund can be solved in general as follows:

Suppose that investors are organized into M utility (risk-preference) classes U_i , indexed by i , each with wealth W_i at time $t=0$. Total wealth of the aggregated fund is W , where $W = \sum W_i$.

Suppose that the possible portfolio outcomes are classified

into N mutually exclusive possible states of the world S_j , indexed by j , each having probability Z_j , where $\sum Z_j=1$. The total return including principal (net of management fees/transaction costs) of the aggregate portfolio is R_j at time $t=1$.

Each investor class may use its wealth to purchase dollars to be delivered at time $t=1$ in state S_j . The implicit price at $t=0$ of such dollar deliveries at $t=1$ is P_j . The total dollars delivered to U_i in S_j is I_{ij}/P_j , where I_{ij} is a dollar amount invested at $t=0$. Each P_j and I_{ij} is determined endogenously as the solution to a set of simultaneous equations.

The problem has three sets of equations: 1) Optimization: Each investor class must choose the I_{ij} 's to maximize its expected utility. 2) Budget Constraint: The sum of the I_{ij} 's invested at $t=0$ by utility class U_i must equal W_i for each utility class. 3) Market Constraint: The purchase of dollars in S_j by each U_i class must sum to $R_j W$ for each S_j .

Investor class i maximizes $\sum U_i(I_{ij}/P_j)Z_j$ subject to $\sum I_{ij}=W_i$. Investor i 's choice variables are: I_{ij} , $j=1$ to N . The Lagrangean and first-order conditions are:

$$\mathcal{L}_i = \sum U_i(I_{ij}/P_j)Z_j + \Omega_i[W_i - \sum I_{ij}] \quad (2)$$

$$\delta \mathcal{L}_i / \delta I_{ij} = U_i'(I_{ij}/P_j)Z_j/P_j - \Omega_i = 0, \text{ for all } MN \text{ } i, j.$$

$$\delta \mathcal{L}_i / \delta \Omega_i = W_i - \sum I_{ij} = 0, \text{ for all } M \text{ } i.$$

The market constraint is:

$$\sum I_{ij}/P_j = R_j W, \text{ for all } N \text{ } j. \quad (3)$$

Hence, there are $MN+M+N$ equations in the following $MN+M+N$ endogenous variables: I_{ij} , Ω_i , and P_j .

If we think in terms of numerous (or even a continuum of) states of the world, it may be easier to think in terms of probability-normalized variables. Define $Q_j = P_j/Z_j$ (probability-normalized price of dollar in state j), and define $J_{ij} = I_{ij}/Z_j$ (probability-normalized investment by i at $t=0$ for dollar deliveries in state j). The three equations then become:

$$U_i'(J_{ij}/Q_j) = \Omega_i Q_j \quad (4)$$

$$W_i = \sum J_{ij} Z_j$$

$$\sum J_{ij} = Q_j R_j W$$

If there were only one utility class, we could eliminate the i subscript and solve as follows:

$$Q_j = U'(R_j W) / \Omega \quad (5)$$

$$\Omega = \sum U'(R_j W) R_j Z_j$$

$$J_j = Q_j R_j W$$

The Q_j variables trace out an implicit or "aggregate utility function" which can serve as an objective function which the aggregated portfolio manager(s) can be asked to maximize. Though it may be useful as an objective function, if there is more than one utility class of investors, the aggregate utility function will not function as a true utility function, since it will change if the distribution of wealth among investor classes changes, or if the expected risks and returns of the aggregated portfolio changes.

IV. The VMP Forecasting Incentive Method

As indicated in the introduction, one of the functions

required by the aggregated mutual fund is an objective appraisal of the risks and returns of the aggregated portfolio, so as to fairly allocate risks and returns among investor risk classes. To avoid conflicts of interest, it is proposed that a group of forecasters, separate and distinct from the portfolio managers, be chosen to perform the function of forecasting expected returns and risks. The remainder of this paper shows how a system of general forecasting incentives, described in Lundgren (1994), can be adapted to resolve the specific problem of obtaining predictions useful for implementing an aggregated mutual fund.

Lundgren (1994) describes a method of paying forecasting incentives based on value marginal product (VMP). This technique assumes that there is a principal (forecast requisitioner) whose goal is to obtain an accurate prediction concerning the future realization of a random variable X . This goal is to be accomplished indirectly, rather than directly, by hiring a set of agents (forecasters) who will do the actual forecasting. The method requires at least three elements: 1) a criterion value to determine the accuracy of predictions, 2) a loss function, and 3) a prediction aggregator function.

The forecast requisitioner aggregates the predictions of individual forecasters to obtain a collective prediction suitable for further action. A typical method of aggregation might be to take an average or weighted average of forecasters' predictions, such as an arithmetic mean or a geometric mean. Let X_c represent the vector of individual predictions, X_1, X_2, \dots, X_n of

forecasters 1, 2, ..., n. A prediction aggregator function can be generalized as follows:

$$G(X_c) = G(X_1, X_2, X_3, \dots, X_n) \quad (6)$$

Let X_{ci} represent the vector of predictions of all forecasters except forecaster i . $G(X_{ci})$ represents a "secondary collective prediction," which would presumably be issued in the absence of forecaster i 's prediction.

The loss function, $L(X_a, G(X_c))$, tells us the lost benefits which occur when $G(X_c)$ is the collective prediction of X , while X_a is an actual or estimated value of X which is later observed. The value X_a can be used as a "criterion value"--a variable value which is used to judge the accuracy or inaccuracy of forecasters' predictions. If the actual value of the variable being predicted is observed within a reasonable period of time, it is natural to use the actual variable value as the criterion value. Otherwise, it is necessary to use a proxy.

The VMP incentive scheme attempts to measure forecaster VMP and pay in accordance therewith. In Lundgren (1994) it is shown that the VMP incentive scheme has various favorable properties. For instance, the scheme results in unbiased forecasts when forecasters are risk neutral, and nearly unbiased forecasts when forecasters are risk averse. The scheme results in optimal effort levels by each forecaster, and attracts nearly optimal numbers of forecasters to the forecasting task.

The VMP incentive method uses a proxy for VMP which may be termed "marginal contribution." The marginal contribution asks

how the value of a collective forecast changes, when the prediction of a particular forecaster is either contributed or withheld. The marginal contribution of forecaster i towards the accuracy of the collective forecast can be given by the equation:

$$MC_i = L(X_a, G(X_{ci})) - L(X_a, G(X_c)) \quad (7)$$

The marginal contribution for a particular forecaster in a particular instance might well be positive, zero, or negative, depending on whether X_i moves the collective forecast towards or away from X_a .

Hence, the VMP incentive method uses the following type of pay schedule for a forecaster:

$$P_i(X_i, X_{ci}, X_a) = F + kL(X_a, G(X_{ci})) - kL(X_a, G(X_i, X_{ci})), \quad (8)$$

where $k > 0$.

The payment schedule in (8) is simply a linear (affine) transformation of the VMP formula in equation (7).

V. Criterion Values for Predictions of the Excess Return

Open-end mutual funds typically value their portfolios at the end of each business day, so that investors may deposit or withdraw assets at that time. An aggregated mutual fund would want to know the likely distribution of expected returns from one business day to the next. Ideally, one might wish to obtain a prediction concerning the whole probability distribution of returns. Practically speaking, it would be simpler to assume that the returns follow some standard distribution, and then ask forecasters to predict the values of certain moments of that

distribution.

The probability distribution of actual returns is said to most closely follow the log-normal distribution. If this is a sufficiently good approximation for prediction of daily returns, then the parameters of interest are the mean of the log distribution and either the variance or the second moment of the log distribution. By a well-known statistical identity, the second moment is equal to the mean squared plus the variance. Daily predictions concerning the expected values of at least two of these three numbers can then be solicited from expert forecasters.

Because neither the mean, nor the variance, nor the second moment of a distribution are directly observed on any business day, it is necessary to use proxies for the criterion values which are used to judge the accuracy of forecaster predictions. The most appropriate criterion value for predictions concerning the mean of the log distribution is the observed logarithm of ex post actual excess returns, since the observed logarithm will on average equal the mean. In mathematical notation, if r is the unobserved mean of the hypothetical distribution, r_e is the collective estimate of r , and r_a is the observed excess return over the riskless rate, then r_a is used as the criterion value to determine the accuracy of r_e .

For predictions concerning the log-variance, the most appropriate criterion value is the square of the observed logarithmic deviation of actual returns from its predicted mean.

This is because, on average, we expect the square of this deviation to equal the log-variance. In mathematical notation, if σ^2 is the unobserved variance of the hypothetical distribution, σ_e^2 is the collective estimate of σ^2 , and $\sigma_a^2 = (r_a - r_e)^2$ is the square of the deviation of observed excess return from the estimated excess return, then σ_a^2 is used as the criterion value to determine the accuracy of σ_e^2 .

For predictions concerning the second moment, the square of the observed logarithm of actual excess returns is the correct criterion value, since the square will on average equal the second moment. In mathematical notation, if $S = r^2 + \sigma^2$ is the unobserved second moment of the hypothetical distribution, S_e is the collective estimate of S , and $S_a = r_a^2$ is the square of the observed excess return over the riskless rate, then S_a is used as the criterion value to determine the accuracy of S_e .

VI. Format of Loss Function

There are at least three possible types of losses which might occur from inaccurate prediction of the mean and variance of the aggregated portfolio. One is an efficiency loss to the investors in the fund, if there is as a result either too little or too much investment in risky assets by the aggregated portfolio. Two is an equity loss (which might or might not have efficiency components) which results if misprediction of the likely probability distribution of returns leads to a different distribution of fund assets between investors of different risk

classes. There is a social efficiency loss if errors of prediction cause investments in risky assets to be lower than might otherwise be the case.

Of these three losses, only the first type of loss is likely to be of unanimous concern to the private investors in an aggregated mutual fund. To determine the approximate size and form of this loss, assume that the aggregated utility function can be approximated by a quadratic utility function of the normalized form, $U(W)=W-CW^2$. The safe asset has excess return of zero; the risky asset (i.e., the aggregated portfolio) has expected excess return of r , with a variance of σ^2 .

Each investor with wealth W may purchase k units of the risky asset and $W-k$ units of the riskless asset. The investor's expected utility is:

$$E(U(k)) = W + kr - CW^2 - 2CWkr - Ck^2r^2 - Ck^2\sigma^2 \quad (9)$$

Solving for the first-order condition in k , we find that $k=\Phi Z$, where $\Phi=(1-2CW)/(2C)$ and $Z=r/(r^2+\sigma^2)$. Suppose r and σ are estimated to be equal to r_e and σ_e . This will cause the amount invested in the risky asset by the portfolio manager to equal $k_e=\Phi Z_e$, where $Z_e=r_e/(r_e^2+\sigma_e^2)$. The expected utility to the fund holders is therefore:

$$E(U(k_e)) = W + k_e r - CW^2 - 2CWk_e r - Ck_e^2 r^2 - Ck_e^2 \sigma^2 \quad (10)$$

The expected utility loss to fund investors is $E(U(k))-E(U(k_e))$. In order to get the dollar value of this loss, we subtract (10) from (9) and divide through by the marginal utility of income, $1-2CW$:

$$\text{Loss} = (k - k_e)r - (k^2 - k_e^2)(r^2 + \sigma^2)/2\Phi \quad (11)$$

Define $S = (r^2 + \sigma^2)$ and $S_e = (r_e^2 + \sigma_e^2)$, then $Z = r/S$ and $Z_e = r_e/S_e$. S is the second moment of the distribution of excess returns, and S_e is the estimated second moment. After substitution for k and k_e and a little manipulation, it can be shown:

$$\text{Loss} = \Phi(rS_e - r_eS)^2 / (2SS_e^2) \quad (12)$$

If we assume that $S = S_e$, then (12) reduces to:

$$\text{Loss} = \Phi(r - r_e)^2 / (2S) \quad (13)$$

Since the loss function is quadratic in r_e , the goal of the forecasters should be to predict the mean of the distribution of r .

If we assume that $r = r_e$, then (12) reduces to:

$$\text{Loss} = \Phi r^2 (S_e - S)^2 / (2SS_e^2) \quad (14)$$

If we assume that S and S_e are close together in value, then the loss function is essentially quadratic in S_e . Therefore, the goal of the forecasters should be to predict the mean of the distribution of S (the second moment of r).

Given that r and S are never directly observed, it is necessary to use the proxies, r_a and $S_a = r_a^2$ as criterion values to judge the accuracy of collective estimates of r and S . Let r_c and S_c represent the vector of individual predictions of r and S by all n forecasters, and let r_{ci} and S_{ci} represent the vector of predictions of all forecasters except forecaster i . $G(r_c)$ and $G(S_c)$ represent the collective predictions of r and S , while $G(r_{ci})$ and $G(S_{ci})$ represent the secondary collective predictions that would be issued in the absence of forecaster i 's prediction.

Given r_a and S_a as criterion values, we can define proxy loss functions for r_e and S_e as follows:

$$L(r_a, G(r_c)) = \Phi(r_a - G(r_c))^2 / (2G(S_c)) \quad (15)$$

$$L(S_a, G(S_c)) = \Phi G(r_c)^2 (S_a - G(S_c))^2 / (2G(S_c)^3) \quad (16)$$

Similarly, we can define secondary loss functions as follows:

$$L(r_a, G(r_{ci})) = \Phi(r_a - G(r_{ci}))^2 / (2G(S_{ci})) \quad (17)$$

$$L(S_a, G(S_{ci})) = \Phi G(r_{ci})^2 (S_a - G(S_{ci}))^2 / (2G(S_{ci})^3) \quad (18)$$

The above loss functions can be plugged into the payment formula in (8) to motivate unbiased predictions from forecasters. It only remains to determine the nature of the aggregation functions, $G(r)$ and $G(S)$.

VII. Aggregating Forecaster Predictions

One method for aggregating predictions is to take a weighted mean of each forecaster's prediction of a particular value. In Lundgren (1994), an example is worked out whereby risk-neutral forecasters attempt to predict the expected value of a random variable, X , having a normal distribution with unknown mean. Each forecaster is assumed to have both common information and private information. Each forecaster's private information has a different level of precision.

In the example, each forecaster submits to the requisitioner both a precision weight, T_i , and a prediction function, $X_i(T)$, where T represents a sum of precision weights. The collective prediction is computed as a weighted mean of the $X_i(T)$ predictions, where the weight T_i is applied to each X_i , and the

value T is set equal to the sum of all precision weights. Each secondary collective prediction is computed as a similar weighted mean of the $X_i(T)$ predictions, except that one of the predictions is excluded and the value T is set equal to the sum of all precision weights except for the precision weight of the excluded prediction.

In Lundgren (1994) this method of aggregation is shown to have favorable properties. In Nash equilibrium, each forecaster is motivated to submit the correct precision weight and prediction function, which when aggregated produces the correct collective prediction, given the combined information of all the forecasters. Given an appropriately weighted loss function in the payment formula, each forecaster also has incentive to exert the socially optimal amount of effort.

A similar method of aggregation would presumably work well to elicit predictions of the expected value of r , since r is essentially the mean of a probability distribution. However, it is not immediately clear whether the same method of aggregation would work well to elicit correct predictions of the standard deviation, variance, precision, or second moment of a probability distribution, such as σ , σ_2 , $1/\sigma_2$, or $S=r_2+\sigma_2$ in our example.

To get a handle on this issue, we can invent a metaphor for the forecasting task. Each forecaster is assumed able to imagine different scenarios for the future. Each scenario has associated with it a particular rate of excess return, r_k . Each scenario is assumed equally likely to come true, regardless of which

forecaster imagines it.⁶ Some forecasters may imagine more scenarios than others. The combined sample of imagined scenarios is assumed to be representative of the population of all scenarios, both imagined and unimagined.⁷

For initial simplicity, assume that each scenario is imagined by only one forecaster as private information. Each forecaster submits a precision weight which is equal to the number of privately imagined scenarios. For purposes of computing a mean value of the distribution, each forecaster submits the mean value of r_k from his private set of scenarios. For purposes of computing a second moment of the distribution (S), each forecaster submits the mean value of r_k^2 from his private set of scenarios. The weighted means of r_k and r_k^2 for all forecasters will correctly aggregate to give the correct mean of r_k and r_k^2 for all scenarios combined.

Using the statistical identity, that the variance is equal to the second moment minus the mean squared, one can compute the implied variance, standard deviation, and precision of the distribution (σ_2 , σ , and $1/\sigma_2$). However, the variance is not capable of being directly computed as a weighted mean of the implied variances observed by individual forecasters. This is because the variance is computed relative to the mean of the observations, and different forecasters may have observed different mean values for their private sets of r_k 's.

Hence, even if the statistic of interest to the forecast requisitioner is the expected value of the variance, the standard

deviation, or the precision, rather than the second moment, it is necessary either to solicit values of the second moment or to use more complicated formulas for aggregating predictions of the variance or its functional derivatives. Either course can be pursued, but it is simpler to think in terms of soliciting expected values of the mean and the second moment, using a weighted mean as the preferred method for aggregating the mean and the second moment, and then computing the variance, standard deviation, or precision as desired. Perhaps coincidentally, the particular application of this paper requires only a determination of the mean and the second moment, which are the easiest statistics to aggregate.

For added complexity in the metaphor, assume that a number (T_z) of scenarios are observed or imagined in common by all forecasters. In addition to the commonly observed scenarios, there are privately imagined scenarios, each of which is imagined by only one forecaster each. Again, each forecaster submits a precision weight (T_i) which is equal to the number of privately imagined scenarios. For purposes of computing a mean value (r_z) and a second moment (S_z) for the commonly observed scenarios, each forecaster computes the mean value of r_k and r_k^2 from the common set of scenarios. For purposes of computing a mean value (r_{pi}) and a second moment (S_{pi}) for the privately imagined scenarios, each forecaster i computes the mean value of r_k and r_k^2 from his private set of scenarios.

It was shown in Lundgren (1994) that risk-neutral

forecasters have incentive under the VMP incentive scheme to submit a set of predictions which aggregate to bring about the best collective prediction, given the total information at their collective disposal. Information will not be wasted or withheld. This assumes, however, that the aggregation procedure is adequate to the task of optimally combining forecaster information. To test the sufficiency of the aggregation procedure, we determine whether the aggregation procedure can achieve the first-best aggregation of the total information available. If it can, forecasters under the VMP incentive method will have incentive to submit predictions which properly aggregate.

The best method of aggregation is to take a weighted mean of all the imagined scenarios, both commonly imagined and privately imagined, as follows:

$$r_T = (T_z r_z + \sum T_i r_{pi}) / (T_z + \sum T_i) \quad (19)$$

$$S_T = (T_z S_z + \sum T_i S_{pi}) / (T_z + \sum T_i) \quad (20)$$

r_T and S_T represent the best possible aggregation of the information available to the forecasters. This best possible aggregation can be obtained if each forecaster i submits the precision weight, T_i , and the following prediction functions:

$$r_i(T) = (T_z r_z + T r_{pi}) / (T_z + T) \quad (21)$$

$$S_i(T) = (T_z S_z + T S_{pi}) / (T_z + T) \quad (22)$$

The forecast requisitioner then aggregates the individual predictions by taking a weighted mean, where $T_c = \sum T_i$:

$$G(r_c) = (\sum T_i r_i(T_c)) / T_c \quad (23)$$

$$G(S_c) = (\sum T_i S_i(T_c)) / T_c \quad (24)$$

It can be verified algebraically that $G(r_c)=r_T$ and $G(S_c)=S_T$. Since the best method of aggregation is feasible under this simple aggregation procedure, the efforts of forecasters to provide the collectively best set of predictions are not hampered. Hence, under the VMP forecaster payment method shown in (4), forecasters will endeavor to provide this collectively best set of predictions.⁸

VIII. Conclusion

This paper identifies an empirical need and a theoretical prescription for an aggregated mutual fund. The empirical need arises from the inadequate diversification of currently existing mutual funds. An aggregated mutual fund which holds other mutual funds can effectuate a certainty-equivalent improvement in expected annual returns of 1.19% (a 37% improvement) simply through added diversification. 0.93% (78%) of this diversification benefit comes from aggregation of mutual funds within investor classes, and 0.26% (22%) of this benefit comes from aggregation between investor classes.

The theory of how risks and returns of an aggregated portfolio are properly allocated among investor classes with different risk preferences was discussed. Implementation of the theory requires some estimate of probable returns and returns volatility. An application of the VMP forecasting incentive method described in Lundgren (1994) was suggested as a good way to elicit unbiased predictions of the necessary information.

This application requires that there be a group of independent forecasters whose incentives are separate and distinct from any incentives provided to the portfolio managers of the aggregated mutual fund.

Application of the VMP incentive scheme requires three elements: 1) choice of a criterion value for judging the accuracy of forecaster predictions, 2) choice of a loss function to use in the VMP formula, and 3) choice of a method for aggregating the predictions of different forecasters. It was shown that the most appropriate criterion values were ex post observed excess returns, R_a , and the square of excess returns, R_a^2 . The most appropriate loss functions were shown to be quadratic in the expected mean, r , and second moment, $S=r^2+\sigma^2$, of the distribution of excess returns. Finally, it was shown that an optimal method for aggregating forecaster predictions is to take a weighted mean of forecasters' predictions of the mean and second moment of the returns distribution, where each forecaster submits the weights to be used.

ENDNOTES

1. The method described in Lundgren (1994) is the subject of a patent (U.S. Patent 5,608,620). Contact author for details.
2. Although no-load mutual funds do not charge sales commissions or withdrawal fees, most such funds do have minimum investment requirements (typically around \$1,000-\$2,000). For some individuals, it may not be feasible to diversify a portfolio of mutual funds. Even where feasible, it is certainly not convenient.
3. Risk-free returns are calculated from the one-month treasury bill annual rates provided by Ibbotson Associates, Stocks, Bonds, Bills, and Inflation (SBBI).
4. Given the assumed linear structure of expectations, this is equivalent to asking which fund an investor of the given risk class will choose to invest in, if each investor must place all his assets into only one fund.
5. The fund classes were chosen from the following categories grouped by CDA/Wiesenberger: Maximum Capital Gains (MCG), Small Company Growth (SCG), International Equity (INT), Long Term Growth (LTG), Growth and Current Income (GCI), and Equity Income (IEQ). Three of the categories were sorted into high/low or high/middle/low classes based on variance of returns. The ten classes and number of funds in each class were: MCG-1,2 (24/25); SCG (16); INT (23); LTG-1,2,3 (43/43/43); GCI-1,2 (41/40); and IEQ (15).
6. The assumption is intended only for simplicity. One can derive the same conclusions of this section, even if scenarios are to be weighted according to some measure of probability or plausibility, or if scenarios come equipped with a probability distribution for the r_k 's rather than having a fixed r_k .
7. If the imagined sample is not representative, there is presumably no humanly possible way to correct the situation. Under the incentive scheme, forecasters have every incentive to correct for any known deficiency in the representativeness of scenario imagination.
8. See Proposition 2 in Lundgren (1994). Lest this seem like a trivial feat, it should be realized that a collection of individually best predictions is not equivalent to a set of collectively best predictions. A collectively best prediction must optimally aggregate the information available to all forecasters combined, whereas an individually best prediction relies only on the information available to a particular individual.

Table 1(a). The Diversification Value of Aggregating Mutual Funds.

Investor Class	Mean Φ	Certainty Equivalent (1)	Excess Returns (2)	Returns (3)
MCG-1	0.682	4.816	6.355	6.614
SCG	0.637	4.497	5.444	6.176
LTG-1	0.585	4.136	5.491	5.680
INT	0.563	3.976	5.380	5.460
MCG-2	0.449	3.175	4.466	4.360
LTG-2	0.421	2.971	3.605	4.080
GCI-1	0.400	2.828	3.627	3.884
IEQ	0.330	2.331	2.883	3.202
LTG-3	0.324	2.289	2.963	3.143
GCI-2	0.308	2.175	2.684	2.987
All	0.452	3.196	4.131	4.389

Certainty Equivalent Returns: In Scenario (1) each investor type holds a single mutual fund. In Scenario (2) each investor type holds an aggregated mutual fund consisting of the 15-43 funds within his investor class. In Scenario (3) each investor type holds an aggregated mutual fund consisting of all 313 mutual funds.

Table 1(b). Diversification Gain from Aggregating Mutual Funds.

Investor Class	Mean Φ	Certainty Equivalent Total	Annual Gains Intra-Class	Annual Gains Inter-Class
MCG-1	0.682	1.798	1.538	0.260
SCG	0.637	1.679	0.947	0.732
LTG-1	0.585	1.544	1.356	0.188
INT	0.563	1.484	1.404	0.080
MCG-2	0.449	1.185	1.291	-.106
LTG-2	0.421	1.109	0.634	0.475
GCI-1	0.400	1.056	0.799	0.257
IEQ	0.330	0.870	0.552	0.318
LTG-3	0.324	0.854	0.674	0.180
GCI-2	0.308	0.812	0.509	0.303
All	0.452	1.193	0.935	0.258

Certainty Equivalent Annual Gains computed from Table 1(a) as follows: Total Gains = (3) - (1).

Intra-Class Gains = (2) - (1).

Inter-Class Gains = (3) - (2).

REFERENCES

- Goetzmann, William N. and Roger G. Ibbotson (1994), "Do Winners Repeat? Patterns in Mutual Fund Behavior," Journal of Portfolio Management, vol. 20, no. 2, Winter, pp. 9-18.
- Grinblatt, Mark and Sheridan Titman (1992), "The Persistence of Mutual Fund Performance," Journal of Finance, vol. 47, no. 5, December, pp. 1977-1984.
- Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser (1993), "Hot Hands in Mutual Funds: Short-Run Persistence of Relative Performance, 1974-1988," Journal of Finance, vol. 48, no. 1, March, pp. 93-130.
- Ippolito, Richard A. (1989), "Efficiency with Costly Information: A Study of Mutual Fund Performance," Quarterly Journal of Economics, February, pp. 1-23.
- Jensen, Michael (1968), "The Performance of Mutual Funds in the Period 1945-1964," Journal of Finance, vol. 23, no. 2, May, pp. 389-416.
- Lundgren, Carl (1994), "Forecasting Incentives Based on Value Marginal Product," Working Paper, September.