

Where Do Alphas Come From?: A New Measure of the Value of Active Investment Management*

Andrew W. Lo[†]

First Draft: March 5, 2007

Latest Revision: May 8, 2007

Abstract

The value of active investment management is traditionally measured by alpha, beta, tracking error, and the Sharpe and information ratios. These are essentially static characteristics of the marginal distributions of returns at a single point in time, and do not incorporate dynamic aspects of a manager's investment process. In this paper, I propose a new measure of the value of active investment management that captures both static and dynamic contributions of a portfolio manager's decisions. The measure is based on a decomposition of a portfolio's expected return into two distinct components: a static weighted-average of the individual securities' expected returns, and the sum of covariances between returns and portfolio weights. The former component measures the portion of the manager's expected return due to static investments in the underlying securities, while the latter component captures the forecast power implicit in the manager's dynamic investment choices. This measure can be computed for long-only investments, long/short portfolios, and asset allocation rules, and is particularly relevant for hedge-fund strategies where both components are significant contributors to their expected returns, but only one should garner the high fees that hedge funds typically charge. Several analytical and empirical examples are provided to illustrate the practical relevance of these new measures.

*I thank Nicholas Chan, John Cox, Arnout Eikeboom, Lisa Goldberg, Stephanie Hogue, Rajnish Kamat, Philippe Luedi, and participants at the Le Club B 2006 Conference, and the JOIM 2007 Spring Conference for helpful comments and discussion. Research support from AlphaSimplex Group is gratefully acknowledged.

[†]Harris & Harris Group Professor, MIT Sloan School of Management, and Chief Scientific Officer, AlphaSimplex Group, LLC. Please direct all correspondence to: Andrew W. Lo, AlphaSimplex Group, One Cambridge Center, Cambridge, MA 02142.

Contents

1	Introduction	1
2	Literature Review	3
3	The AP Decomposition	5
3.1	The General Result	6
3.2	A New Definition of Passive Investing	10
3.3	Alpha vs. Beta	11
4	Some Analytical Examples	15
4.1	A Numerical Example	15
4.2	Mean Reversion and Momentum Strategies	19
4.3	A Stop-Loss Policy	20
5	Implementing the AP Decomposition	22
5.1	Population vs. Sample Moments	22
5.2	GMM Estimation	24
5.3	The Sampling Interval	24
6	An Empirical Application	27
6.1	Trading NASDAQ Size-Deciles	29
6.2	Daily vs. Monthly Transparency	30
7	Conclusion	32
A	Appendix	34
A.1	GMM Estimators for the AP Decomposition	34
	References	37

1 Introduction

With the growing popularity of hedge funds and other absolute-return investment strategies, there is a widening gap between the performance metrics of traditional investment management and alternatives. While alpha, beta, volatility, tracking error, the Sharpe ratio, and the information ratio have become the standard tools for gauging the value-added of long-only portfolio managers, they have not had as much impact among investors of absolute-return strategies. Part of this gap is no doubt cultural in origin; the growth of the mutual-fund industry was accelerated by the broad acceptance of portfolio theory and the benefits of diversification. This, in turn, led to the push for indexation and benchmark-based performance attribution, from which many of the current performance measures emerged.

However, another possible reason for the lack of impact of traditional performance measures for alternative investments is the fact that such measures are static, and do not capture the dynamic and predictive nature of active investment strategies. Specifically, measures such as alpha, beta, tracking error, and the information ratio are all functions of parameters of the portfolio-return and benchmark-return distributions at a single point in time, e.g., expected returns, covariances, and variances. None of these measures involves the relation between returns at *multiple points in time*, yet such multi-point statistics are often the central focus of active investment strategies. For example, the most admired portfolio managers of our time are revered for their ability to foresee certain market trends well in advance of the public, or to detect mispriced securities and exploit them ahead of the market, or to enter or exit certain investments before others recognize the opportunities. In every case, these investment skills involve forecasts or predictions, yet the standard performance measures listed above do not depend explicitly on the forecast power of the portfolio manager.

In this paper, I propose a new measure of the economic value of active management—an active/passive or “AP” decomposition—that takes into account forecast power explicitly. It is a simple decomposition of the expected return of a portfolio into two components: one that depends only on the average values of portfolio weights and asset returns, and another that depends on the correlation between portfolio weights and returns. It is this latter component that measures directly the value of active management—the portfolio weights of a successful manager will generally be positively correlated with returns, yielding a positive

contribution to the portfolio's expected return. This correlation is directly affected by the manager's forecasting abilities because portfolio weights are functions of the manager's prior information. Therefore, the correlation between portfolio weights and returns at date t is a measure of the predictive power of the information used by the manager to select his date- t portfolio weights. In short, it is a measure of the manager's asset-timing ability.

Of course, it is possible to generate positive expected returns without any variability in portfolio weights: a buy-and-hold strategy in assets with positive risk premia such as the S&P 500 will yield a positive expected return. In this case, the active component described above will contribute nothing to the portfolio's expected return, hence the portfolio can be said to be passive. This is a novel definition of passive and active investing, and has little to do with the standard definitions involving deviations from a benchmark portfolio. I show that a more natural definition for a passive portfolio is one where the portfolio weights are uncorrelated with returns. If weights have no forecast power, then active management is adding no value and the only source of expected return is risk premia, which can usually be generated by a buy-and-hold portfolio.

The AP decomposition is a simple consequence of the definition of covariance, and is trivial for active managers to implement. In fact, position-level information is not necessary—only average portfolio weights and individual-asset average returns are needed to perform the decomposition. Moreover, because the decomposition is based on an identity, the empirical version holds exactly, allowing us to attribute realized or ex-post performance accurately and exhaustively to active and passive components.

Finally, if asset returns are assumed to satisfy a linear K -factor model such as the Capital Asset Pricing Model (CAPM), the Arbitrage Pricing Theory (APT), or any other linear pricing model, the AP decomposition yields several additional insights. In particular, a portfolio's expected return can be decomposed into three distinct components when returns exhibit a linear factor structure: security selection, factor-timing ability, and risk premia. The first two components may be interpreted as the result of active management, and the last component is passive. This decomposition provides one explanation for the seemingly persistent differences between long-only and alternative investments—the long-only constraint imposes a limit to the amount of factor timing that can be accomplished, and this limit may be a severe handicap in environments where factor risk premia change sign, i.e., periods of

time-varying expected returns. A factor-based AP decomposition also addresses a recent concern of many institutional investors, that they are paying hedge funds for alpha but are getting beta instead. The relevant question is whether the beta exposure is time-varying or fixed—if it is the former, then it may be considered a genuine source of active value, but if it is the latter, it may be possible to achieve the same exposures in a more passive manner.

I begin in Section 2 with a brief review of the performance attribution literature, and present the main results of the paper in Section 3. I provide several analytical examples in Section 4, and then show how to implement the decomposition in Section 5. Section 6 contains a detailed empirical example of the AP decomposition applied to a statistical arbitrage strategy using daily data for NASDAQ size-deciles from January 2, 1990 to December 29, 1995. I conclude in Section 7.

2 Literature Review

The origins of performance attribution can be traced back to the Capital Asset Pricing Model of Sharpe (1964) and Lintner (1965), who derived a linear relation between the excess return of an investment and its systematic risk or market beta, i.e., the security market line:

$$R_{pt} - R_f = \beta_p(R_{mt} - R_f) + \epsilon_{pt} \quad , \quad E[\epsilon_t | R_{mt}] = 0 \quad . \quad (1)$$

Departures from this linear relation were generically termed “alpha”,

$$R_{pt} - R_f = \alpha_p + \beta_p(R_{mt} - R_f) + \epsilon_{pt} \quad (2)$$

and Treynor (1965), Sharpe (1966), and Jensen (1968, 1969) applied this measure to gauge the economic value-added of mutual-fund managers. Since then, a variety of related measures have been proposed, including

$$\frac{E[R_{pt}] - R_f}{\sigma_p} = \text{Sharpe Ratio} \quad (3)$$

$$\frac{E[R_{pt}] - R_f}{\beta_p} = \text{Treynor Ratio} \quad (4)$$

$$\frac{\alpha_p}{\sigma(\epsilon_{pt})} = \text{Information Ratio} \quad (5)$$

where σ_p and $\sigma(\epsilon_{pt})$ denote the standard deviations of R_{pt} and the residual ϵ_{pt} in (2), respectively. Graham and Harvey (1997) and Modigliani and Modigliani (1997) have derived risk-adjusted transformations of these basic measures, and Sharpe (1992) has proposed a constrained-regression framework for performing style analysis.

All of these measures are essentially static in nature because they are based on characteristics of the marginal distributions of returns at a single date t , e.g., means, variances, and contemporaneous covariances of the portfolio and market returns.¹ Even Sharpe’s (1992) regression-based decomposition—which is a conditional measure by construction—is static, because the conditioning information is contemporaneous to portfolio returns at date t . None of these measures captures the time-series dependence between dates t and $t+1$, which should be central focus of any forecasting measure of investment skill.

In contrast to these static measures, Treynor and Mazuy (1966) were among the first to propose a dynamic measure of active management. To pick up market-timing skills, Treynor and Mazuy (1966) augmented the linear framework of (1) with a quadratic term $(R_{mt} - R_f)^2$. Formally, this is still a contemporaneous regression, hence its parameters do not involve any time-series properties, but the motivation—to detect asymmetries in up and down markets—is distinctly dynamic.

In a recent comment on the fundamental indexation approach of Arnott, Hsu, and Moore (2005), Treynor (2005) does focus on the covariance between portfolio weights and returns, but only in the specific context of explaining the potential improvements of fundamental indexation over market-capitalization weights. The AP decomposition described in this paper provides a considerably more general framework for thinking about the benefits and costs of any indexation algorithm.

A more explicit measure of market-timing skill is given by Merton (1981), Henriksson and Merton (1981), and Henriksson (1984), who showed that perfect market-timing is equivalent to a buy-and-hold investment in the market that is fully protected by put options on the market with a riskless-rate strike. Imperfect market-timing skill may then be modeled as

¹Of course, data spanning many dates are used to estimate these parameters, but the parameters themselves are single-point statistics, as opposed to multi-point statistics such as autocorrelations. Moreover, the estimates for the standard performance metrics are almost always computed under the assumption that the data are independently and identically distributed, which rules out any non-random-walk behavior, the starting point of any active investment strategy.

a partially insured investment in the market. Although this measure may also seem like a static one because it involves parameters at a single point in time of the multivariate distribution of market returns and corresponding put-option prices, it is, in fact, dynamic due to the multi-period nature of options prices.² In other words, the dynamics of market-timing have been compressed into the put-option price, which is then used as the “numeraire” for market-timing skill in Merton’s (1981) framework.

Finally, Grinold and Kahn (2000, Chapter 17) describe a bottoms-up approach to performance attribution within a linear-factor model framework in which benchmark-timing is a component. Although they do not focus explicitly on the covariance between portfolio weights and returns, it is implicit in their time-series expression of “active benchmark timing return”,³ and in a footnote, they acknowledge that the same type of decomposition can be applied to individual securities.⁴

The measure of active management proposed in Section 3 is a more direct version of these last two notions of market timing, in which each portfolio weight is viewed as a bet on the future return of a given asset, and bets that yield positive profits over time, i.e., asset-timing ability, are indications of investment skill.

3 The AP Decomposition

Consider a portfolio P invested in n securities indexed by $i = 1, \dots, n$, and defined by its weights $\{\omega_{it}\}$ in those securities at date t . Denote by R_{pt} the portfolio’s return between dates $t-1$ and t , which is given by:

$$R_{pt} = \sum_{i=1}^n \omega_{it} R_{it} \tag{6}$$

where R_{it} is the date- t return of security i . Assuming that the means, variances, and covariances of individual securities’ returns are well-defined, the main result of the paper follows almost immediately from the definition of expected return, and is presented in Section 3.1.

²In particular, one of the option-pricing parameters is the maturity date, which renders the option price a multi-point statistic.

³Equation 17.27 of Grinold and Kahn (2000, p. 504) is proportional to the covariance between factor betas and returns, and is the sample counterpart to (17b) below.

⁴See footnote 9 of Grinold and Kahn (2000, p. 504). I am grateful to Lisa Goldberg for this citation.

This very general result implies a new way of defining passive and active investing, which is described in Section 3.2. And by imposing additional structure on asset returns—a linear factor structure, in particular—we show in Section 3.3 that our decomposition provides a new way of distinguishing between alpha and beta.

3.1 The General Result

We start with the following assumptions:

(A1) *The returns $\{R_{it}\}$ for each security i forms a stationary and ergodic stochastic process with finite moments up to order 4.*

(A2) *Date- t portfolio weights $\{\omega_{it}\}$ are stationary and ergodic stochastic processes that are functions of state variables \mathbf{X}_{t-1} .*

Assumption (A1) is standard, and implies that the means, variances, and covariances for all asset returns are well-defined, and that estimators of those parameters will have well-behaved limiting distributions with finite variances, which is useful for conducting statistical inferences.

Assumption (A2) requires more discussion. Beginning with Markowitz (1952), much of the investments literature has assumed that portfolio weights are non-stochastic. The motivation for this starting point is the focus on portfolio optimization, in which the portfolio weights are the choice variables with respect to which a mean-variance objective function is optimized. From this perspective, the expected return of a portfolio is simply a weighted average of the expected returns of each component security, weighted by the fraction of the portfolio devoted to that security.

But in practice, portfolio weights are not fixed parameters—they represent decisions taken by a portfolio manager, and, as such, depend on a number of inputs. Therefore, from an investor’s perspective, the statistical properties of a portfolio’s return is determined not just by the return distributions of the component securities, but also by the characteristics of the portfolio weights. In fact, the investor in an active investment product pays dearly for the services of the portfolio manager, and is paying largely for that manager’s portfolio weights! Accordingly, unless a manager’s portfolio weights depend on inputs that are non-stochastic

and known to us, and we have complete knowledge of the functional relationship between inputs and weights, we must view portfolio weights as random variables. And if portfolio weights are random variables, it is not surprising that the stochastic relation between weights and returns can have significant implications for the properties of their product, which we shall demonstrate shortly in Proposition 1 below.

Assumption (A2) provides a simple but important set of restrictions for the kind of randomness that a well-defined set of portfolio weights can exhibit. Perhaps the most important restriction is the fact that the date- t weights $\{\omega_t\}$ can depend on information only *prior* to date t , otherwise unlimited arbitrage opportunities will abound. To see why, suppose that we remove this restriction. Then consider the following portfolio:

$$\omega_{it} = \frac{\text{Max}[0, R_{it}]}{\sum_{j=1}^n \text{Max}[0, R_{jt}]} \quad , \quad R_{pt} = \sum_{i=1}^n \omega_{it} R_{it} \geq 0 \quad (7)$$

which places positive weight at date t on only those securities with positive return at date t . This is clearly an arbitrage if returns are not all degenerate. Assumption (A2) eliminates such possibilities, and is equivalent to ruling out “look-ahead” bias in the construction of portfolio weights.

Given these assumptions, we have the following general decomposition for a portfolio’s expected returns:

Proposition 1 *Under assumptions (A1)–(A2), the expected return of any portfolio P satisfies the following decomposition:*

$$\text{E}[R_{pt}] = \sum_{i=1}^n \text{E}[\omega_{it} R_{it}] = \sum_{i=1}^n \left(\text{Cov}[\omega_{it}, R_{it}] + \text{E}[\omega_{it}] \text{E}[R_{it}] \right) \quad (8)$$

$$= \sum_{i=1}^n \text{Cov}[\omega_{it}, R_{it}] + \sum_{i=1}^n \text{E}[\omega_{it}] \text{E}[R_{it}] \equiv \delta_p + \nu_p \quad (9)$$

$$\text{where } \delta_p \equiv \sum_{i=1}^n \text{Cov}[\omega_{it}, R_{it}] \quad (\text{Active Component})$$

$$\nu_p \equiv \sum_{i=1}^n \text{E}[\omega_{it}] \text{E}[R_{it}] \quad (\text{Passive Component})$$

$$\theta_p \equiv \frac{\delta_p}{\delta_p + \nu_p} \quad (\text{Active Ratio}) \quad (10)$$

Proposition 1 is a simple decomposition of a portfolio’s expected return into two components: the sum of the covariances between portfolio weights and returns, and the sum of the products of expected portfolio weights and expected returns. Although this active/passive or AP decomposition follows trivially from the definition of covariance, it has some surprisingly useful implications for identifying the relative contributions of active and passive investment management.

The motivation for attributing the covariance terms δ_p in (9) to active management is the fact that, by definition, “active” implies a conscious decision on the part of the portfolio manager to buy, sell, or avoid a security—summarized in the portfolio weights $\{\omega_{it}\}$ —and the impact of those decisions on a portfolio’s total expected return $E[R_{pt}]$ is captured by the covariances in (9). In particular, if a manager has positive weights when security returns are positive and negative weights when security returns are negative on average, this implies positive covariances between portfolio weights and returns, and will have a positive impact on the portfolio’s expected return. In effect, the covariance term captures the manager’s timing ability, asset by asset, and while perfect timing of the sort described in (7) is not possible, imperfect timing certainly is.

However, there is another source of positive expected return that has nothing to do with asset timing: the manager may be holding passive long positions in securities with positive expected returns, and passive short positions in securities with negative expected returns. For example, buy-and-hold investors in the S&P 500 should expect a positive expected return over time because of the equity risk premium. This passive source of expected return is captured by the second term ν_p in (9), which has no timing component whatsoever, but involves only the first moments of the marginal distributions of returns and weights. Similarly, a dedicated shortseller whose sole investment mandate is to identify and short over-valued companies faces a significant challenge from the negative equity risk premium implicit in any passive short equity position.

The distinction between timing and risk premia is more obvious from a slightly modified version of (9):

$$E[R_{pt}] = \sum_{i=1}^n \sigma(\omega_{it})\sigma(R_{it})\text{Corr}[\omega_{it}, R_{it}] + \sum_{i=1}^n E[\omega_{it}]E[R_{it}] . \quad (11)$$

If the portfolio weights $\{\omega_{it}\}$ are constant through time, it is difficult to argue that such a portfolio is being actively managed, and (11) supports this intuition: a constant portfolio implies zero variance for the portfolio weights, which implies no active component according to (11). Nevertheless, constant portfolios can still generate positive expected returns simply by holding securities with positive expected returns, as captured by the passive component. But (11) shows that the larger the correlation between weights and returns, and the larger the variances of returns and weights, the bigger the contribution from active management.

Alternatively, (11) implies that negative correlation between weights and returns detracts from a portfolio’s expected return, and higher return- and weight-volatility makes this worse. This suggests one possible approach to improving the active component of a portfolio strategy: decreasing (increasing) the volatility of weights for those securities where the correlations between weights and returns are negative (positive).

The quantity θ_p defined in (10)—which we call the “active ratio”—provides a useful summary statistic for the degree to which a portfolio is actively managed. Unlike the traditional static measures of active performance such as the Sharpe or information ratios (see Section 2), θ_p is dynamic, unit-free, and independent of the choice of benchmark. The dynamic nature of θ_p is clear from the decomposition (11)—the time series properties of returns affect θ_p directly, in contrast to static performance measures. The fact that θ_p is unit-free is obvious from a simple dimensional analysis of this ratio: both the numerator and denominator are defined as returns per unit time, hence the ratio is invariant to any unit of measurement. In other words, if θ_p is estimated to be 28%, then this is the fraction of the portfolio’s return due to active management, irrespective of whether we estimate θ_p with daily, weekly, or monthly data. Contrast this with the information ratio—if an investment has an IR of 0.28, it matters a great deal whether this is a daily, monthly, or annual estimate. Finally, θ_p requires no benchmark for its definition, although certainly it is trivial to compute θ_p relative to a benchmark by simply applying (9) to the excess returns of a portfolio. However, one of the most important implications of (9) is that active investment management is not simply adding value in excess of a passive benchmark, which can be done passively by taking on non-benchmark factor exposures in a multi-factor world. We shall return to this issue in Section 3.3 when we impose a linear factor structure on asset returns.

3.2 A New Definition of Passive Investing

The AP decomposition (9) implies that constant portfolios have no active component, but what about portfolios with time-varying weights that are traditionally considered to be passive? For example, a value-weighted portfolio of all securities is clearly a passive portfolio, yet the portfolio weights change according to the market values of the component securities hence $\sigma(\omega_{it})$ is generally positive for all securities in this case. Does this portfolio's expected return have a non-zero active component?

The answer is no, as long as the individual securities' returns are serially independent, i.e., as long as the Random Walk Hypothesis holds for all securities. To see why, let P_{it} and S_{it} denote the price and shares outstanding of asset i at date t , respectively, and observe that the value-weighted portfolio is given by:

$$\omega_{it} = \frac{P_{it-1}S_{it-1}}{\sum_{j=1}^n P_{jt-1}S_{jt-1}} = \frac{P_{it-2}(1 + R_{it-1})S_{it-1}}{\sum_{j=1}^n P_{jt-2}(1 + R_{jt-1})S_{jt-1}}. \quad (12)$$

which depends on returns at date $t-1$. If returns are serially independent, then the correlations $\text{Corr}[\omega_{it}, R_{it}]$ are all zero, hence the active component in (11) is zero despite the fact that the volatilities of the weights are non-zero.

If, on the other hand, returns are not serially independent, then it is possible for the AP decomposition to yield a non-zero active component to a buy-and-hold portfolio. This is a sensible outcome because the presence of serial correlation implies that past returns do contain forecast power for future returns; hence, a portfolio strategy with weights that change as a function of past returns may benefit from (or be hurt by) such serial correlation.

This suggests a broader but more precise definition for a passive portfolio:

Definition 1 *A passive portfolio is any portfolio with weights ω_{it} that are uncorrelated with its corresponding returns R_{it} for all $i = 1, \dots, n$.⁵*

Under this definition, a portfolio is passive if its weights do not contain any information related to future returns. Whether or not the portfolio is benchmarked is irrelevant. For example, consider a portfolio with a market beta of 1.00, achieved through S&P 500 futures

⁵We define constant-weight portfolios to be uncorrelated with returns since their covariances with returns are zero, even though in this case the correlation coefficient is undefined because the variance of the portfolio weight is zero.

contracts, that also happens to include a diversified buy-and-hold basket of commodities.⁶ The expected return of such a portfolio is likely to exceed the S&P 500 because of the positive risk premium associated with the commodities component, but should the excess return be attributed to active management? Using traditional performance measures, this portfolio is likely to exhibit positive “alpha”, but the AP decomposition will yield a very different conclusion.

The crucial characteristic of an active portfolio is the deliberate and successful use of information for forecasting returns—this is usually what we have in mind when we speak of “investment skill”. In the next section, we sharpen the distinction between alpha and other sources of expected return by assuming a linear factor structure for individual asset returns.

3.3 Alpha vs. Beta

To distinguish explicitly between alpha and beta, we have to impose additional structure on the return-generating processes for individual assets:

(A3) *For each asset i , the return R_{it} satisfies a linear K -factor model:⁷*

$$R_{it} = \alpha_i + \beta_{i1}F_{1t} + \cdots + \beta_{iK}F_{Kt} + \epsilon_{it} \quad (13)$$

$$0 = E[\epsilon_{it} | F_{1t}, \dots, F_{Kt}] \quad (14)$$

where the factors F_{kt} are stationary and ergodic stochastic processes.

Although the linear K -factor structure is presented as an assumption, several authors have derived theories to support such a specification, including Merton’s (1973) Intertemporal Capital Asset Pricing Model, Ross’s (1976) Arbitrage Pricing Theory, and Lo and Wang’s (2006) dynamic equilibrium model of returns and trading volume. However, in Assumption (A3), we leave room for the presence of an intercept α_i , which most other equilibrium asset-pricing models rule out because the presence of non-zero α_i may create arbitrage op-

⁶Commodities have historically exhibited very little correlation to the S&P 500, hence including a buy-and-hold portfolio of commodities is unlikely to affect the overall market beta of the portfolio.

⁷For notational simplicity, we omit the riskless rate R_f from this specification, but without any loss in generality, some readers may prefer to interpret both the asset return R_{it} and the factor returns F_{kt} as excess returns, in excess of R_f .

portunities.⁸ Since the AP decomposition is considerably more general than any particular asset-pricing model or linear factor structure, for our purposes, we remain agnostic about whether or not α_i is zero for all assets, and merely leave it as a possibility within our framework.

Under Assumption (A3), the return R_{pt} of an arbitrary portfolio of assets may be written as:

$$R_{pt} = \sum_{i=1}^n \omega_{it} R_{it} \quad (15a)$$

$$= \sum_{i=1}^n \omega_{it} \alpha_i + \left(\sum_{i=1}^n \omega_{it} \beta_{i1} \right) F_{1t} + \cdots + \left(\sum_{i=1}^n \omega_{it} \beta_{iK} \right) F_{Kt} + \sum_{i=1}^n \omega_{it} \epsilon_{it} \quad (15b)$$

$$= \alpha_{pt} + \sum_{k=1}^K \beta_{pk,t} F_{kt} + \epsilon_{pt} \quad (15c)$$

$$\text{where } \beta_{pk,t} \equiv \sum_{i=1}^n \omega_{it} \beta_{ik} \text{ , } \alpha_{pt} \equiv \sum_{i=1}^n \omega_{it} \alpha_i \text{ .} \quad (15d)$$

Armed with this return decomposition, the corresponding AP decomposition for expected returns follows immediately:

Proposition 2 *Under Assumptions (A1)–(A3), the expected return of any portfolio P satisfies the following decomposition:*

$$\begin{aligned} \mathbb{E}[R_{pt}] &= \sum_{i=1}^n \alpha_i \mathbb{E}[\omega_{it}] + \sum_{k=1}^K \text{Cov}[\beta_{pk,t}, F_{kt}] + \sum_{k=1}^K \mathbb{E}[\beta_{pk,t}] \mathbb{E}[F_{kt}] \\ &= \text{Security Selection} + \text{Factor Timing} + \text{Risk Premia} \end{aligned} \quad (16)$$

where

$$\text{Security Selection} \equiv \sum_{i=1}^n \alpha_i \mathbb{E}[\omega_{it}] \quad (17a)$$

$$\text{Factor Timing} \equiv \sum_{k=1}^K \text{Cov}[\beta_{pk,t}, F_{kt}] \quad (17b)$$

$$\text{Risk Premia} \equiv \sum_{k=1}^K \mathbb{E}[\beta_{pk,t}] \mathbb{E}[F_{kt}] \quad (17c)$$

⁸In particular, if the magnitudes of α_i are too large, or there are too many assets with non-zero α_i , then it may be possible to construct a portfolio that has no factor risks and no idiosyncratic risk but a non-zero intercept, which implies arbitrage.

and

$$\mathbb{E}[\beta_{pk,t}] = \sum_{i=1}^n \beta_{ik} \mathbb{E}[\omega_{it}] \quad , \quad k = 1, \dots, K \quad (18a)$$

$$\text{Cov}[\beta_{pk,t}, F_{kt}] = \sum_{i=1}^n \beta_{ik} \text{Cov}[\omega_{it}, F_{kt}] \quad , \quad k = 1, \dots, K \quad (18b)$$

Proposition 2 provides a more refined decomposition than Proposition 1, thanks to the linear K -factor structure assumed in (A3). Expected returns are now the sum of three components: a security-selection component (17a) that depends on the α_i 's, a factor-timing component (17b) that depends on the covariance between the portfolio betas and factors, and a risk-premia component (17c) that represents the expected return from passive exposures to factor risks. The first two components can be viewed as active sources of expected return, with the third being the passive component as in Proposition 1.

This factor-based AP decomposition clarifies several issues surrounding the measurement of active management. The first is the well-known implication of any multi-factor model that expected returns in excess of a benchmark need not be due to investment acumen, but may simply be passive exposures to non-market risk factors (see, for example, Merton, 1973 and Ross, 1976). For example, if a portfolio with unit market-beta has a passive credit-factor beta of 0.5 (in other words, no time-variation in this beta), then (17c) shows that this portfolio will yield half the risk premium of the credit factor in addition to the equity risk premium. Such outperformance can be achieved purely through passive means.

A second issue that Proposition 2 resolves is the historical disparity between the expected returns of long-only investments and hedge funds. Some have argued that the distinction is due to the outsize rewards from hedge-fund fee structures, which tend to draw the most talented portfolio managers from the long-only world, along with their alpha sources. Whether or not this exodus of talent from mutual funds to hedge funds is real is debatable, but (16) provides another explanation which is at least as compelling: the long-only constraint is a severe handicap from the perspective of expected returns, because it hampers a portfolio manager's ability to engage in factor timing. By allowing weights to be negative, a hedge-fund manager can create negative factor-betas for his portfolio, which will yield positive expected returns if the corresponding factor risk-premia are negative. Long-only managers

cannot easily do this if factor-betas are largely positive, hence one source of positive expected return is unavailable to them. Therefore, we should not expect long-only managers to be able to generate the level of expected returns that unconstrained managers can produce.

A third issue on which Proposition 2 sheds significant light is the recent concern of many institutional investors that they are paying hedge funds for alpha but are getting beta instead. Proposition 2 shows how to assess the value-added of hedge funds that make directional bets: check the covariances of their portfolio betas with the factors. If hedge funds possess factor-timing ability or “allocation alpha” (Leibowitz, 2005), this should be viewed as “alpha” or active management. But if the beta exposures are being generated passively, i.e., not with covariances (17b) but with relatively stable weights in (17c), then there are cheaper alternatives such as the “beta grazers” in Leibowitz (2005) or the futures-based hedge-fund beta replication strategies described in Hasanhodzic and Lo (2006). This suggests the following definition of passive investing in the context of a linear K -factor model:

Definition 2 *A portfolio is said to have passive factor exposure to factor k if its portfolio beta $\beta_{pk,t}$ is uncorrelated with the factor F_{kt} .*

Note that Definition 2 leaves open the possibility that a portfolio is passive with respect to one factor but active with respect to another. This is particularly important for the hedge-fund industry which is extraordinarily heterogeneous, and where highly specialized investment expertise can and should thrive.

Of course, passive beta exposures to non-traditional betas such as credit, liquidity, volatility, and yield-curve twists may be more difficult to generate and manage for a typical pension plan sponsor than S&P 500 beta, in which case higher management fees may be justifiable. But unless the manager is providing some form of active management—unique sources of α_i ’s or factor-timing ability—there should be downward pressure on the corresponding management fees, and certainly incentive fees cannot be justified.

One final insight provided by Proposition 2 is that investment expertise can manifest itself in two distinct formats: identifying untapped sources of expected return (the α_i ’s) and creating additional expected return through factor timing (the time-varying $\beta_{pk,t}$ ’s). Even if all α_i ’s are zero, as some academic studies claim, there can still be substantial value-added

from active management as long as risk premia vary over time and as functions of market conditions.

4 Some Analytical Examples

In this section I provide three examples of the AP decomposition where active and passive components can be evaluated analytically. Section 4.1 contains a simple numerical example that highlights the basic computations involved in the decomposition. In Section 4.2, I consider mean-reversion and momentum strategies, and in Section 4.3 I provide an analysis of a stop-loss strategy.


4.1 A Numerical Example

Consider a portfolio of two assets, one which yields a monthly return that alternates between 1% and 2% (Asset 1) and the other which yields a fixed monthly return of 0.15% (Asset 2). Let the weights of this portfolio, called A1, be given by 75% in Asset 1 and 25% in Asset 2. The chart and table in Figure 1 illustrates the dynamics of this portfolio over a 12-month period, where the expected return of the portfolio is 1.16% per month, none of which is due to the active component. Therefore $\theta_p = 0\%$ in this case.

Now consider portfolio A2, which differs from A1 only in that the portfolio weight for Asset 1 alternates between 50% and 100% as Asset 1's return alternatives between 1% and 2% (see the chart in Figure 2). In this case, the total expected return is 1.29% per month, of which 0.13% is due to the positive correlation between the portfolio weight for Asset 1 and its return. This yields an active ratio θ_p of 9.71%. Note that this ratio does not depend on the sampling interval of the returns, in contrast to the figures 1.29% and 0.13%, both of which are monthly returns.


Finally, consider a third portfolio A3 which also has alternating weights for Asset 1, but exactly in the opposite direction to Asset 1's returns: when the return is 1%, the portfolio weight is 100%, and when the return is 2%, the portfolio weight is 50%. This is obviously counterproductive, and Figure 3 confirms this intuition. Portfolio A3 loses 0.13% per month from its active component, and its active ratio $\theta_p = -12.05\%$.

Note that in all three cases, the passive components are identical at 1.16% per month




Month	ω_1	R_1	ω_2	R_2	R_p
Strategy A1					
1	75%	1.00%	25%	0.15%	0.79%
2	75%	2.00%	25%	0.15%	1.54%
3	75%	1.00%	25%	0.15%	0.79%
4	75%	2.00%	25%	0.15%	1.54%
5	75%	1.00%	25%	0.15%	0.79%
6	75%	2.00%	25%	0.15%	1.54%
7	75%	1.00%	25%	0.15%	0.79%
8	75%	2.00%	25%	0.15%	1.54%
9	75%	1.00%	25%	0.15%	0.79%
10	75%	2.00%	25%	0.15%	1.54%
11	75%	1.00%	25%	0.15%	0.79%
12	75%	2.00%	25%	0.15%	1.54%
Mean:	75.00%	1.50%	25.00%	0.15%	1.16%
SD:	0.00%	0.52%	0.00%	0.00%	0.39%
Corr[ω, R]:	0.00%		0.00%		
Cov[ω, R]:	0.00%	+	0.00%	=	0.00%
E[ω]E[R]:	1.13%	+	0.04%	=	1.16%
Total:					1.16%
	θ:				0.00%

Figure 1: The expected return of a constant portfolio does not contain any active component.



Month	ω_1	R_1	ω_2	R_2	R_p
Strategy A2					
1	50%	1.00%	50%	0.15%	0.58%
2	100%	2.00%	0%	0.15%	2.00%
3	50%	1.00%	50%	0.15%	0.58%
4	100%	2.00%	0%	0.15%	2.00%
5	50%	1.00%	50%	0.15%	0.58%
6	100%	2.00%	0%	0.15%	2.00%
7	50%	1.00%	50%	0.15%	0.58%
8	100%	2.00%	0%	0.15%	2.00%
9	50%	1.00%	50%	0.15%	0.58%
10	100%	2.00%	0%	0.15%	2.00%
11	50%	1.00%	50%	0.15%	0.58%
12	100%	2.00%	0%	0.15%	2.00%
Mean:	75.00%	1.50%	25.00%	0.15%	1.29%
SD:	26.11%	0.52%	26.11%	0.00%	0.74%
Corr[ω, R]:	100.00%		0.00%		
Cov[ω, R]:	0.13%	+	0.00%	=	0.13%
E[ω]E[R]:	1.13%	+	0.04%	=	1.16%
Total:					1.29%
	θ:				9.71%

Figure 2: The portfolio weights are positively correlated with returns, which adds value to the portfolio and yields a positive active component.



Month	ω_1	R_1	ω_2	R_2	R_p
Strategy A3					
1	100%	1.00%	0%	0.15%	1.00%
2	50%	2.00%	50%	0.15%	1.08%
3	100%	1.00%	0%	0.15%	1.00%
4	50%	2.00%	50%	0.15%	1.08%
5	100%	1.00%	0%	0.15%	1.00%
6	50%	2.00%	50%	0.15%	1.08%
7	100%	1.00%	0%	0.15%	1.00%
8	50%	2.00%	50%	0.15%	1.08%
9	100%	1.00%	0%	0.15%	1.00%
10	50%	2.00%	50%	0.15%	1.08%
11	100%	1.00%	0%	0.15%	1.00%
12	50%	2.00%	50%	0.15%	1.08%
Mean:	75.00%	1.50%	25.00%	0.15%	1.04%
SD:	26.11%	0.52%	26.11%	0.00%	0.04%
Corr[ω, R]:	-100.00%		0.00%		
Cov[ω, R]:	-0.13%	+	0.00%	=	-0.13%
E[ω]E[R]:	1.13%	+	0.04%	=	1.16%
Total:					1.04%
	θ:				-12.05%

Figure 3: The portfolio weights are negatively correlated with returns, which subtracts value to the portfolio and yields a negative active component.

because the average weight for each asset was the same across all three portfolios. The only differences among A1, A2, and A3 were the dynamics of the portfolio weights, and these differences gave rise to different values for the active component δ_p and the active ratio θ_p .

4.2 Mean Reversion and Momentum Strategies

Consider a simple asset-allocation example where one asset is the risk-free asset that yields R_f and the other is a risky asset with return R_t that satisfies a stationary autoregressive process with one lag or AR(1):

$$R_t = \mu + \rho(R_{t-1} - \mu) + \epsilon_t, \quad \epsilon_t \text{ IID WN}(0, \sigma_\epsilon^2) \quad (19)$$

$$\omega_t = \gamma_1 + \gamma_2 R_{t-1} \quad (20)$$

$$R_{pt} = \omega_t R_t + (1 - \omega_t) R_f \quad (21)$$

The asset-allocation strategy (20) is a simple linear function of last period's return plus a constant, and covers two important cases: a mean-reversion strategy ($\gamma_2 < 0$), and a momentum strategy ($\gamma_2 > 0$). The expected return of this strategy can be easily derived as:

$$E[R_{pt}] = \text{Cov}[\omega_t, R_t] + E[\omega_t]E[R_t] + (1 - E[\omega_t]) R_f \quad (22a)$$

$$= \gamma_2 \rho \text{Var}[R_t] + (\gamma_1 + \gamma_2 \mu) \mu + (1 - \gamma_1 - \gamma_2 \mu) R_f. \quad (22b)$$

This expression shows that the passive component is a weighted average of the two assets using the expected values of the weights ω_t and $1 - \omega_t$. On the other hand, the active component is a function of γ_2 , ρ , and the return variance of the risky asset. These three parameters represent the sum total of the information content in the strategy (20). If γ_2 is of the same sign as ρ , then the active component is positive—the strategy will be a momentum strategy when returns exhibit momentum ($\rho > 0$) and will be a mean-reversion strategy when returns exhibit mean reversion ($\rho < 0$). If, on the other hand, γ_2 is of opposite sign to ρ , then the active component will subtract value because the strategy will be exactly out of phase with the risky asset. And as the variance of R_t increases the active component becomes larger in absolute value, magnifying both positive and negative timing ability.

4.3 A Stop-Loss Policy

Suppose that a portfolio manager implements a stop-loss policy on an existing portfolio strategy with return R_t so that if the strategy's return falls below a threshold ζ at date $t-1$, the entire portfolio will be invested in the riskfree asset at date t , and if the strategy's return is greater than or equal to ζ , the entire portfolio will be invested in the strategy. Assuming that the strategy's return-generating process is an AR(1), we can specify the dynamics R_{pt} of the combined portfolio-plus-stop-loss strategy as:

$$R_t = \mu + \rho(R_{t-1} - \mu) + \epsilon_t, \quad \epsilon_t \text{ IID } \mathcal{N}(0, \sigma_\epsilon^2) \quad (23)$$

$$\omega_t = \begin{cases} 1 & \text{if } R_{t-1} > \zeta \\ 0 & \text{if } R_{t-1} \leq \zeta \end{cases} \quad (24)$$

$$R_{pt} = \omega_t R_t + (1 - \omega_t) R_f \quad (25)$$

where we have assumed that ϵ_t is Gaussian so that we can solve for expected returns in closed form. Under these specifications, the expected return of R_{pt} is given by:

$$\begin{aligned} \mathbb{E}[R_{pt}] &= \text{Cov}[\omega_t, R_t] + \mathbb{E}[\omega_t]\mathbb{E}[R_t] + (1 - \mathbb{E}[\omega_t]) R_f \\ &= \rho\sigma\phi\left(\frac{\alpha - \mu}{\sigma}\right) + \mu\left(1 - \Phi\left(\frac{\alpha - \mu}{\sigma}\right)\right) + R_f\Phi\left(\frac{\alpha - \mu}{\sigma}\right) \end{aligned} \quad (26)$$

The first term in (26) is the active component, and the second two terms are the passive component. As long as $\rho > 0$, implying momentum for the risky asset, the active component will increase the expected return of the portfolio. This result makes intuitive sense because the stop-loss policy (24) activates when past returns fall below ζ , and a lower past return implies lower future returns when $\rho \geq 0$. If, however, $\rho < 0$, then the stop-loss policy (24) becomes counterproductive because returns are mean-reverting, hence the time to switch to the riskless asset is after a very large positive return in the most recent past.

Table 1 provides a numerical illustration of the AP decomposition for various values of the two parameters ζ and ρ . When $\rho = -25\%$, the active component is negative regardless of the threshold ζ , although higher values of ζ make the active component more costly because more upside is eliminated by the stop-loss policy. When $\rho = 0$, the active component is 0 because returns follow a random walk in this case, and any portfolio strategy based on

α	ρ	Annualized					
		$E[R_{p,t}]$	Active	Passive	%Active	$E[\omega_t]$	$1-E[\omega_t]$
-1.0%	-25%	1.6%	-6.6%	8.1%	-423.1%	62.5%	37.5%
-0.5%	-25%	1.2%	-6.7%	8.0%	-547.6%	59.1%	40.9%
0.0%	-25%	0.9%	-6.8%	7.8%	-720.9%	55.7%	44.3%
0.5%	-25%	0.7%	-6.9%	7.6%	-962.5%	52.3%	47.7%
1.0%	-25%	0.5%	-6.9%	7.4%	-1290.0%	48.8%	51.2%
-1.0%	0%	8.1%	0.0%	8.1%	0.0%	62.5%	37.5%
-0.5%	0%	8.0%	0.0%	8.0%	0.0%	59.1%	40.9%
0.0%	0%	7.8%	0.0%	7.8%	0.0%	55.7%	44.3%
0.5%	0%	7.6%	0.0%	7.6%	0.0%	52.3%	47.7%
1.0%	0%	7.4%	0.0%	7.4%	0.0%	48.8%	51.2%
-1.0%	25%	14.7%	6.6%	8.1%	44.7%	62.5%	37.5%
-0.5%	25%	14.7%	6.7%	8.0%	45.8%	59.1%	40.9%
0.0%	25%	14.6%	6.8%	7.8%	46.8%	55.7%	44.3%
0.5%	25%	14.5%	6.9%	7.6%	47.5%	52.3%	47.7%
1.0%	25%	14.3%	6.9%	7.4%	48.1%	48.8%	51.2%
-1.0%	50%	21.3%	13.1%	8.1%	61.8%	62.5%	37.5%
-0.5%	50%	21.4%	13.5%	8.0%	62.8%	59.1%	40.9%
0.0%	50%	21.5%	13.7%	7.8%	63.7%	55.7%	44.3%
0.5%	50%	21.4%	13.8%	7.6%	64.4%	52.3%	47.7%
1.0%	50%	21.3%	13.8%	7.4%	65.0%	48.8%	51.2%

Table 1: Measures of the value of a stop-loss policy for a portfolio with monthly returns that follow an AR(1), and monthly parameter values $R_f = 5\%/12$, $E[R_t] = 10\%/12$, and $\text{Var}[R_t] = (20\%)^2/12$.

prior returns will be uncorrelated with current returns. And when ρ is positive, the active component becomes positive as well, and as predicted by (26), the larger the threshold and the larger the ρ , the more valuable is the stop-loss and the larger is the active ratio.

See Kaminski and Lo (2007) for a more detailed analysis of the performance of stop-loss policies.

5 Implementing the AP Decomposition

The AP decomposition in Section 3 is a simple function of means, variances, and covariances of portfolio weights and returns, hence its implementation is a straightforward exercise in estimating unconditional first and second moments. In fact, in Section 5.1 we show that the properly defined sample moments of portfolio weights and returns are related in exactly the same way as their population counterparts in (9), hence the decomposition must hold exactly as an identity when applied to the data. However, there is one subtlety in implementing (9) that is not addressed by the population version, which is the selection of the sampling interval, and this is considered in Section 5.3.

5.1 Population vs. Sample Moments

Assumptions (A1) and (A2) are generally sufficient to ensure that the usual sample means, variances, and covariances of portfolio weights and returns will be well-behaved estimators for their population values (see, for example, White, 1984). Therefore, implementing the fundamental decomposition (9) involves nothing more challenging than estimating first and second moments. Moreover, the fact that sample moments are related to each other in the same way that population moments are related makes the application of (9) to the data almost trivial. For example, consider the sample covariance between portfolio weights and returns for asset i :

$$\frac{1}{T} \sum_{t=1}^T (\omega_{it} - \bar{\omega}_i)(R_{it} - \bar{R}_i) = \frac{1}{T} \sum_{t=1}^T \left(\omega_{it} R_{it} - \bar{\omega}_i R_{it} - \omega_{it} \bar{R}_i + \bar{\omega}_i \bar{R}_i \right) \quad (27a)$$

$$= \frac{1}{T} \sum_{t=1}^T \omega_{it} R_{it} - \frac{\bar{\omega}_i}{T} \sum_{t=1}^T R_{it} - \frac{\bar{R}_i}{T} \sum_{t=1}^T \omega_{it} + \bar{\omega}_i \bar{R}_i \quad (27b)$$

$$= \frac{1}{T} \sum_{t=1}^T \omega_{it} R_{it} - \bar{\omega}_i \bar{R}_i . \quad (27c)$$

Therefore,

$$\frac{1}{T} \sum_{t=1}^T \omega_{it} R_{it} = \frac{1}{T} \sum_{t=1}^T (\omega_{it} - \bar{\omega}_i)(R_{it} - \bar{R}_i) + \bar{\omega}_i \bar{R}_i \quad (28)$$

which is the sample-moment counterpart to (9) for a single asset i . Repeated applications of (27) to the sample average of a portfolio will result in the sample version of (9), hence the AP decomposition must also hold in-sample. With properly defined first and second moments, (9) is an identity for sample moments as well as for population moments.

However, there is an even simpler approach to estimating the AP decomposition that eliminates the need for second moments altogether, and that is to infer the active component by subtracting the passive component from the total expected return:

$$\delta_p = E[R_{pt}] - \nu_p . \quad (29)$$

Because the right side of (29) involves only first moments, the active component δ_p and the active ratio θ_p may be computed solely from the average weights and average returns of the portfolio, hence:

$$\hat{\delta}_p = \frac{1}{T} \sum_{t=1}^T R_{pt} - \sum_{i=1}^n \bar{\omega}_i \bar{R}_i \quad (30)$$

$$\text{where } \bar{\omega}_i \equiv \frac{1}{T} \sum_{t=1}^T \omega_{it} , \quad \bar{R}_i \equiv \frac{1}{T} \sum_{t=1}^T R_{it}$$

$$\hat{\theta}_p = \frac{\hat{\delta}_p}{\frac{1}{T} \sum_{t=1}^T R_{pt}} . \quad (31)$$

In particular, it is not necessary for a manager to provide position-level transparency to give investors a clear sense of the value of his investment process. This is a particularly important characteristic of the fundamental decomposition (9) when applied to hedge-fund strategies because of secrecy surrounding such strategies. Few hedge-fund managers should balk at divulging average weights and returns, unless, of course, the manager's average weights do not vary much over time, in which case the investors should reconsider paying hedge-fund fees to such a manager.

5.2 GMM Estimation

The simplicity of the estimators (30) and (31) may suggest that statistical inference is not needed, but there is still the issue of assessing the accuracy of these estimators to address. Fortunately, given Assumptions (A1) and (A2), it is possible to derive the asymptotic distributions of (30) and (31) under fairly general conditions. These results are summarized in the following proposition (see the Appendix for a proof):

Proposition 3 *Under (A1) and (A2), the active component δ_p and active ratio θ_p of any portfolio P may be estimated consistently by their sample counterparts (30) and (31), and both estimators are asymptotically normal with variances that may be consistently estimated via the Generalized Method of Moments.*

5.3 The Sampling Interval

One important practical issue surrounding the implementation of the AP decomposition is the choice of sampling interval for weights and returns. This obviously does not affect the population version of the decomposition, but is of critical importance in practical applications. However, even for the population version, the relevance of the sampling interval can be understood in the context of time aggregation. Consider a portfolio strategy in which the weights $\{\omega_{it}(\mathbf{X}_{t-1})\}$ vary each period, but we only observe portfolio weights and returns every q periods. If we now apply the AP decomposition to this subset of weights and returns, would we arrive at the same values for δ_p and ν_p , appropriately scaled? The answer is no. To see why, we must first define the precise set of observables at each q -period interval, and then ask whether these observables satisfy the same relation (9) as their one-period counterparts.

First consider aggregating the portfolio return R_{pt} over q periods and ignore the effects of compounding for simplicity:

$$R_{pt} = \sum_{i=1}^n \omega_{it} R_{it} \quad (32)$$

$$R_{p\tau}(q) = \sum_{t=(\tau-1)q+1}^{\tau q} R_{pt} = \sum_{t=(\tau-1)q+1}^{\tau q} \sum_{i=1}^n \omega_{it} R_{it} \quad (33)$$

$$= \sum_{i=1}^n \sum_{t=(\tau-1)q+1}^{\tau q} \omega_{it} R_{it} . \quad (34)$$

Then the expected value of the q -period portfolio return is given by:

$$\mathbb{E}[R_{p\tau}(q)] = \sum_{i=1}^n \sum_{t=(\tau-1)q+1}^{\tau q} \mathbb{E}[\omega_{it} R_{it}] \quad (35)$$

$$= \sum_{i=1}^n q \mathbb{E}[\omega_{it} R_{it}] \quad \text{from stationarity} \quad (36)$$

$$= \sum_{i=1}^n q \text{Cov}[\omega_{it}, R_{it}] + \sum_{i=1}^n q \mathbb{E}[\omega_{it}] \mathbb{E}[R_{it}]. \quad (37)$$

The equality (37) shows that (9) aggregates linearly over time, so that the q -period expected return of a portfolio is equal to the sum of q times the active component and q times the passive component.

However, the question we started with is whether (9) remains the same with q -period inputs, and (37) has q -period returns only on the left side. To address our original problem, let us rewrite the right side of (33) with q -period variables as well:

$$R_{p\tau}(q) = \sum_{i=1}^n \sum_{t=(\tau-1)q+1}^{\tau q} \omega_{it} \frac{R_{it}}{R_{i\tau}(q)} R_{i\tau}(q) \quad (38)$$

$$= \sum_{i=1}^n \omega_{i\tau}(q) R_{i\tau}(q) \quad \text{where } \omega_{i\tau}(q) \equiv \sum_{t=(\tau-1)q+1}^{\tau q} \omega_{it} \frac{R_{it}}{R_{i\tau}(q)}. \quad (39)$$

In (39), the q -period portfolio return $R_{p\tau}(q)$ is shown to be equal to the product of q -period asset returns $R_{i\tau}(q)$ and q -period weights $\omega_{i\tau}(q)$, but note that these weights are not simply averages of ω_{it} over q periods, but are return-weighted averages. In fact, the only set of q -period portfolio weights that, when multiplied by q -period asset returns, will yield the same q -period portfolio return as the time-aggregated portfolio return $R_{p\tau}(q)$ is the return-weighted average defined in (39).

Now we are prepared to answer the question we first posed regarding the effect of time aggregation on (9). Suppose we apply the decomposition to q -period weights and returns—will we arrive at a scaled version of the one-period decomposition? The q -period version is given by:

$$\mathbb{E}[R_{p\tau}(q)] = \sum_{i=1}^n \mathbb{E}[\omega_{i\tau}(q) R_{i\tau}(q)] \quad (40)$$

$$= \sum_{i=1}^n \text{Cov}[\omega_{i\tau}(q), R_{i\tau}(q)] + \sum_{i=1}^n \text{E}[\omega_{i\tau}(q)]\text{E}[R_{i\tau}(q)] \quad (41)$$

and from (41) it is easy to see that the q -period version does not reduce to q times the one-period version as in (39). In particular, the covariances of q -period weights and returns involve complex cross-products of one-period weights and returns within the q -period interval, and even the expected value of q -period weights involve co-moments between weights and returns across the interval. In fact, much of the complexity of (41) is an artefact of the time aggregation, and has little to do with the underlying investment process.

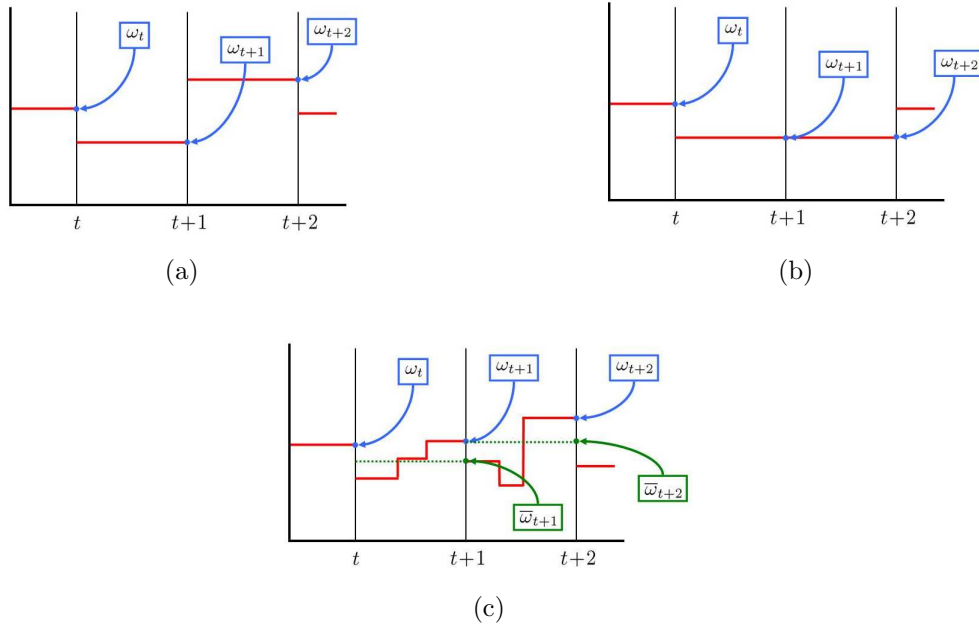


Figure 4: Illustration of all possible configurations of sampling and decision intervals: (a) matched sampling and decision intervals; (b) more frequent sampling intervals; (c) less frequent sampling intervals.

For this reason, it is critical that the sampling interval be at least as fine as the finest decision interval of the investment process (see Figures 4a and b). If portfolio decisions are made on a weekly basis, then (9) must be implemented with data sampled at least once a week. If, on the other hand, the strategy employs daily information, then weekly data will not suffice (see Figure 4c).

Sampling more finely than the finest decision interval is always acceptable because an unchanged portfolio weight may always be viewed as an active decision to maintain a position,

i.e., “let it ride”. Although Warren Buffett changes his portfolio weights rather slowly over time, this is quite deliberate and not due to lethargy or inertia. However, one useful implication of this longer decision interval is that we are unlikely to be misled by applying our decomposition to the monthly returns and weights of his portfolio, whereas the same cannot be said for the legendary day-trader Steven Cohen of SAC Capital.


Moreover, random shifts in portfolio weights in between decision intervals may add noise to the estimates (30) and (31), but will not affect the consistency of the estimators since, by definition, the noise is independent of returns (if they are correlated with returns, then they are considered part of the information in portfolio weights and will count towards the active component). Therefore, more finely sampled data is always preferred, but at a minimum, the data must be sampled at least as finely as the investment decision interval, otherwise the AP decomposition may not accurately reflect the true active and passive contributions of the manager.

6 An Empirical Application

To develop a better understanding of the characteristics of the AP decomposition, we apply our framework to a specific market-neutral equity trading strategy that, by construction, is particularly dynamic. The strategy, first proposed by Lo and MacKinlay (1990), is an example of a “statistical arbitrage” program, consisting of buying losers and selling winners in proportion to their under- or over-performance relative to the cross-section mean. Specifically, let:

$$\omega_{it} = -\frac{1}{n}(R_{it-1} - \bar{R}_{t-1}) \quad , \quad \bar{R}_{t-1} \equiv \frac{1}{n} \sum_{i=1}^n R_{it-1} \quad (42)$$

which implies that the weights sum to 0 at each date t , hence these are arbitrage portfolios. The fact that these weights are so directly tied to returns implies very active trading, and the fact that the weights sum to 0 implies very little market-beta exposure. These two implications suggest that much of this portfolio’s return should be due to active management, and that θ should be quite large.



Month	ω_1	R_1	ω_2	R_2	R_p
Strategy A4					
1	50%	1.00%	50%	0.15%	0.58%
2	50%	2.00%	50%	0.15%	1.08%
3	50%	1.00%	50%	0.15%	0.58%
4	50%	2.00%	50%	0.15%	1.08%
5	50%	1.00%	50%	0.15%	0.58%
6	50%	2.00%	50%	0.15%	1.08%
7	50%	1.00%	50%	0.15%	0.58%
8	100%	2.00%	0%	0.15%	2.00%
9	100%	1.00%	0%	0.15%	1.00%
10	100%	2.00%	0%	0.15%	2.00%
11	100%	1.00%	0%	0.15%	1.00%
12	100%	2.00%	0%	0.15%	2.00%
Mean:	70.83%	1.50%	29.17%	0.15%	1.13%
SD:	25.75%	0.52%	25.75%	0.00%	0.57%
Corr[ω, R]:	16.90%		0.00%		
Cov[ω, R]:	0.02%	+	0.00%	=	0.02%
E[ω]E[R]:	1.06%	+	0.04%	=	1.11%
Total:					1.13%
	θ:				1.85%

Figure 5: Sampling more frequently than the decision interval does not create any biases in the active ratio estimator.

6.1 Trading NASDAQ Size-Deciles

We apply this strategy to the daily returns of the five smallest size-decile portfolios of all NASDAQ stocks, as constructed by the University of Chicago’s Center for Research in Security Prices (CRSP), from January 2, 1990 to December 29, 1995.⁹ Table 2 reports the cross-autocorrelations of these five decile returns, and the lead/lag pattern that emerges from this table—the fact that larger-decile returns today are more highly correlated with smaller-decile returns tomorrow than smaller-decile returns today with larger-decile returns tomorrow—underscores the likely profitability of the contrarian strategy (42). Of course, trading NASDAQ size-deciles is obviously unrealistic in practice, but our purpose is to illustrate the performance of our expected-return decomposition (9), not to derive an implementable trading strategy.

	\mathbf{R}_{1t+1}	\mathbf{R}_{2t+1}	\mathbf{R}_{3t+1}	\mathbf{R}_{4t+1}	\mathbf{R}_{5t+1}
\mathbf{R}_{1t}	10.0%	21.5%	15.8%	18.1%	16.7%
\mathbf{R}_{2t}	23.4%	15.4%	20.2%	19.7%	15.8%
\mathbf{R}_{3t}	26.2%	25.0%	15.2%	23.9%	21.6%
\mathbf{R}_{4t}	25.4%	27.0%	24.3%	18.2%	18.7%
\mathbf{R}_{5t}	25.4%	26.6%	26.5%	26.2%	19.4%

Table 2: First-order cross-autocorrelation matrix of the daily returns of the five smallest CRSP-NASDAQ size deciles, from January 2, 1990 to December 29, 1995. Decile 1 is the smallest-capitalization decile.

Figure 6 illustrates the remarkable performance of the contrarian strategy (42) over the 1990 to 1995 sample period, and Table 3 contains summary statistics for the daily returns of the five deciles and the strategy. With an annualized average return of 31.4% and standard deviation of 7.9%, this strategy’s performance is considerably better than that of any of the five deciles, which is one indication that active management is playing a significant role in this case.

This intuition is confirmed by the AP decomposition of the strategy’s expected return into active and risk-premia components in Table 4. On an annualized basis, the active component yields 32% which exceeds the strategy’s total expected return of 31.4%, implying

⁹We selected this time period purposely because of the emergence of day-trading in the early 1990’s, an important source of profitability for statistical arbitrage strategies.




Figure 6: Cumulative return of a daily mean-reversion strategy of Lo and MacKinlay (1990) applied to the five smallest CRSP-NASDAQ size deciles from January 2, 1990 to December 29, 1995.

an active ratio of 101.9% and a slightly negative risk-premia component. In this case, more than all of the strategy’s expected return is coming from active management, and the risk-premia component is subtracting value. The explanation for this rather unusual phenomenon was provided by Lo and MacKinlay (1990, Section 5.3.1), who observed that because the contrarian strategy holds losers and shorts winners, on average it will be buying the low-mean assets and shorting the high-mean assets. Therefore, the risk-premia component—which is the sum of average portfolio weights multiplied by average returns—will consist of positive average weights for low-mean stocks and negative average weights for high-mean stocks for this strategy. Fortunately, the positive correlation between weights and returns is more than sufficient to compensate for the negative risk-premia, as Table 4 confirms.

6.2 Daily vs. Monthly Transparency

To illustrate the importance of the sampling interval in the measurement of the active ratio, suppose that we are given only the following monthly information for this daily strategy: month-end portfolio weights, and the monthly returns of each of the five deciles. Table 5 shows that the strategy’s expected return cannot be correctly computed from this information

Statistic	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	R_{pt}
Mean*250	27.4%	17.5%	14.0%	13.7%	12.8%	31.4%
SD*sqrt(250)	12.2%	9.8%	8.9%	9.1%	9.5%	7.9%
SR*sqrt(250)	2.25	1.78	1.58	1.50	1.35	3.95
Min	-2.9%	-2.7%	-2.7%	-3.3%	-3.5%	-2.2%
Median	0.1%	0.1%	0.1%	0.1%	0.1%	0.1%
Max	6.7%	3.6%	2.0%	2.1%	2.3%	2.4%
Skew	0.6	0.0	-0.5	-0.7	-0.9	-0.1
XSKurt	5.1	2.4	2.1	3.1	3.9	1.7
ρ_1	10.0%	15.4%	15.2%	18.2%	19.4%	4.7%
ρ_2	10.3%	7.7%	10.1%	13.9%	10.5%	0.9%
ρ_3	5.7%	4.2%	7.5%	9.2%	11.0%	7.5%

Table 3: Summary statistics of the daily returns of the five smallest CRSP-NASDAQ size deciles, and the daily returns of a mean-reversion strategy of Lo and MacKinlay (1990) applied to those decile returns, from January 2, 1990 to December 29, 1995. Decile 1 is the smallest-capitalization decile.

Statistic	Estimate	SE	t-stat
Daily			
Portfolio Mean * 250	31.4%	0.3%	91.00
Risk Premia * 250	-0.6%	3.5%	-0.17
Active Component * 250	32.0%	3.5%	9.24
Active Ratio	101.9%	0.3%	354.40
Monthly			
Portfolio Mean * 12	-4.0%	1.0%	-3.98
Risk Premia * 12	0.1%	4.0%	0.03
Active Component * 12	-4.1%	4.1%	-1.01
Active Ratio	102.6%	11.8%	8.66

Table 4: GMM estimates of the active ratio θ (in percent) of the daily and monthly returns of a mean-reversion strategy of Lo and MacKinlay (1990) applied to the five smallest CRSP-NASDAQ size-decile returns, from January 2, 1990 to December 29, 1995. Daily estimates employ a truncation lag of 6 in the computation of the Newey-West asymptotic standard errors, and monthly estimates employ a truncation lag of 3.

because much of the return comes from daily bets that are not captured by month-end portfolio weights. The decomposition of the incorrect expected return of -4.0% is reported in the lower panel of Table 4 and also has no bearing to the daily strategy—the estimates imply that the active component yields -4.1% , which is clearly false. By observing only month-end weights and cumulative returns, we have no way of inferring the profitability of decisions made at a daily frequency.

Statistic	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	R_{pt}
Mean*12	27.5%	17.4%	13.9%	13.7%	12.8%	-4.0%
SD*sqrt(12)	20.6%	17.7%	15.6%	15.0%	15.9%	8.8%
SR*sqrt(12)	1.34	0.98	0.89	0.91	0.80	-0.45
Min	-8.0%	-11.3%	-9.0%	-9.7%	-11.4%	-6.6%
Median	1.6%	1.0%	1.2%	1.0%	1.5%	-0.5%
Max	26.4%	21.5%	18.1%	16.8%	16.2%	9.7%
Skew	1.2	0.9	0.7	0.5	0.1	1.2
XSKurt	3.3	3.0	2.1	2.0	1.2	4.7
ρ_1	36.4%	43.7%	43.2%	41.4%	45.0%	8.2%
ρ_2	17.3%	16.6%	18.9%	10.1%	13.7%	15.7%
ρ_3	-5.6%	-2.7%	-3.1%	-7.8%	-7.0%	-3.2%

Table 5: Summary statistics of the monthly returns of the five smallest CRSP-NASDAQ size deciles, and the monthly returns of a mean-reversion strategy of Lo and MacKinlay (1990) applied to those decile returns, from January 1990 to December 1995. Decile 1 is the smallest-capitalization decile.

7 Conclusion

The investment management industry has developed a series of measures to gauge the performance of portfolio managers. Most of these measures are based on characteristics of the marginal distributions of asset and portfolio returns at a single point in time. Also, the typical statistical procedures used to estimate these characteristics implicitly assume that the data are independently and identically distributed, which eliminates any possibility of dynamic effects such as intertemporal correlations and forecast power.

In contrast to these static measures, the AP decomposition and active ratio—which are based on the definition of covariance—are multi-point statistics that capture the very essence

of active management: time-series predictability. A successful portfolio manager is one whose decisions are more often right than wrong, but “right” and “wrong” have specific meanings in this context: positive correlation between portfolio weights and returns. Given that portfolio weights are functions of a manager’s prior and proprietary information, positive correlation between weights and returns is a clear indication of forecast power and, consequently, active investment skills.

This definition of active management provides a natural dichotomy between active and passive investing. A passive portfolio has no forecast power, and can therefore be implemented more easily than an active portfolio. This is a sensible generalization of the standard definition of a passive portfolio, i.e., a constant mix of liquid benchmarks.

If asset returns are assumed to satisfy a linear multi-factor model, the AP decomposition shows that a portfolio’s expected returns can be decomposed into three components: deviations from the linear factor model or “alpha”, positive correlation between portfolio betas and factor realizations or “factor timing”, and static factor exposures times risk premia. This decomposition provides a clear and simple framework for resolving the question of whether hedge-fund investors are paying for alpha and getting beta from their investments. Moreover, the AP decomposition provides one explanation for the historical differences in performance between long-only and long/short portfolios: the long-only restriction severely limits a manager’s factor-timing ability, which can be a substantial performance drag during periods where risk premia change sign.

Finally, because the AP decomposition is an identity in any given sample of data, it can be used to conduct detailed performance attributions, factor by factor, and asset by asset. By separating the active and passive components of a portfolio, it should be possible to study and improve the performance of both.

A Appendix

Throughout the Appendix, the following conventions are maintained: (1) all vectors are column vectors unless otherwise indicated; (2) vectors and matrices are always typeset in boldface, i.e., X and μ are scalars and \mathbf{X} and $\boldsymbol{\mu}$ are vectors or matrices.

A.1 GMM Estimators for the AP Decomposition

Denote by \mathbf{X}_t the vector of period- t portfolio weights and returns $[\omega_{1t} \cdots \omega_{nt} R_{1t} \cdots R_{nt}]'$, and let $\{\mathbf{X}_t\}$ be a stochastic process that satisfies the following conditions:

- (H1) $\{\mathbf{X}_t : t \in (-\infty, \infty)\}$ is stationary and ergodic.
- (H2) $\gamma_o \in \Gamma$ is an open subset of \mathfrak{R}^k .
- (H3) $\forall \gamma \in \Gamma$, $\varphi(\cdot, \gamma)$ and $\varphi_\gamma(\cdot, \gamma)$ are Borel-measurable and $\varphi_\gamma(\mathbf{X}, \cdot)$ is continuous on Γ for all \mathbf{X} .
- (H4) φ_γ is first-moment continuous at γ_o , $E[\varphi_\gamma(\mathbf{X}, \cdot)]$ exists, is finite, and is of full rank.
- (H5) Let $\varphi_t \equiv \varphi(\mathbf{X}_t, \gamma_o)$ and

$$\mathbf{v}_j \equiv E[\varphi_0 | \varphi_{-1}, \varphi_{-2}, \dots] - E[\varphi_0 | \varphi_{-j-1}, \varphi_{-j-2}, \dots]$$

and assume:

- (i) $E[\varphi_0 \varphi_0']$ exists and is finite
- (ii) \mathbf{v}_j converges in mean square to 0
- (iii) $\sum_{j=0}^{\infty} E[\mathbf{v}_j' \mathbf{v}_j]^{1/2}$ is finite

which implies $E[\varphi(\mathbf{X}_t, \gamma_o)] = 0$.

- (H6) Let $\hat{\gamma}$ solve $\frac{1}{T} \sum_{t=1}^T \varphi(\mathbf{X}_t, \gamma) = 0$.

Then Hansen (1982) shows that:

$$\sqrt{T}(\hat{\gamma} - \gamma_o) \stackrel{a}{\sim} \mathcal{N}(0, \mathbf{V}_\gamma) \quad , \quad \mathbf{V}_\gamma \equiv \mathbf{H}^{-1} \boldsymbol{\Sigma} \mathbf{H}^{-1'} \quad (\text{A.1})$$

where

$$\mathbf{H} \equiv \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \varphi_{\gamma}(\mathbf{X}_t, \gamma_o) \right] \quad (\text{A.2})$$

$$\mathbf{\Sigma} \equiv \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \sum_{s=1}^T \varphi(\mathbf{X}_t, \gamma_o) \varphi(\mathbf{X}_s, \gamma_o)' \right]. \quad (\text{A.3})$$

and $\varphi_{\gamma}(R_t, \gamma)$ denotes the derivative of $\varphi(R_t, \gamma)$ with respect to γ .¹⁰ Specifically, let $\varphi(R_t, \gamma)$ denote the following vector function:

$$\varphi(\omega_t, \mathbf{R}_t, \gamma) \equiv [\omega_{1t} - \mu_{\omega_1} \cdots \omega_{nt} - \mu_{\omega_n} \quad R_{1t} - \mu_1 \cdots R_{nt} - \mu_n]'. \quad (\text{A.4})$$

where $\gamma \equiv [\mu_{\omega_1} \cdots \mu_{\omega_n} \quad \mu_1 \cdots \mu_n]'$. The GMM estimator of γ , denoted by $\hat{\gamma}$, is given implicitly by the solution to:

$$\frac{1}{T} \sum_{t=1}^T \varphi(\omega_t, \mathbf{R}_t, \gamma) = 0 \quad (\text{A.5})$$

which yields the standard estimators $\bar{\omega}_i$ and \bar{R}_i given in (30). For the moment conditions in (A.4), \mathbf{H} is given by:

$$\mathbf{H} = -\mathbf{I} \quad (\text{A.6})$$

where \mathbf{I} is a $(2n \times 2n)$ identity matrix. Therefore, $\mathbf{V}_{\gamma} = \mathbf{\Sigma}$ and the asymptotic distribution of the active ratio estimator $\hat{\theta}$ follows from the delta method:

$$\sqrt{T} \left(\hat{\theta} - \theta \right) \overset{a}{\sim} \mathcal{N}(0, V_{\text{GMM}}) \quad , \quad V_{\text{GMM}} = \frac{\partial g}{\partial \gamma} \mathbf{\Sigma} \frac{\partial g}{\partial \gamma'} \quad (\text{A.7})$$

where $\partial g / \partial \gamma$ is given by:

$$\frac{\partial g}{\partial \gamma} = [\mu_1 \cdots \mu_n \quad \mu_{\omega_1} \cdots \mu_{\omega_n}]. \quad (\text{A.8})$$

¹⁰See Magnus and Neudecker (1988) for the specific definitions and conventions of vector and matrix derivatives of vector functions.

An estimator for $\partial g/\partial \boldsymbol{\gamma}$ may be obtained by substituting $\hat{\boldsymbol{\gamma}}$ into (A.8), and an estimator for $\boldsymbol{\Sigma}$ may be obtained using Newey and West's (1987) procedure:

$$\hat{\boldsymbol{\Sigma}} = \hat{\boldsymbol{\Omega}}_0 + \sum_{j=1}^m \omega(j, m) (\hat{\boldsymbol{\Omega}}_j + \hat{\boldsymbol{\Omega}}_j'), \quad m \ll T \quad (\text{A.9})$$

$$\hat{\boldsymbol{\Omega}}_j \equiv \frac{1}{T} \sum_{t=j+1}^T \boldsymbol{\varphi}(R_t, \hat{\boldsymbol{\gamma}}) \boldsymbol{\varphi}(R_{t-j}, \hat{\boldsymbol{\gamma}})' \quad (\text{A.10})$$

$$\omega(j, m) \equiv 1 - \frac{j}{m+1} \quad (\text{A.11})$$

and m is the truncation lag, which must satisfy the condition $m/T \rightarrow \infty$ as T increases without bound to ensure consistency. An estimator for V_{GMM} can then be constructed as:

$$\hat{V}_{\text{GMM}} = \frac{\partial g(\hat{\boldsymbol{\gamma}})}{\partial \boldsymbol{\gamma}} \hat{\boldsymbol{\Sigma}} \frac{\partial g(\hat{\boldsymbol{\gamma}})}{\partial \boldsymbol{\gamma}'}. \quad (\text{A.12})$$

References

- Arnott, R., Hsu, J. and P. Moore, 2005 “Fundamental Indexation”, *Financial Analysts Journal* 61, 83–99.
- Campbell, J., Lo, A., and C. MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.
- Graham, J. and C. Harvey, 1997, “Grading the Performance of Market Timing Newsletters”, *Financial Analysts Journal* 53, 54–66.
- Grinold, R. and R. Kahn, 2000, *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*. New York: McGraw-Hill.
- Hansen, L., 1982, “Large Sample Properties of Generalized Method of Moments Estimators”, *Econometrica* 50, 1029–1054.
- Harvey, A., 1981, *Time Series Models*. New York: John Wiley & Sons.
- Hasanhodzic, J. and A. Lo, 2006, “Can Hedge-Fund Returns Be Replicated?: The Linear Case”, to appear in *Journal of Investment Management*.
- Henriksson, R., 1984, “Market Timing and Mutual Fund Performance: An Empirical Investigation”, *Journal of Business* 57, 73–96.
- Henriksson, R. and R. Merton, 1981, “On Market Timing and Investment Performance II: Statistical Procedures for Evaluating Forecast Skills”, *Journal of Business* 54, 513–533.
- Jensen, M., 1968, “The Performance of Mutual Funds in the Period 1945–1964”, *Journal of Finance* 23, 389–416.
- Jensen, M., 1969, “Risk, the Pricing of Capital Assets, and the Evaluation of Investment Performance”, *Journal of Business* 42, 167–247.
- Kaminski, K. and A. Lo, 2007, “When Do Stop-Loss Rules Stop Losses?”, unpublished working paper, MIT Laboratory for Financial Engineering.
- Leibowitz, M., 2005, “Alpha Hunters and Beta Grazers”, *Financial Analysts Journal* 61, 32–39.
- Lintner, J., 1965, “The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets”, *Review of Economics and Statistics* 47, 13–37.
- Lo, A., 2002, “The Statistics of Sharpe Ratios”, *Financial Analysts Journal* 58, 36–52.
- Lo, A. and C. MacKinlay, 1999, *A Non-Random Walk Down Wall Street*. Princeton, NJ: Princeton University Press.
- Lo, A. and J. Wang, 2006, “Trading Volume: Implications of an Intertemporal Capital Asset Pricing Model”, *Journal of Finance* 61, 2805–2840.
- Magnus, J. and H. Neudecker, 1988, *Matrix Differential Calculus: With Applications in Statistics and Economics*. New York: John Wiley & Sons.

- Lo, A. and C. MacKinlay, 1990, "When Are Contrarian Profits Due To Stock Market Overreaction?", *Review of Financial Studies* 3, 175–206.
- Markowitz, H., 1952, "Portfolio Selection", *Journal of Finance* 7, 77–91.
- Merton, R., 1973, "An Intertemporal Capital Asset Pricing Model", *Econometrica* 41, 867–887.
- Merton, R., 1981, "On Market Timing and Investment Performance I: An Equilibrium Theory of Value for Market Forecasts", *Journal of Business* 54, 363–406.
- Modigliani, F. and L. Modigliani, 1997, "Risk-Adjusted Performance", *Journal of Portfolio Management* Winter, 45–54.
- Newey, W. and K. West, 1987, "A Simple Positive Definite Heteroscedasticity and Autocorrelation Consistent Covariance Matrix", *Econometrica* 55, 703–705.
- Ross, S., 1976, "The Arbitrage Theory of Capital Asset Pricing", *Journal of Economic Theory* 13, 341–360.
- Sharpe, W., 1964, "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk", *Journal of Finance* 19, 425–442.
- Sharpe, W., 1966, "Mutual Fund Performance", *Journal of Business* 39, 119–138.
- Sharpe, W., 1991, "The Arithmetic of Active Management", *Financial Analysts Journal* 47, 7–9.
- Sharpe, W., 1992, "Asset Allocation: Management Style and Performance Measurement", *Journal of Portfolio Management* 18, 7–19.
- Sharpe, W., 1994, "The Sharpe Ratio", *Journal of Portfolio Management* 21, 49–58.
- Treynor, J., 1965, "How to Rate Management of Investment Funds", *Harvard Business Review* 43, 63–75.
- Treynor, J., 2005, "Why Market-Valuation-Indifferent Indexing Works", *Financial Analysts Journal* 61, 65–69.
- Treynor, J. and F. Black, 1973, "How to Use Security Analysis to Improve Portfolio Selection", *Journal of Business* 46, 66–86.
- Treynor, J. and K. Mazuy, 1966, "Can Mutual Funds Outguess the Market?", *Harvard Business Review* 44, 131–163.
- White, H., 1984, *Asymptotic Theory for Econometricians*. New York: Academic Press.