The Dynamics of Active Portfolios

Leigh Sneddon

Westpeak Global Advisors

July, 2005

Abstract

Active portfolio management is dynamic: new information arrives continually, and trading based on this information repeatedly changes the portfolio. With turnover controlled, the impact of today's trading persists into the future, where it will influence future trading and commingle past and future information. The best strategy for a single month is generally not the best strategy for a horizon of many months. Not surprisingly, important questions remain unanswered. What performance can investors reasonably expect over time? How can active managers improve multi-period performance? What are the best combinations of signals that predict returns immediately and signals that predict returns later? How much trading should be done? Put simply, active management currently delivers a multi-period product built with a single-period theory. This paper gives a solution for the time dimension of active management, taking the underlying theory from the single period to full multi-period performance. The solution includes expressions for active risk, turnover, and active return before and after trading costs. The paper contains the "Fundamental Law of Active Management" as a special case, and shows that actual performance can deviate significantly from its predictions. Examples illustrate the solution, showing how to improve the multi-period performance delivered to investors. With this information, the time dimension of active management becomes an opportunity we can exploit.

Introduction

Active managers build strategies that exploit predictive information or signals. The process is dynamic: the signals change through time and managers repeatedly trade or “rebalance” their portfolios to take advantage of the new information. With turnover controlled, the impact of today’s trading persists into the future, where it will influence future trading and commingle past and future signals.

There has been great progress in the understanding of single-period portfolio properties: see for example Markowitz (1959), Elton and Gruber (1995), and Grinold and Kahn (1999). The information ratio or IR is the ratio of active return to active risk and is the standard risk-adjusted measure of active management performance. The portfolio’s active return is its return in excess of the benchmark return, and the active risk, or tracking error, is the standard deviation over time of that active return. Grinold (1989) showed that the annualized before-costs IR with monthly rebalancing is given approximately by $IR \approx IC\sqrt{12N}$ where IC is the correlation between alpha[1] and the single-period return, and N is the number of assets in the investment universe. This result is known as the Fundamental Law of Active Management (FL). One of its approximations is treating the periods between successive optimizations as independent of each other. In other words, the FL is a single-period result. Turning to multi-period dynamics, many questions remain. Grinold and Kahn (1999) concludes with a list of open issues, the first of which is portfolio dynamics: “Active portfolio management is a dynamic problem…This full dynamic problem is both complicated and important.”

Numerical simulations address multi-period effects (see for example Lobosco, 2003 and Qian et al, 2004). Each simulation, however, is typically based on the signals and returns over a specific historical time period. By contrast, the analytic solution presented here gives the weighted averages - of portfolio return, risk, turnover[2] and other quantities - over all possible time series of signals and returns, with the weights determined by the variances and correlations.[3] This averaging over infinitely many possible scenarios adds robustness relative to an historical simulation, which includes only one scenario. Additionally, an investment strategy has many variables, such as the selection of signals, the weights with which the signals are combined, and portfolio turnover. Selecting the best configuration requires comparing many configurations, and for each one it is faster to evaluate an analytic solution than to run a simulation.

This paper has two goals. Active management delivers a multi-period product that is in essence built using a single-period model. The first goal is to develop and solve a multi-period model of active management. The second goal is to use the model to answer the following questions. What performance can investors reasonably expect over an interval containing multiple optimizations? How can active managers build their strategies to deliver the best performance over time? What are the properties, when turnover is controlled, of repeatedly rebalanced portfolios? How do signals with different predictive horizons interact and affect performance?[4] How should we combine them? How much trading should we do?

Section 1 compares this paper with other investment dynamics literature. Section 2 defines the model, and Section 3 describes the concepts of portfolio memory and the steady state of portfolio dynamics. The solutions for multi-period portfolio return, its standard deviation risk, turnover and other properties are in Section 4. Section 5 shows the impact on investment performance of the time horizon of predictive power. The results depend strongly on turnover. Section 6 combines a signal having short-term predictive power with one having long-term power, and shows the variation of investment performance with their relative weight, and with turnover. The Appendix provides additional information on the basic solution.

1. Comparative Survey of Related Literature

Markowitz and van Dijk (2004) describe an ideal portfolio theory: "Ideally, financial analysts would like to be able to optimize a consumption-investment game with many securities, many time periods, transaction costs, and changing probability distributions. We cannot."[5] Important multi-period literature has addressed portfolio problems where the probability distributions of return do not change over time. Examples include Morton and Pliska (1995), Pliska and Selby (1995), Atkinson, Pliska and Wilmott (1997), Leland (2000), Pliska and Suzuki (2002), and Oksendal and Sulem (2002). In these cases, the target asset weights are fixed over time.

Active management differs in important ways from the rich multi-period literature. A first difference is that the predictive information, and hence the return probability distributions, change through time. This means the target asset weights also change through time.

Grinold (1997, 2005) uses the framework of dynamic programming to approach multi-period investing with changing forecasts.[6] Like this paper, Grinold (1997, 2005) presents results on signals with predictive power beyond the current rebalance period and explores the impact of transaction costs and predictive lifetimes. Markowitz and van Dijk (2004) also study multi-period investing with changing forecasts. They provide a closed-form solution for a two-asset model, and use it demonstrate the accuracy with two assets of a trial or approximate solution that can be evaluated when there are more than two assets. In spite of the presence of changing forecasts, these contributions still differ significantly from this paper.[7]

A second difference between active management and the multi-period literature lies in the objective function.[8] A common feature of the literature, for example the papers just discussed, is an objective function consisting of a sum or integral of utilities, one for each rebalance. The utility function contains the predictive alpha, and a variance measure of predicted risk, and/or a penalty for trading. Active performance, however, is not judged on predictive alpha and predicted risk, but on realized portfolio returns and realized risk.[1] The objective functions for active management therefore depend not on alpha, but on return. Accordingly, in this paper "utility" refers to the function the optimizer maximizes at the beginning of each period, while “objective” refers to the multi-period active performance metric that the manager wishes to maximize. The utility deals with expected returns and the objective deals with realized returns. A practical model must tell us what portfolio risk and return we can reasonably expect to be realized for a given set of signals and specified portfolio construction parameters. This paper does this by first deriving and solving a dynamic model of active management, including alpha and return, and then giving the expected values of future realized portfolio risk, active return, turnover, and trading costs. These solutions allow us to compute objective functions or performance metrics used to assess active managers. The objective used here to illustrate the results is the active management standard: the information ratio (IR).

A third way in which this paper differs from the multi-period literature is the nature of the dynamic model. Here it includes additional features that are specific to the design and maintenance of active strategies. Beyond the changing forecasts, and return-based performance metrics, these features include randomly varying realized returns, explicit optimizations, and combinations of multiple, serially correlated[9] signals, with different, time-dependent predictive powers. In these ways this paper seeks to focus yet more closely than others on the specific demands of active management.[10, 11]

2. The Model

An active strategy consists of sequential optimizations, and repeated rebalances to new target weights over time, based on randomly evolving signals. The model consists of two major components: one describing the returns and the signals, and one describing portfolio construction. This Section describes each component and the overall information flow, and concludes with a list of the underlying assumptions. We use some notational conventions: the subscripts m and j label assets and signals respectively, and t is the rebalance date.

Returns and Signals

The residual return[12] of asset m between rebalance dates t and $t+1$ is $r_m(t)$. We assume that its mean is zero: $\mathrm{E}\{r_m\}=0$, and that it is uncorrelated over time.

New information flows into the signals continually, and we refer to each piece of new information, each “news flash”, as an “innovation”. These innovations make up the signals, and the signals are correlated over time. To capture this, the model treats each signal as depending both on its prior value and on the latest innovation. The symbol $a_{jm}(t)$ refers to the value of the j'th signal for security m at time t. The model for its evolution over time is[13]

$$a_{jm}(t)=f_j a_{jm}(t-1)+e_{jm}(t) \tag{1}$$

where $e_{jm}(t)$ is the innovation or news flash flowing into the j'th signal for security m at time t. The model parameter f_j is just the serial correlation of the j'th signal, that is the correlation of the signal with itself from one rebalance to the next. It is less than one, and can be thought of as a decay parameter: it is the discount factor describing how quickly the signal “forgets” prior information. We can readily estimate the f_j's by simply regressing the signal at time t against the same signal at time $t-1$. New information is by definition a surprise, so the means of the innovations are zero: $\mathrm{E}\{e_{jm}\}=0$.

The active manager combines the signals into an alpha, a composite measure of the relative attractiveness of each asset.[14] Often this is done through a simple linear weighting of the signals. The alpha for stock m is:

$$a_m(t)=\sum_j c_j a_{jm}(t) \tag{2}$$

where the multipliers c_j are fixed in time and uniform across assets, but can vary across signals, so some signals receive more weight than others. Notice that a with a subscript j is an individual signal, while a without a subscript j is the composite alpha.

We can set these multipliers to anything we choose, so nothing is lost by constraining the raw signals each to have an average variance of one: that is $\sum_m \mathrm{Var}\{a_{jm}\}=N$.

Randomness enters the model through the returns and through the innovations. At each rebalance date, and for each security, there is a set of random variables: a return, and one innovation for each signal. These variables are assumed to be jointly normal.[15] Their joint distribution is then defined by their variances and their correlations, both across assets and over time.

It is the correlation between innovations and future returns that gives the signals their predictive power. The correlation on security m, between return from t to $t+1$ and the information flowing to signal j at time t', is

$$T_j(t-t') \equiv \text{Corr}\{r_m(t), e_{jm}(t')\} \qquad (3)$$

For $t \geq 0$ the function $T_j(t)$ determines the forward predictive power of the innovations flowing into signal j, that is each innovation's correlation with the current period's return $(t=0)$, the next period's return $(t=1)$, etc. These correlations in turn give the signals their predictive power.

Before each return is realized, it is random and its distribution, when conditioned on the latest information, changes as new information arrives. This information is itself random so in this sense the returns are "randomly random."

Quadratic Portfolio Construction

The second component of the model describes how the portfolio manager can use alpha to set active weights.[16] Each rebalance is to a new portfolio that maximizes a quadratic utility function. The utility in turn consists of three parts: reward, a risk penalty, and a trading penalty. The reward is the active portfolio aggregate alpha:

$$\text{Reward} = \sum_m a_m w_m \qquad (4)$$

where w_m is the active weight of security m, that is security m's weight in the portfolio less its weight in the benchmark.

The second part of utility is risk aversion times tracking variance. The most general quadratic form for this component is $-g\sum_{m,n} w_m V_{mn} w_n$, where g is risk aversion and V_{mn} is the covariance between the returns of assets m and n. The model here considers the simple case of $V_{mn} = V_m$ if $m = n$ and 0 otherwise, so the risk component of utility is

$$\text{Risk Penalty} = -g \sum_m V_m w_m^2 \tag{5}$$

The third part of the utility function reflects the fact that trading is costly and so needs to be controlled. The most natural form for this component would be a penalty proportional to the turnover. Turnover is always positive and takes the form $\sum_m |\Delta w_m|$ where Δw_m is the change in active weight of security m : $\Delta w_m(t) \equiv w_m(t) - w_m(t-1)$. The absolute value operation, | . |, means that this penalty is neither linear nor quadratic in the security weights. It is piecewise linear because it "turns a corner" to remain positive as a buy becomes a sell. This piecewise linearity disrupts the analytic tractability of the model. The approach here is to set the third component instead to a quadratic penalty given by

$$\text{Trading Penalty} = -\frac{1}{2h} \sum_m (\Delta w_m)^2 \tag{6}$$

This penalty has the same general effect as one proportional to turnover: it controls turnover by reducing utility as the optimized holdings move further away from the prior holdings. The turnover tolerance $2h$ sets the scale of the penalty, with higher values allowing more turnover. Importantly, the model's solution includes the actual turnover, and hence also the trading costs.[17]

The utility is simply the sum of the three parts: (4), (5) and (6). Each portfolio optimization finds the new active weights that maximize this utility, and in so doing performs the trade-offs between alpha, risk, and trading. The result is a set of new active weights that, with the returns and trading costs, determine the portfolio's turnover, predicted risk, and realized return in that period. These quantities will vary from one optimization to the next. We will see in Section 3 that over time these variations are random fluctuations about the values given by the steady-state solution.

Figure 1 shows an information flow diagram of the model. At each rebalance new information updates the signals, which then combine to form the alpha. Portfolio construction takes alpha and current portfolio positions, and determines new positions. This process repeats over multiple rebalances. The resulting trading costs combine with security positions and returns to produce the multi-period investment performance.

Summarizing, the model includes a number of important features of practical active management. The solution reflects repeated optimizations that use a risk model and a trading penalty, and so control both risk and turnover. Importantly, the model includes not only predictive alpha but also realized return. Signals, returns and predicted return volatilities all vary across assets. The predictive alpha consists of multiple signals, and each of these signals can be correlated with the returns of multiple forward periods. The predictive power of each signal varies with the forward horizon, creating a "term structure" of predictive power, and different signals can have different term structures. Each signal is serially correlated, and the correlation strength varies across signals.

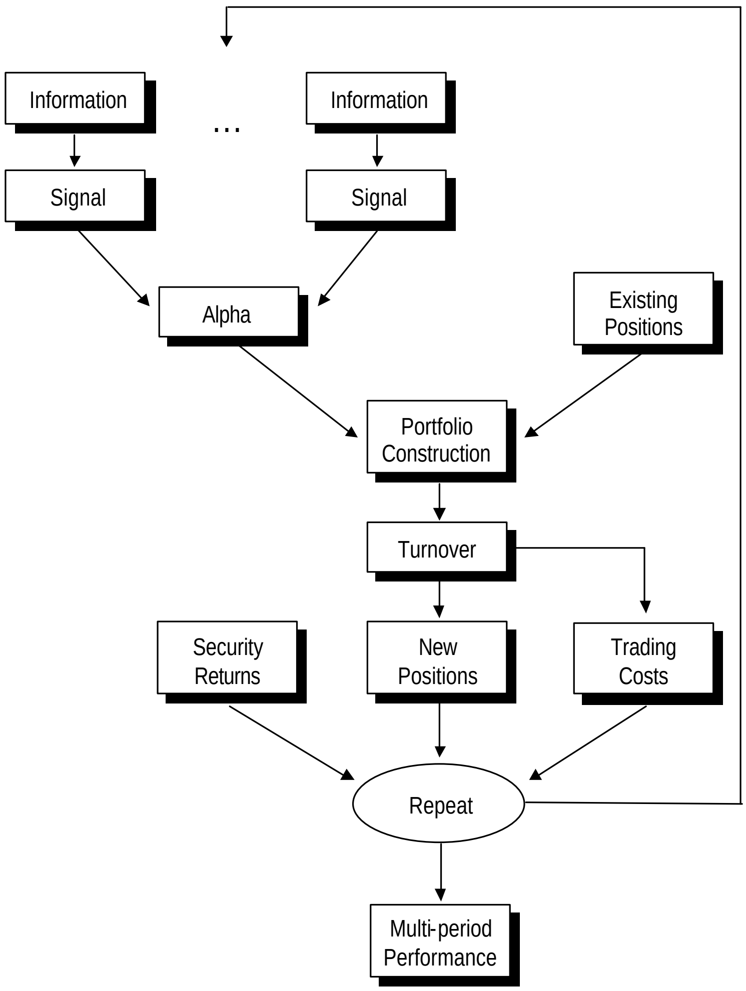

Figure 1: Information Flow of the Model

Assumptions

The world is complicated and the model reflects simplifying assumptions. For transparency we list them here: (i) short sales and leverage are allowed;[18,19] (ii) a signal for a security has predictive power only for that security, and the correlation between new information and forward return is uniform across assets; (iii) returns grow with simple interest; (iv) the unconditional mean of residual returns is zero, and the portfolio has the same beta as the benchmark; (v) returns and signal innovations are jointly normally distributed; (vi) returns are uncorrelated across time;[20] (vii) the risk model is constant over time;[21] (viii) trading costs are directly proportional to turnover;[22] (ix) drift in active weights between rebalances is ignored.[23]

Assumptions (i) through (v), or variations of them,[24] are standard in analytic single-period portfolio theory. Assumptions (vi) through (ix) arise only in the multi-period context and they, or variations of them, are also present in the multi-period literature. While it is not necessary to do so, this paper ignores correlations across stocks and across signals.[25]

3. Memory and Steady States

Signals accumulate information from the past, adding memory to active management: today's signal values are related to yesterday's. This is perhaps the simplest form of memory in active management.

A second form of memory arises because trading is costly, and many investors are averse to high turnover. Portfolios therefore typically have turnover control, which changes their behavior significantly. Without it, rebalancing is memory-less in the sense that changing the starting portfolio weights would leave the optimal weights unaffected. This latter scenario is the focus of most of analytic active portfolio theory. With turnover control, however, each set of optimal weights depends directly on the prior weights, which in turn depend on their prior weights, and so on back through time. With turnover controlled, portfolios have memory. With memory, positions taken in this period affect positions in future periods.

Additionally, performance is affected by the power of signals to predict return beyond the current period. This period's return can be predicted not only by this period's signals, but also by past signals as well. This is a third form of memory in active management.

Because they have memory, active portfolios are dynamic systems and, like many dynamic systems, can exhibit two regimes over time, the transient and the steady state. The transient depends on the starting configuration, but gives way to the steady state, which lasts indefinitely and is independent of the starting point.[26]

Figure 2 shows portfolio alpha over 50 consecutive rebalances in a Monte Carlo simulation[27] of the model with only one signal. Each point shows the post-rebalance portfolio alpha. The simulation starts with no active weights, the rebalance frequency is monthly, the universe has 200 assets, and the serial correlation of alpha from one month to the next is 0.6. The transient lasts for the first 6 or so rebalances after which the portfolio evolves into its steady state. The alpha then fluctuates around the steady state solution.[28,29]

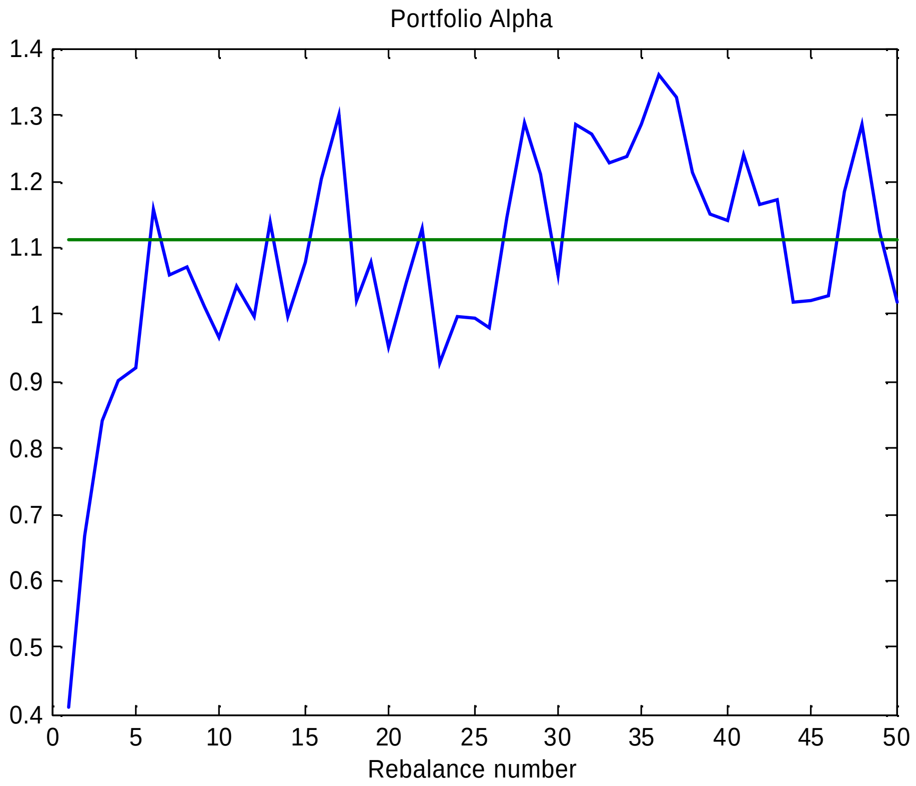

Figure 2: Portfolio alpha. The fluctuating line is the result of a simulation and shows the transient regime, lasting here about 6 rebalances, and the subsequent steady state regime, which persists indefinitely. The horizontal line is the steady state solution.

The dynamic model in this paper exhibits both the transient and the steady state regimes. The focus, however, is on the steady state because steady state behavior is long-lived, and its expected performance is systematically predictable. It is independent of how the portfolio starts, and is the aspect of active management ultimately delivered to investors.

4. The Solution

This Section provides the steady state expected values of the IC, active position size, portfolio alpha, active return, tracking error, turnover, trading costs and information ratio. For readers wishing simply to use the results, Table 1 contains the solutions and Table 2 contains the definitions of the symbols. The remainder of this Section provides, not a complete mathematical derivation, but some background to the origin and nature of the solution.[30] The Appendix displays and explains the behavior of active position size, turnover and portfolio alpha.

In Section 3 we saw three forms of memory in active management. Each of them means that a quantity we often think of as single-period in nature in fact has a multi-period character. The first is that, with information arriving over time, the signals themselves are multi-period in nature. In this model each signal is a discounted sum of the innovations, or pieces of information, that have formed it.

$$a_j(t) = \sum_{t=0}^{\infty} f_j^t e_j(t-t) \quad (7)$$

The discounting means that more recent innovations have greater weight.[31] The discount factor f_j is also equal to the signal's serial correlation. To reduce clutter in (7), (8), (9) and (14), we omit the subscript m that labels individual assets.[32]

The second form of memory comes from portfolio construction, specifically from turnover control, which causes active positions to depend directly on prior active positions, and so on back through time. At each time, active position also depends on the alpha at that time. The result of putting these pieces together is that each active position is proportional, not simply to the most recent alpha, but to the discounted sum of all prior alphas.

$$w(t) = hA\left[a(t) + Aa(t-1) + A^2 a(t-2) + \ldots\right] \tag{8}$$

A is shorthand for

$$A \equiv \left(1 + 2hgV\right)^{-1} \tag{9}$$

Recall that $2h$ is the turnover tolerance, g is the risk aversion and V is the risk model's forecast of return variance. These are all positive so A is less than one. We see from (8) that A is the discount factor controlling the weight of prior alphas in the active positions. When turnover tolerance is high, the relative weight on prior alphas decays more quickly with their age.

The third form of memory is that returns are correlated, not only with the most recent information innovations, but also with prior innovations. As we noted in the first form of memory, each signal is a discounted sum of prior innovations. Correspondingly, the IC of a signal is proportional to a discounted sum of correlations of return with each prior innovation.

$$IC_j = \sqrt{1 - f_j^2} \sum_{t=0}^{\infty} T_j(t)\, d^t \tag{10}$$

Recall from Section 2 that $T_j(t)$ is the correlation between current return and the innovation t periods earlier. A useful shorthand for the discounted sum of predictive correlations is “script T”:

$$\Im_j(d) \equiv \sum_{t=0}^{\infty} T_j(t)\, d^t \tag{11}$$

where d can be any discount factor. With uncorrelated signals, the total IC is just the average of the signal IC's, weighted by the signal multipliers c_j.[33] Using "script T" in (10) therefore gives

$$IC = \frac{\sum_j c_j \sqrt{1-f_j^2}\,\Im_j\left(f_j\right)}{\sqrt{\sum_j c_j^2}} \tag{12}$$

Now that we see the role of memory in the signals, active weights, and the IC, we can see how important features of the solution arise. How big, for example, is alpha? Section 2 shows that it is a combination of signals, each multiplied by a constant, c_j. The innovations have means equal to zero, so the signals, and hence alpha, also have means of zero. A good measure of the size of alpha is then its standard deviation, or the square root of its variance. We are setting the average variance of each signal to 1 so, with uncorrelated signals, the variance of alpha is then just the sum of the squares of the signal multipliers.

$$\mathrm{Var}\{a\} = \sum_j c_j^2 \tag{13}$$

How big are the active weights? We saw from the memory discussion that each active weight is hA times a discounted sum of current and prior alphas. The alphas have mean zero, so the active weights do too. Their variance will be $\left(hA\right)^2$ times the variance of the discounted sum of alphas. The signals making up alpha are independent of each other and so contribute separately to the variance, with each contributing an amount proportional to the square of its multiplier c_j. The values at different times, however, are not independent of each other because each signal is serially correlated. The variance of the discounted sum of each signal will reflect its serial correlation f_j, and the discount rate A. The result for the variance of active weight is

$$\mathrm{Var}\{w\} = (hA)^2 \sum_j c_j^2 \frac{1+f_j A}{(1-f_j A)(1-A^2)} \tag{14}$$

Turnover is the result of changing active weights. At each rebalance, each asset contributes an amount equal to its absolute change in active weight. Adding up the contributions of all assets, and multiplying by the number of rebalances per year, means that turnover is $f_{\mathrm{R}} \sum_m abs\left(change\left(w_m\right)\right)$. The absolute change in active weight will be related to the size of active weight, that is to $\sqrt{\mathrm{Var}\{w\}}$; and $\mathrm{Var}\{w\}$ is given by (14). Working out the details, the round-trip turnover is

$$Turnover = f_{\mathrm{R}} \sum_m \sqrt{(hA_m)^2 \sum_j c_j^2 \frac{4}{p} \frac{1-f_j}{(1-f_j A_m)(1+A_m)}} \tag{15}$$

This result reveals the expected similarity to (14). Knowing the turnover is important for two reasons. It allows us to see how other portfolio properties vary with turnover, and it allows us to compute the costs of trading.

As is common in practice, this paper assumes that trading costs are proportional to turnover. This means that trading costs are simply the product of tc and TN, where tc is the cost of trading measured in portfolio return per unit of round-trip turnover.

$$TradingCost = tc \times Turnover \tag{16}$$

If, for example, buys and sells each cost 25 bps of asset price, and 10% of the portfolio is bought and sold, then tc is 0.0025, $Turnover$ is 0.2, and the cost to the portfolio's return is 0.0025 * 0.2 = 5 bps.

How large is the tracking variance? With uncorrelated assets each asset independently contributes an amount equal to the variance of the product of its active weight and its return. As a result, the annualized tracking variance is simply

$$TrackingVariance = \sum_m s_m^2 \mathrm{Var}\{w_m\} \tag{17}$$

(14) gives us the value of $\mathrm{Var}\{w\}$, and tracking error is simply the square root of tracking variance.

Portfolio active return is the sum over assets of the product of active weight times return. We saw in the memory discussion that each active weight is hA times a discounted sum of current and prior alphas. Let's consider first the case of predictive power in the current period only. Return is then only correlated with current alpha and not with prior alphas. To determine the expected value of active return, we can therefore ignore the contribution of prior alphas to active weight. All we need then is hA times the expected value of the product of return and current alpha. This expected value is just the IC times the volatilities of single-period return and alpha. (12) and (13) give us the IC and the volatility of alpha, and the volatility of single-period return is the volatility of annual return s_m divided by the square root of f_{R}, the number of periods per year. Multiplying by f_{R} to annualize then gives portfolio active return before trading costs as $\sqrt{f_{\mathrm{R}}}\sum_m hA_m s_m \sum_j c_j\sqrt{1-f_j^2}\Im\left(f_j\right)$. Including multi-period predictive power, the third form of memory, involves the prior alphas and interactions between the discount rate A and the serial correlations f_j. Including these effects, the result for annualized active portfolio return before trading costs becomes

$$ActiveReturn = \sqrt{f_{\mathrm{R}}}\sum_m hA_m s_m \sum_j c_j\sqrt{1-f_j^2}\,\frac{f_j\Im_j\left(f_j\right)-A_m\Im_j\left(A_m\right)}{f_j - A_m} \qquad (18)$$

With solutions in hand for the steady state expected values of active return, trading cost, tracking variance and turnover, we can now determine many different performance measures. We use the industry-standard after-trading-cost information ratio as an example of a risk-adjusted performance metric. In terms of the above results, it is

$$IR = \frac{ActiveReturn - TradingCost}{\sqrt{TrackingVariance}} \qquad (19)$$

Sections 5 and 6 give examples of IR in particular situations.

The predictive alpha of the portfolio as a whole is of also of interest. It is the average of alpha over the portfolio, weighted by the active asset weights, and we seek its steady state expected value. In the memory discussion we saw that active weight is hA times a discounted sum of prior alphas. If the signals were not serially correlated, only the most recent alpha's contribution to the active weight would matter. We would then need the expected value of alpha in the active weight times the alpha we are averaging, in other words, the expected value of alpha times itself. Since alpha has mean zero, this is the variance of alpha, given in (13), and the portfolio alpha would be $\sum_m hA_m \sum_j c_j^2$. With serial correlation, the discount rate A_m and the serial correlation again interact, and the result is

$$PortfolioAlpha = \sum_m hA_m \sum_j c_j^2 \frac{1}{1 - A_m f_j} \tag{20}$$

The next two Sections give examples that use these solutions to determine and enhance investment performance.

Quantity	Solution
Active weight at time t	$hA_m\left[a(t)+A_m a(t-1)+A_m^2 a(t-2)+\ldots\right]$
Information Coefficient, IC	$\sum_j c_j\sqrt{1-f_j^2}\,\Im_j(f_j)\Big/\sqrt{\sum_j c_j^2}$
Variance of predictive alpha	$\sum_j c_j^2$
Variance of active weight	$\mathrm{Var}\{w_m\}=(hA_m)^2\sum_j c_j^2\frac{1+f_jA_m}{(1-f_jA_m)(1-A_m^2)}$
Turnover, round trip	$f_{\mathrm{R}}\sum_m\sqrt{(hA_m)^2\sum_j c_j^2\frac{4}{p}\frac{1-f_j}{(1-A_mf_j)(1+A_m)}}$
Tracking variance	$\sum_m s_m^2\mathrm{Var}\{w_m\}$
Active return	$\sqrt{f_{\mathrm{R}}}\sum_m hA_m s_m\sum_j c_j\sqrt{1-f_j^2}\,\frac{f_j\Im_j(f_j)-A_m\Im_j(A_m)}{f_j-A_m}$
Trading Cost	$tc\times Turnover$
Information ratio	$\frac{ActiveReturn-TradingCost}{\sqrt{TrackingVariance}}$
Portfolio alpha	$\sum_m hA_m\sum_j c_j^2\frac{1}{1-A_mf_j}$

Table 1: Solutions. Turnover, tracking variance, active return and information ratio are annualized.

Symbol	Definition
A, A_m	A_m is shorthand for $(1+2hgV_m)^{-1}$ and is the factor discounting prior alphas in asset m's active weight.
a_m	Predictive alpha for asset m
a_{jm}	Signal j on asset m
c_j	Manager-selected multiplier of signal j in predictive alpha
f_{R}	Annual rebalance frequency
g	Risk aversion
$2h$	Turnover tolerance
IR	Information ratio
N	Number of assets in the universe
j_j r_m	Serial correlation of signal j Return of asset m
s_m	Volatility of asset m's return, annualized
$T_j(t)$	Correlation of new information for signal j with the single-period return starting t periods after the information arrives.
$\Im_j(d)$	"Script T", a discounted sum of multi-period predictive correlations: $\Im_j(d) \equiv \sum_{t=0}^{\infty} T_j(t)\, d^t$
tc	The cost of trading, measured in units of portfolio return, per unit of turnover.
V_m	Risk model's prediction of the variance of asset m's return
w, w_m	The active weight of asset m

Table 2: Definitions of symbols

5. Term Structure and Turnover Control

Some signals have more power predicting immediate returns, while others have more power predicting returns for periods starting in the future. Figure 3 shows the term structure of two simple signals. Term structure here means the values of the Information Coefficient (IC) of the signal when predicting the return for the current period (at zero on the horizontal axis) and for subsequent periods. [34] The information flowing into the first signal predicts return in the current period only: it is a “just-in-time signal”. The information flowing into the second signal only predicts returns exactly four periods ahead.[35,36] The length of the period is whatever has been selected for the rebalance interval. In the examples in this paper it is one month.

Figure 4 shows the multi-period IR as a function of turnover[37], before and after transaction costs, for the just-in-time signal. The horizontal line in the Figure is the prediction of the FL. When turnover is at its maximum value and trading costs are zero, the IR predicted by the FL agrees with the real performance.[38] The turnover at which this occurs here exceeds 500% each way per year. The before-trading-costs (BTC) line touches the FL line at that maximum turnover point. Individual points labeled “… Sim” in the Figures are the results of Monte Carlo simulations run as checks of the solution.

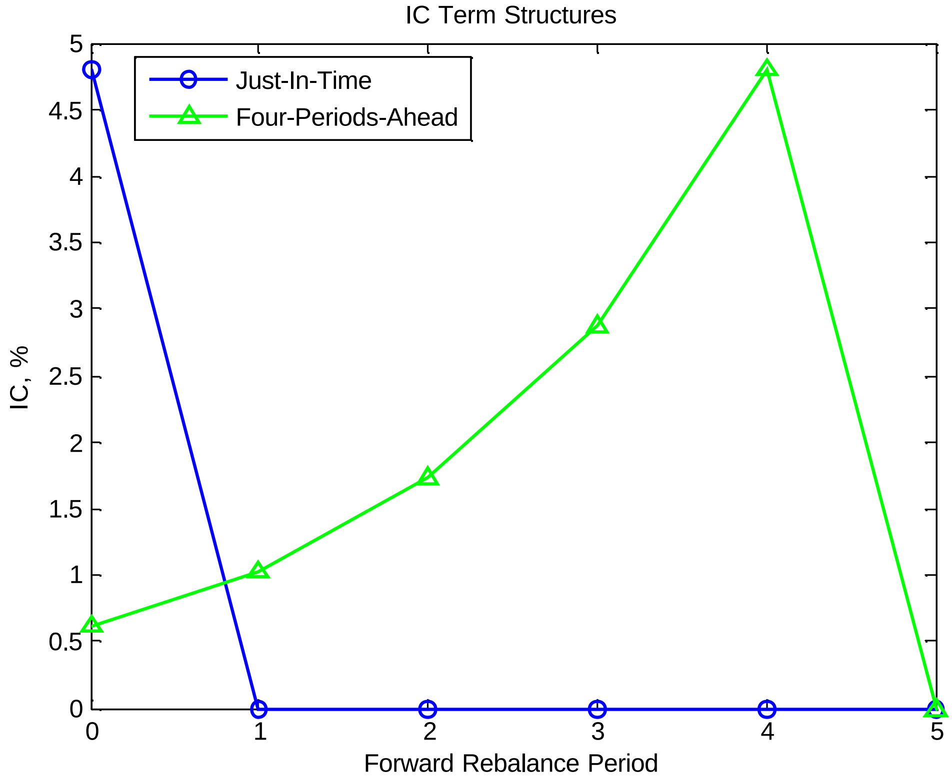

Figure 3: Term structures: IC's for Just-In-Time and Four-Periods-Ahead signals.

The BTC line also shows that the effect of controlling turnover is to reduce the investment performance significantly: as the turnover approaches zero, so does the IR. This is because the active positions become essentially fixed in time, while the alpha is fluctuating around zero, so the portfolio alpha averages to zero over time. The after-trading-costs (ATC) line in Figure 4 shows the extent to which trading costs of 0.25% each way reduce the IR.

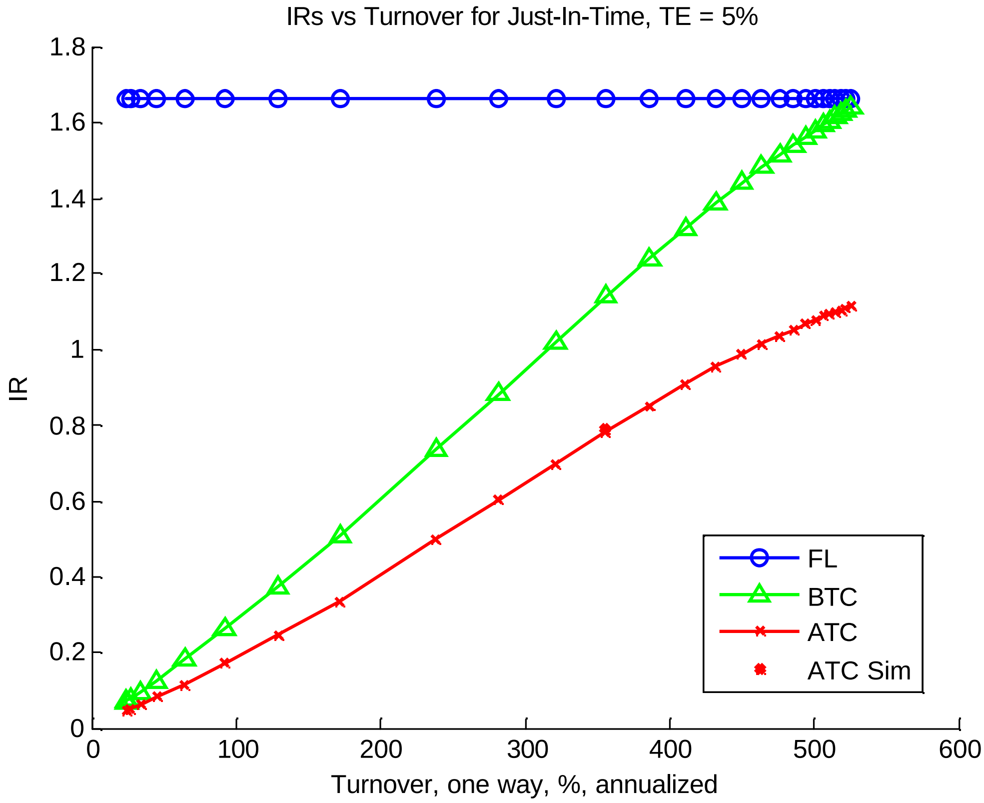

Figure 4: IR for just-in-time signal, plotted against turnover. The FL curve is the prediction of the Fundamental Law, and the BTC and ATC curves are the true performance, before and after trading costs respectively.

Why is the IR a strictly increasing function of turnover here? The reason is that for this signal incoming information only has predictive power in the current period. The faster the positions can change, the more able they are to keep up with the changing alpha, and so the larger is the portfolio's alpha. It is impossible for them to change too quickly because it is only the alpha itself that makes them change, and the alpha's impact is immediate. This effect even persists after taking into account reasonable costs of trading. We see that the solution quantifies the benefits of high turnover in the presence of signals with short-term predictive power. The after-cost multi-period IR for the just-in-time signal is always less than the FL prediction.

For signals that predict returns further ahead, the situation is different, as Figure 5 shows. This Figure shows the multi-period results, and the FL, for the signal whose predictive power occurs further in the future. Again the FL, based on the current period IC, is correct when there is no turnover control. At the maximum turnover point the BTC curve again touches the FL line, as happens for the just-in-time signal. In the case of a longer predictive horizon, however, there is a benefit to reduced turnover, with IR peaking here at a turnover of around 200%. This is because the active positions can then remain in place long enough to exploit the delayed predictive power of the signal. Higher turnover removes the positions before they have paid off.

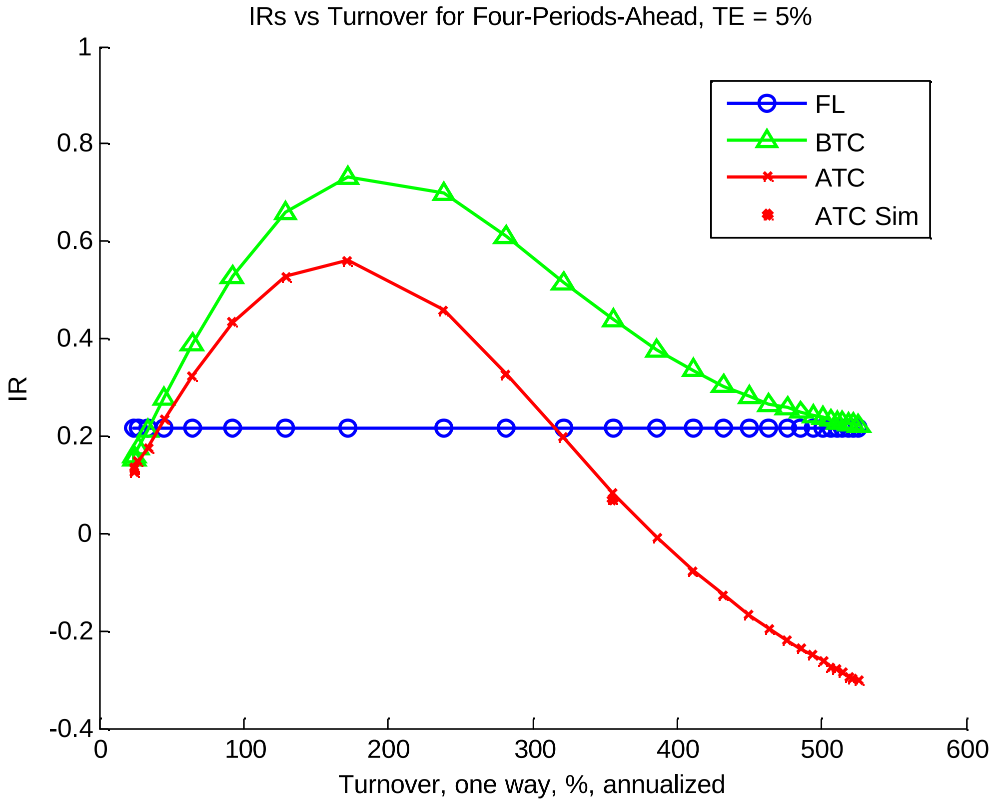

Figure 5: IR for 4-period-ahead predictive power, plotted against turnover. The FL curve is the prediction of the Fundamental Law, and the BTC and ATC curves are the actual performance before and after transaction costs respectively.

In this case of longer horizon predictive power the IR can exceed the prediction of the FL. This is because the FL is sensitive to the current period only. With delayed predictive power, reducing the rate of change of the active positions harvests that power more fully, enhancing performance beyond the FL prediction.

As seen in Figure 4 and Figure 5, the solution shows how to optimize investment performance by selecting turnover to match the alpha's term structure.

6. The Tortoise and the Hare

In practice, active managers often combine signals with different term structures. At the same time as selecting the correct turnover, we can also improve performance by selecting the best relative weighting of the different signals. Grinold (1997) writes, “We hope in the future to deal with a … thornier issue called the tortoise and the hare problem. How, in the presence of transaction costs, do we optimally combine two strategies with short (the hare) and long (the tortoise) information horizons?” This Section solves the problem of the tortoise and the hare.

The hare is a just-in-time signal like the one studied in Section 5. We blend it with each of two different tortoises. Both have predictive power that extends into the future and decays exponentially as the time lapse increases. For one, the “faster” tortoise, the decay half-life is three months and for the other, “slower” tortoise, the half-life is one year. [39] Figure 6 shows the term structures for these three signals. Their strengths are chosen so that the signals separately deliver comparable peak IR’s.

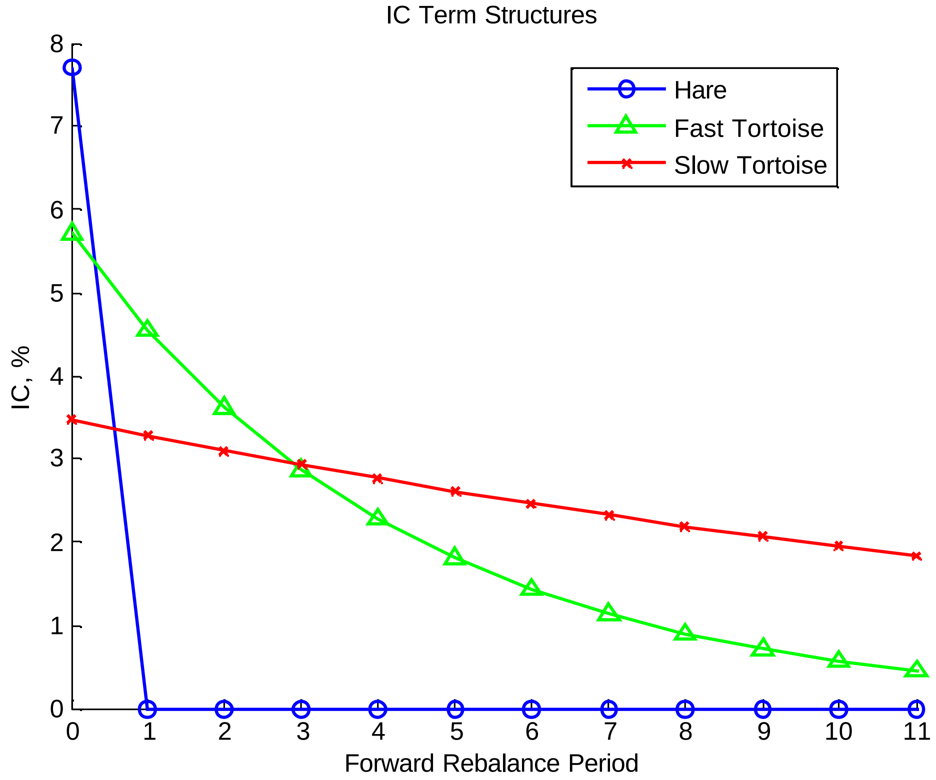

Figure 6: Term Structures for the Hare and two Tortoises

Figure 7 shows the FL predictions for IR as a function of turnover for various combinations of the hare with the "three-month" or faster tortoise. The curves are flat because the FL does not take turnover into account. Figure 8 shows that true performance is quite different.[40] The FL predicts that a 40% weight on the tortoise gives the greatest IR. Figure 8 shows that this is the correct choice at annual turnovers of over 400% each way, but that at moderate turnovers, putting 60% to 80% of the weight on the tortoise significantly increases the IR. Figure 8 also shows that controlling turnover to values below 100% significantly reduces the performance of all the combinations.

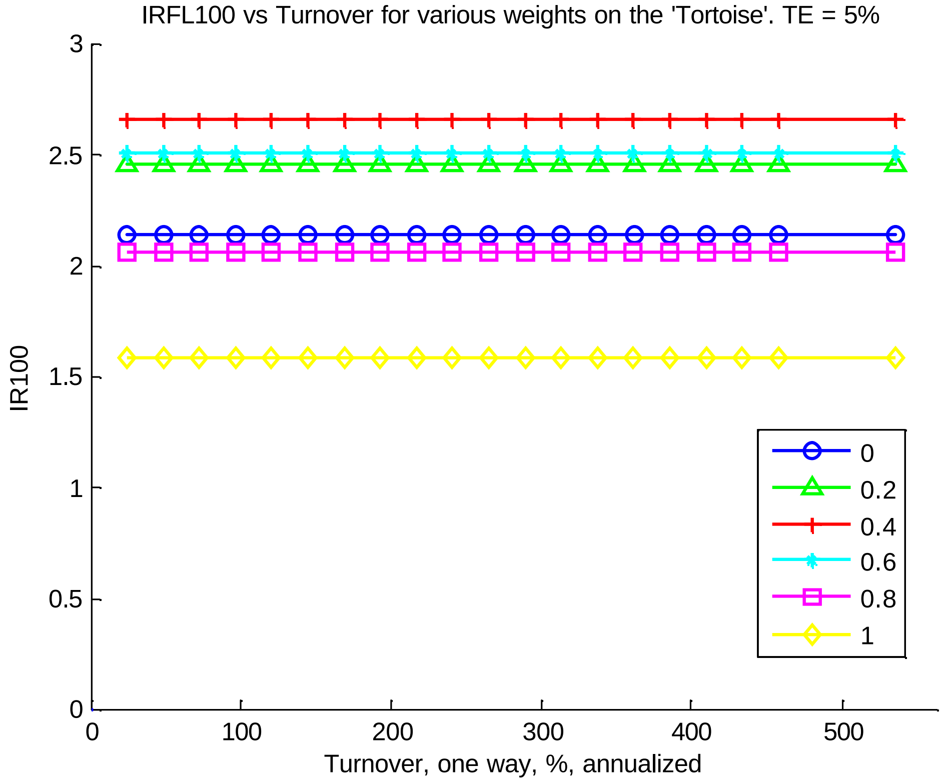

Figure 7: FL predictions for various combinations of the hare and the three-month or faster tortoise, plotted against turnover. The legend shows the weights on the tortoise.

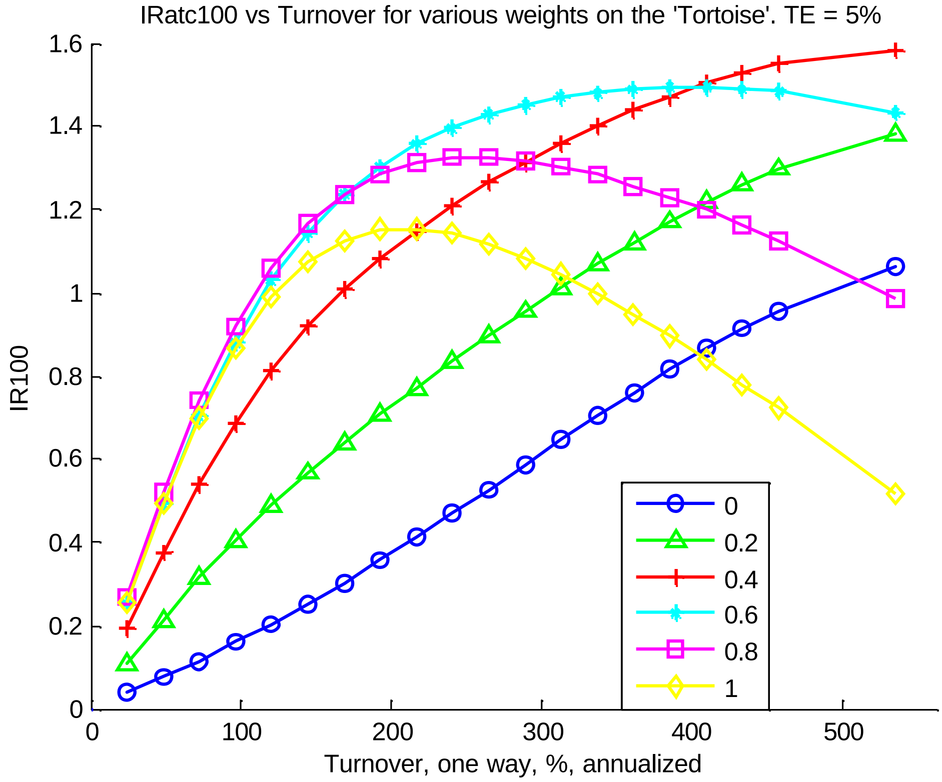

Figure 8: IR for various combinations of the hare and the three-month, or faster, tortoise, plotted against turnover. The legend shows the weights on the tortoise.

When the hare mixes with a slower tortoise, one with a predictive half-life of for example one year, another feature emerges. Figure 9 shows that peak performance then occurs in two scenarios: maximum turnover with most of the weight on the hare, and controlled turnover with most of the weight on the tortoise.[40] For this case the FL indicates giving the tortoise 20% to 40% of the weight, no matter what the turnover is. Figure 9 again shows that this is correct when turnover is very high. At turnovers of around 100% each way, however, putting most of the weight on the tortoise strongly improves performance.

Conversely, putting most of the weight on the tortoise at the highest turnovers significantly reduces the IR.

If turnover differs from what is best for the alpha, a lot of performance can be 'left on the table'.

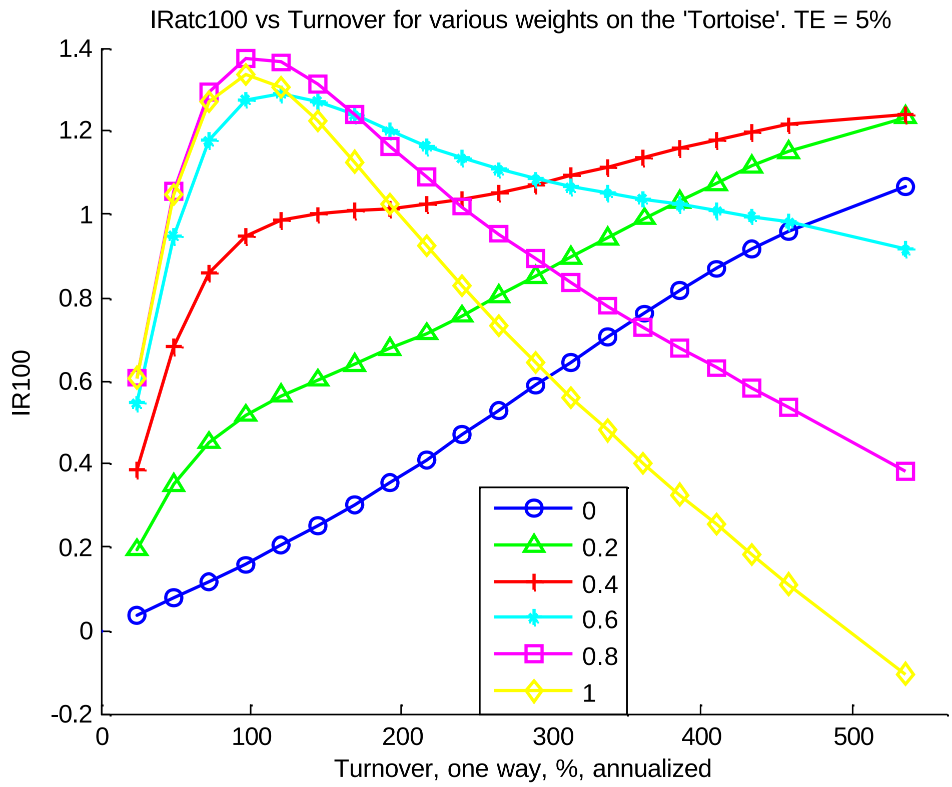

Figure 9: IR for various combinations of the hare and the twelve-month, or slower, tortoise, plotted against turnover. The legend shows the weights on the tortoise.

7. Conclusion

We can see at least two dimensions in active management: the cross-section of assets, and repeated trading sessions across time. The traditional underlying theory has only the first dimension. We know, however, that time is important: new information arrives continually, and trading based on this information repeatedly changes the portfolio. Further, the impact of today's trading persists into the future, where it influences future trading and commingles past and future information. Put simply, active management is a multi-period product built on a single period model. Most progress in multi-period investing has focused on problems that differ in significant ways from active management, and a solution containing multi-period realities as well as the practical specifics of active management has not been available.

This paper gives a solution for the multi-period dynamics of active management. It provides expressions that convert familiar single-period entities into their multi-period counterparts, and it quantifies explicitly multi-period behavior in active portfolios. The solution includes a number of important features of practical active management. Repeated optimizations use a risk model and a trading penalty, controlling both risk and turnover. The solution evaluates practical performance metrics based on realized portfolio return. Signals, returns and predicted return volatilities all vary across assets. The predictive alpha consists of multiple signals, and each of these signals can be correlated with the returns of multiple forward periods. The predictive power of each signal varies with the forward horizon, creating a "term structure" of predictive power, and different signals can have different term structures. Each signal is serially correlated, and the correlation strength varies across signals.

The paper uses examples to illustrate the solution. The first set computes the performance of two signals with simple multi-period properties: each piece of information flowing into the signals is only predictive of return in one period. For the first or "just-in-time" signal it is the period beginning just as the information arrives, while for the second or "four-periods ahead" signal it is the period beginning four periods later. These simple predictive properties can be considered the "atoms" making up more realistic signals. Both examples show that performance depends strongly on turnover. For the just-in-time signal, higher turnover means better the performance, even after reasonable trading costs. The four-periods-ahead signal has a "sweet spot", an optimal turnover which allows the active positions to keep up with the predictive alpha, while staying in place long enough to allow the delayed predictive power to deliver its benefits.

The second set of examples solves the problem of the "tortoise and the hare": how best to combine a signal that has short-range predictive power with one that has long-range power. The decision is then a joint one: selecting both turnover and signal weights. The examples show optimal selections when the long-range power decays moderately quickly and more slowly.

The paper also shows how to determine what performance investors can reasonably expect over an interval containing multiple optimizations. There are significant differences between multi-period performance and the predictions of the widely used Fundamental Law of Active Management. These arise because the Law does not reflect the impact of controlling turnover and this impact is large at low turnover; it does not quantify trading costs and these are large at high turnover; and it does not reflect predictive power beyond the current period.

The examples show how the solution can be used to improve the long-run performance delivered to investors. With this information, the time dimension of active management becomes an opportunity we can exploit.

Acknowledgments

I am grateful to my Westpeak colleagues for their interest in this work. The manuscript has benefited greatly from the insightful suggestions of Nardin Baker, Dan DiBartolomeo, Doug Holmes, Mark Hooker, Angelo Lobosco, John Minahan, Larry Pohlman, Gita Rao, Easton Ragsdale, and Heydon Traub.

.

Appendix: Active positions, Turnover and Portfolio Alpha

This Appendix displays and explains in greater detail the basic properties of the solution, in particular the behavior of active position size, portfolio turnover, and portfolio alpha.

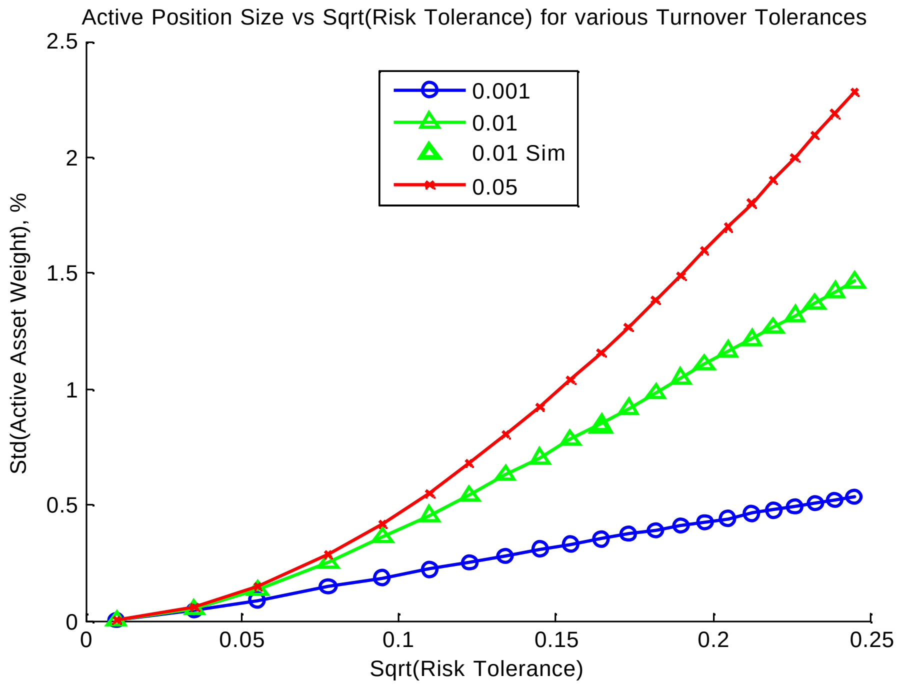

Figure 10: The size of active positions, plotted against the square root of risk tolerance. Each curve is for a different value of turnover tolerance, as shown in the legend.

Figure 10 shows how the active position size, $\sqrt{\mathrm{Var}\{w\}}$, varies with turnover tolerance $2h$ and risk tolerance $1/g$. As expected, increasing risk tolerance increases the size of active positions. For large risk tolerance active position size grows proportionally to the square root of risk tolerance. [41]

Increasing turnover tolerance also increases active position size. This is because active weights fluctuate about zero, so they can only achieve larger values by increasing the differences between successive values, that is by increasing turnover.

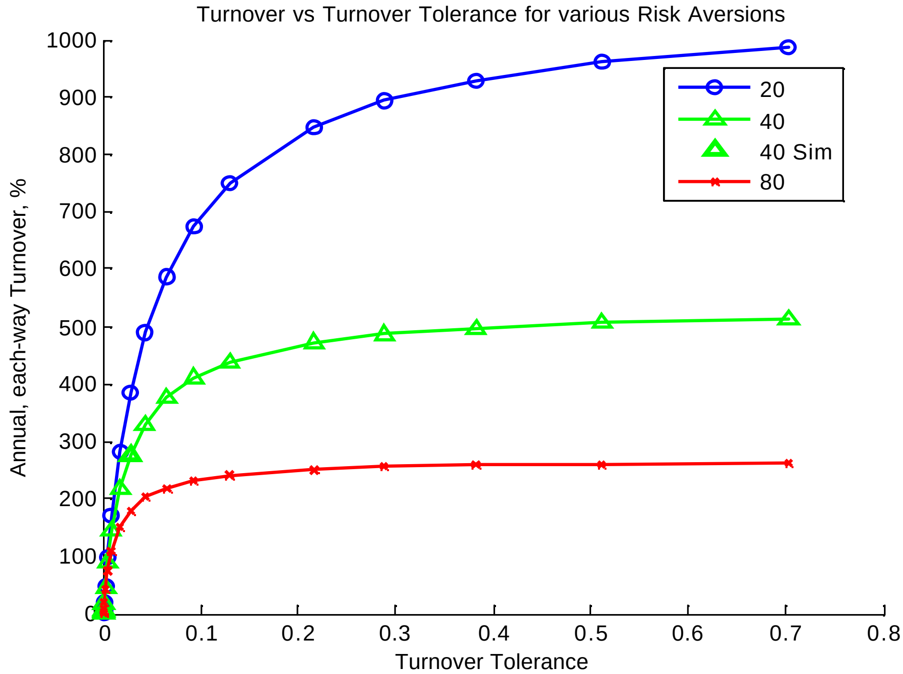

Figure 11: Turnover plotted against turnover tolerance. Each curve is for a different value of risk aversion, as shown in the legend

Figure 11 shows the impact of turnover tolerance $2h$ and risk aversion g on turnover. Turnover is of course an increasing function of h. As turnover tolerance gets larger, however, its impact on turnover continually diminishes and the curves in Figure 11 flatten as h increases. This is because in this paper changes of alpha are the only drivers of turnover. Therefore even when turnover is not controlled, it still has a maximum value that is limited by the volatility of alpha. Turnover also increases with decreasing risk aversion because, with larger risk, alpha drives changes between larger active positions.

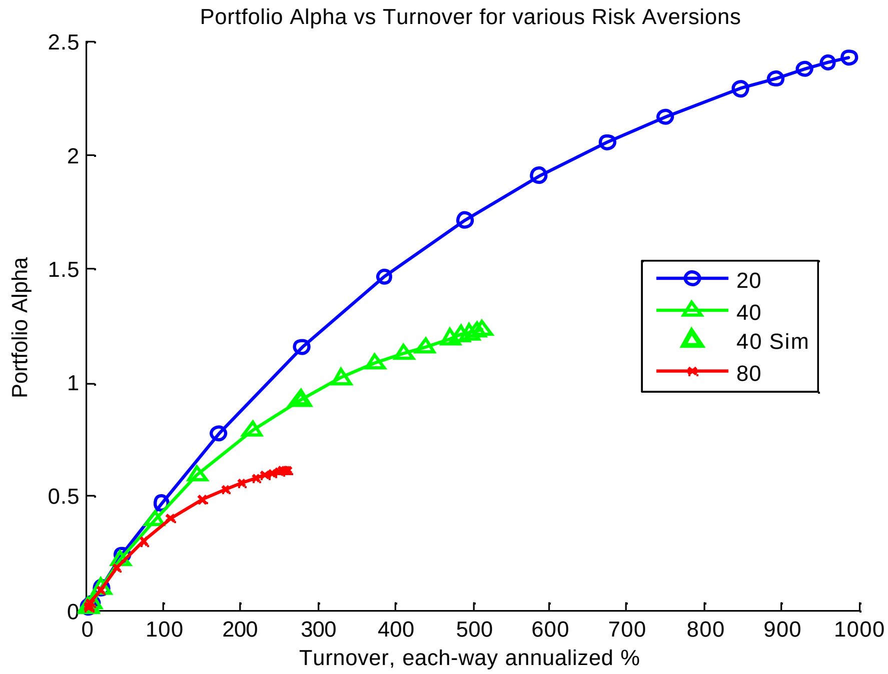

Figure 12: Portfolio alpha plotted against turnover. Each curve is for a different value of risk aversion as shown in the legend

Figure 12 shows how portfolio alpha varies with turnover and with risk aversion g . With no turnover allowed, portfolio alpha averages to zero over time because the holdings are fixed while the security alphas fluctuate about zero. Increasing turnover allows the portfolio to capture more of the changing alpha. The curves in Figure 12 each stop at a maximum turnover value, the same maximum identified in Figure 11.

REFERENCES

Bernstein, P. L. 2004. Proceedings of the CFA Institute Seminar: *Points of Inflection: New Directions for Portfolio Management*: Page 1.

Campbell, J.Y., L.M. Viceira. 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.

Elton, E.J., and M.J. Gruber. 1995, *Modern Portfolio Theory and Investment Analysis*, Fifth Edition, Wiley

Fama, E.F. 1970 "Multi-period Consumption-Investment Decisions." *American Economic Review*, vol. 60 (March):163-174.

Grinold, R.C. 1989. "The Fundamental Law of Active Management." *The Journal of Portfolio Management*, (Spring).

Grinold, R.C., and M. Stuckelman. 1993. "The Value-Added/Turnover Frontier" " *The Journal of Portfolio Management*, (Summer): 8-17.

Grinold, R.C. 1997. "The Information Horizon." *The Journal of Portfolio Management*, (Fall):57-67.

Grinold, R.C. 2005. "A Dynamic Model of Portfolio Management", working paper, March 2005.

Grinold, R.C., and R. N. Kahn. 1999. *Active Portfolio Management*, Second Edition, McGraw Hill.

Lobosco, A. 2003. "The Transfer Coefficient: Implications for Portfolio Management." *Working Paper*.

Markowitz, H.M. 1959. *Portfolio Selection: Efficient Diversification of Investments*. New York: John: Wiley and Sons. 1991. 2nd ed. Cambridge, MA: Basil Blackwell.

Markowitz, H.M., and E.L. van Dijk. 2003. 'Single-Period Mean-Variance Analysis in a Changing World." *Financial Analysts Jou*rnal, vol. 59, no. 2 (March/April):20-44.

Mulvey, J. M., W.R. Pauling, R. E. Madey. 2003 "Advantages of Multi-period Portfolio Models" *The Journal of Portfolio Management*, Winter 2003.

Qian, E., E.H. Sorensen, R. Schoen, and R. Hua. 2004. "Multiple Alpha Sources and Active Management" *The Journal of Portfolio Management*, Winter 2004.

Sneddon, L. 2005. Portfolio Dynamics for Active Managers: Correlated Assets, Correlated Signals, and Imperfect Risk Models. *Working Paper*

Waring, M.B. 2004. Proceedings of the CFA Institute Seminar: *Points of Inflection: New Directions for Portfolio Management*. Page 71.

[1] In this paper “return” refers to actual or realized return: the type that pays the bills. By contrast, “alpha” refers to a predictive signal, or a combination of predictive signals. It does not refer to absolute or relative realized return, or pay the bills.

[2] Turnover is the value of assets purchased plus the value of assets sold, all divided by the value of the portfolio.

[3] These averages are also the results to which a Monte Carlo simulation converges over time.

[4] Grinold (1997) expresses this gastronomically and much more eloquently: “The interplay of information and time is as subtle as the interplay of food and time. Fresh is best is a good rule but not universally accurate; vegetables and baked goods are best when fresh; fruit needs to ripen; red wine benefits with age; and sherry is best as a blend of several vintages.”

[5] Like Markowitz and van Dijk (2004), this paper does not optimize consumption. It does however provide an exact solution containing the other features these authors mention: many assets, many time periods, transaction costs, and changing probability distributions. The “ideal” solution requires that consumption be included, and may in fact be feasible.

[6] A discounted dynamic programming model underlies Exhibit 12 in Grinold (1997). In this model, alpha evolves as an AR(1) process. A mapping determines u_t, the active weight at time t, as a function of the current alpha and the prior active weight. The mapping is determined by solving the equation $V\left(u_{t-1},\alpha_t\right)=\mathrm{Max}_{u_t}\left[U\left(u_{t-1},\alpha_t,u_t\right)+\Delta\mathrm{E}\left\{V\left(u_t,\alpha_{t+1}\right)\mid\alpha_t\right\}\right]$, where Δ is a discount factor between 0 and 1, and $U\left(u_{t-1},\alpha_t,u_t\right)=\alpha_t u_t-\frac{1}{2}\lambda u_t^2-TCost\left(u_{t-1}\rightarrow u_t\right)$. The author is grateful to Richard Grinold for this information.

[7] In addition to the differences mentioned in Section 1, the dynamics of Markowitz and van Dijk (2004) differ in their details from this paper and from other literature mentioned here. The mean and variance of the risky security’s return switch over time among a discrete set of alternate (mean, variance) pairs. The switching is based on a matrix of transition probabilities. At each time

step a table specifies the next portfolio based on the current portfolio and the current (mean, variance) pair. The solution assumes a minimum level of predictive power.

[8] To quote Markowitz and van Dijk (2004, p41) somewhat out of context: “The hard part has to do with the objective”

[9] A signal is serially correlated if its value at one time is correlated with its value at other times.

[10] Other multi-period literature of course addresses active management, for example Grinold (1997, 2005), Lobosco (2003), Mulvey et al (2004) and others. Markowitz and van Dijk (2004) also get at many of the same concepts in different forms.

[11] This paper also differs from the referenced literature in that it allows the predictive term structures to vary independently of signal volatilities; it distinguishes the control of turnover from the cost of turnover; and it handles arbitrarily long horizons without introducing a utility discounting factor.

[12] Residual returns are raw returns less an adjustment equal to beta times the market return. If the portfolio has the same beta as the benchmark, the active residual return is the same as the active raw return.

[13] This is a first order auto-regressive, or AR(1), process.

[14] In the multi-period context, alpha is not the forecast of a return, but a (generally composite) signal that correlates with various returns, each over a different period.

[15] Campbell and Viceira (2002), on page 25, compare the merits of using normal returns and using lognormal returns. The latter is equivalent to replacing “return” everywhere in this paper by “logarithm of one plus return”. The distinction is less important for shorter periods. As a very rough guide, with an annualized return dispersion of 40%, a monthly return hitting –100% is an 8.6 (=100/(40/Sqrt(12))) standard deviation event. As Campbell and Viceira point out, the ultimate solution is a continuous time model.

[16] Active weight or active position is the difference between an asset's weight in the portfolio and its weight in the benchmark.

[17] This leaves for another time of the impact on performance of different methods of turnover control.

[18] Comparing the average position size (see (14)) to the average benchmark weight indicates the likely numerical significance of allowing short sales. Waring (2004) describes the long-only constraint as the "biggest and most important factor in producing inefficient portfolios". Many investment products now allow short sales. If Bernstein (2004) is correct about the "impending death of long-only as a conventional strategy", the theory's allowing short sales will be less of an issue going forward.

[19] Deviations from 100% invested will average to zero, and will be small for low tracking errors.

[20] Predictive signals, however, are based on market and asset properties, for example earnings yield, and so each signal is correlated over time.

[21] In other words we treat risk factor exposures and covariances as fixed at their average value over time. This is not exact for processes that use fresh risk model data each month. It is, however, a reasonable heuristic when the principal objectives are, as here, long term averages of portfolio properties. Moreover, for processes that do not use a commercial risk model in optimization, but rely on static controls for example on assets, industries and sectors, the risk model is fixed over time.

[22] Trading costs can also depend for example on trade size, trading volume, volatility and execution strategy.

[23] This effect is generally small because it depends on the difference between portfolio and benchmark of an effect driven by the differences in returns across assets.

[24] See for example Grinold and Kahn (1995), page 303.

[25] To focus here on dynamics and avoid matrix mathematics, we defer the solution with cross-sectional correlations, and correlations between signals, to a subsequent paper (Sneddon, 2005).

[26] Lobosco (2003) also distinguishes the dynamic steady state from transient behavior.

[27] The Monte Carlo process here is as follows. At each rebalance a random number generator generates an innovation $e_m(t)$ for each asset m. Combining this with the previous alpha according to (1) gives the new alpha for that stock. The portfolio is then optimized using the new alpha values, and the new portfolio alpha is computed and plotted as a point on the chart.

[28] First rebalancing with unrestricted turnover moves the portfolio immediately into the steady state.

[29] The results in this paper are from the closed-form steady-state solution in Section 4. The only exceptions are Figure 2 demonstrating the transient and steady state, and individual points labeled "… Sim" in the charts and used as checks on the steady state solution.

[30] We make another optional assumption here, because it is often made in practice and it simplifies the equations: we treat the variance of each signal as uniform across assets. To avoid this assumption and recover the more general solution, the reader needs only to replace each c_j in Section 4 by $c_j\sqrt{\mathrm{Var}\{a_{jm}\}}$. The quantities IC, $\mathrm{Var}\{a\}$ and $\mathrm{Var}\{w\}$ given by (12), (13) and (14) respectively then vary across assets and so become IC_m, $\mathrm{Var}\{a_m\}$ and $\mathrm{Var}\{w_m\}$.

[31] This is seen by substituting (1) in itself, and repeating back through time.

[32] To recover the full results, append the subscript m to the variables A, a, a_j, e_j, s, V, and w.

[33] The correlations T_j, and hence the summary measure $\Im_j$, measure predictive power and so are small numbers, typically less than 0.1.

[34] The IR's are annualized, and risk aversion is chosen to produce a tracking error of 5%. Except where noted, all IR's are after transaction costs of 0.25% each way, and are based on monthly rebalancing, an investment universe of 100 assets, and signal serial correlations from one month to the next of 0.6.

[35] In the notation of Equation (3), for the 'just-in-time' signal $T(t)$ is zero except for $t = 0$, while for the four-periods-ahead signal $T(t)$ is zero except for $t = 4$.

[36] Figure 3 shows that the four-periods-ahead signal also predicts returns that are less than four periods ahead. This is because it accumulates the information from the past. At time t the innovation that just arrived is correlated with the four-period-ahead return, that is the return for the period beginning $t+4$. The innovation that arrived one period ago, at $t-1$, is also part of the signal and is correlated with the return that is now only three periods ahead. The information that was new two periods ago is still part of the signal now. That information predicts return that was 4 periods ahead then, but is only two periods ahead now. And so on, back to the innovation that arrived at $t-4$, which is correlated with return for the period beginning now, zero periods ahead.

[37] The charts in Sections 5 and 6 show other portfolio properties as functions of turnover and tracking error. This is done, for each point in the charts, by finding the values of turnover tolerance $2h$ and risk aversion g that give the desired turnover and tracking error, and then using those values when evaluating the other portfolio properties.

[38] Our solution contains the Fundamental Law. It can be reached by making four assumptions in addition to those outlined in Section 2. The solution in Section 4 is free of these additional assumptions, but it simplifies to the FL if we are willing to make them. They are that (a) there is no turnover control ($h \to \infty$) (b) trading costs are zero ($tc = 0$) (c) all assets have the same realized volatility ($s_m = s$) (d) the risk model predicts the same volatility for all assets ($V_m = V$).

[39] The half-life is a time interval that gives the speed of the exponential decay. The power is cut in half for every additional half-life that elapses between the information's arrival and the beginning of the return period. For example, if the time lapse is three half-lives long, the predictive power is only one eighth what it is when the time lapse is zero. Mathematically, the predictive correlations are $T_j(t) = T_j(0)2^{-t/halfLife}$ and the IC follows suit: $IC(t) = IC(0)2^{-t/halfLife}$, so both are non-zero for all t.

[40] Figure 8 and Figure 9 both reflect each way transaction costs of 50 bps.

[41] This contrasts with the single-period limit, where the active position size becomes linear in the risk tolerance itself rather than in its square root. The single period limit is achieved by allowing the turnover tolerance $2h$ to become infinitely large.

WESTPEAK
GLOBAL ADVISORS

THE DYNAMICS OF ACTIVE PORTFOLIOS

Leigh Sneddon

Montebello
July 2005

Outline

- Why Dynamics?
- Multi-Period Model
- Memory and the Steady State
- Solving the Model
- Performance Examples

Why Dynamics?

Active Management is Dynamic

- ***New information arrives continually, and trading repeatedly changes the portfolio over time***
- ***Memory: with turnover controlled, today's positions depend on today's information but also on***
  - yesterday's positions and
  - yesterday's information
- ***Memory: today's signal values are correlated with yesterday's***
- ***Memory: returns over future periods are correlated with signals present and past***

Why Dynamics?

Questions to Answer

- What are the properties,when turnover is controlled, of repeatedly rebalanced portfolios?
- How do signals with different predictive horizons interact and affect performance?
- How should we combine them?
- How much trading should we do?
- What performance can investors expect over time?
- How can we deliver the best long-term performance?
- How accurate is the Fundamental Law?

Why Dynamics?

Put simply . . .

Active management is a multi-period product
built on a single period model

Multi-Period Model

- ***Signals and returns***
- ***Portfolio rebalancing***

Multi-Period Model

Signals and Returns

- *Alpha is a linear combination of signals*

$$a_m(t) = \sum_j c_j a_{jm}(t)$$

- *Each signal is an accumulation of new pieces of information, or innovations*

$$a_{jm}(t) = f_j a_{jm}(t-1) + e_{jm}(t)$$

- *Innnovations, and therefore signals, are correlated with forward returns*

$$T_j(t-t') \equiv \mathrm{Corr}\{r_m(t), e_{jm}(t')\}$$

Portfolio Rebalancing

- *Utility = Alpha − Risk Penalty − Trading Penalty*

$$= \sum_m a_m w_m \quad - g \sum_m V_m w_m^2 \quad - \frac{1}{2h} \sum_m (\Delta w_m)^2$$

- *This is the Utility for each rebalance.*
- *The manager's ultimate Objective however depends on realized portfolio return, risk and trading cost from many such rebalances over time*

Assumptions

- Short sales and leverage allowed
- Ignore risk model variations over time
- Optional: no correlations across securities in the risk model or in the signals
- Optional: no correlations between signals
- See paper for other assumptions

Multi-Period Model

Features

- Repeated optimizations that control risk and turnover
- Predictive alpha made up of multiple signals
- Realized returns
- Multi-period predictive power: signals correlate with the returns of multiple forward periods
- Term structure: the power of each signal varies with the interval to the forward return period
- Different signals can have different term structures
- Serial correlation of each signal, with strengths that can differ across signals
- Signals, return volatilities and predicted return volatilities vary across assets

Multi-Period Model

Information Flow

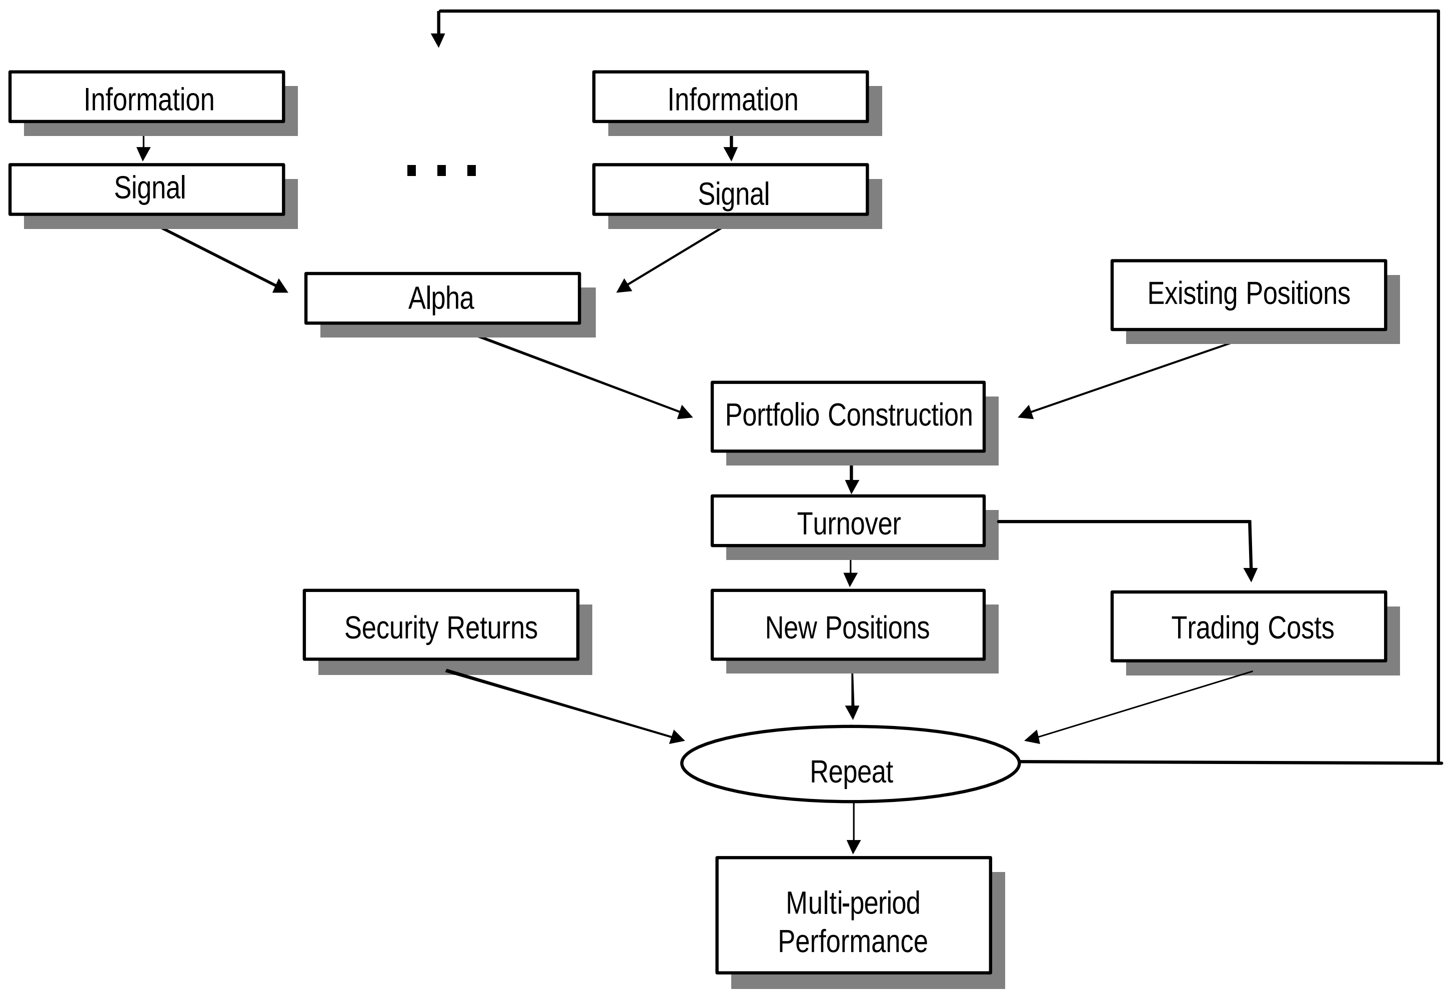

Dynamics and the Steady State

- ***Memory means that active portfolios are dynamic systems and so exhibit two regimes over time***
- ***Transient regime depends on starting condition, gives way to the steady state***
- ***Steady state regime lasts indefinitely and is independent of the starting condition***

Dynamics and the Steady State

Simulation and Solution

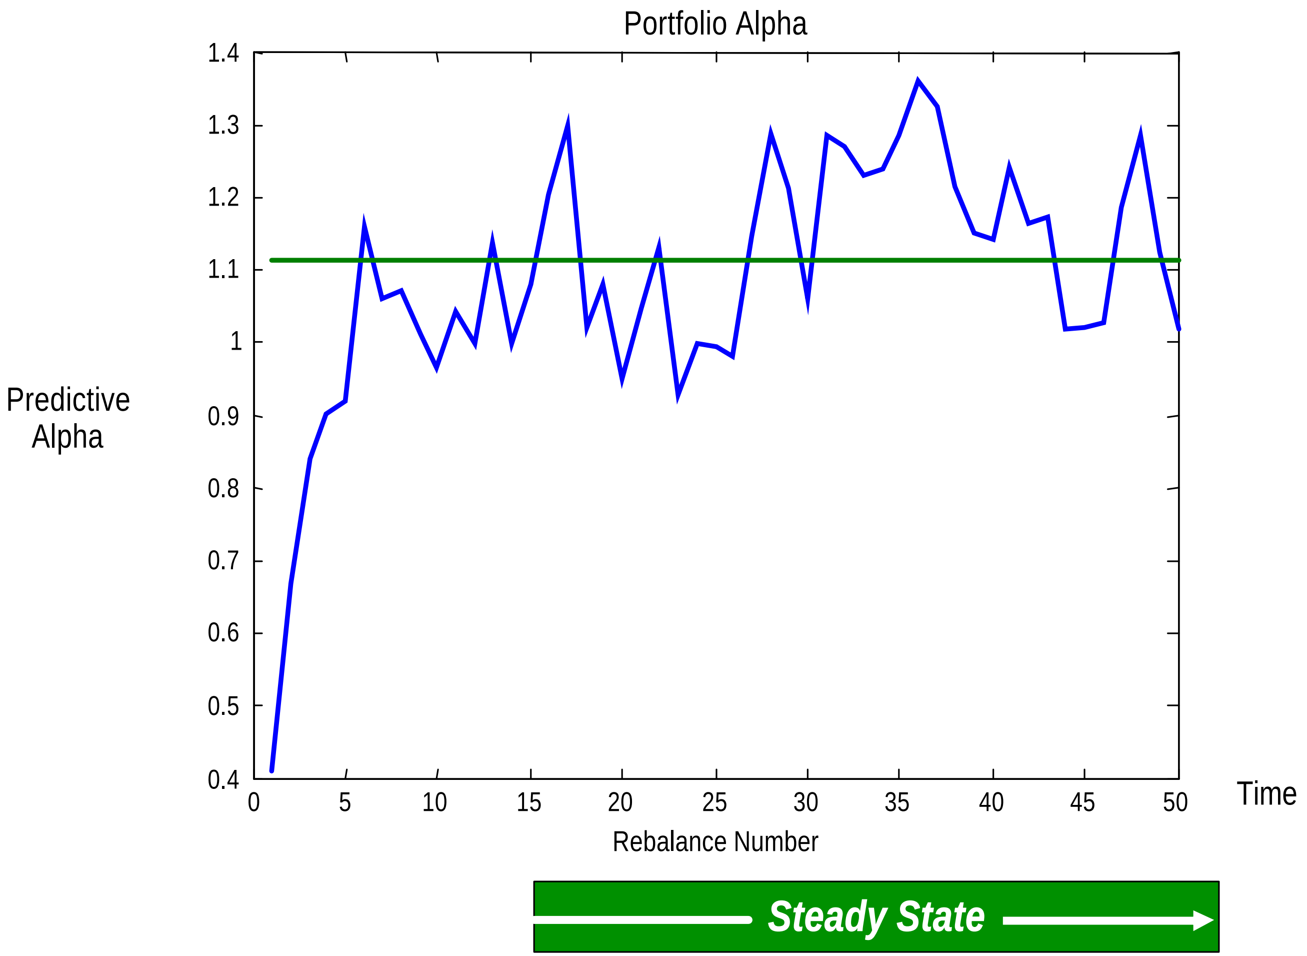

Solving the Model

Memory

- ***Signals are discounted sums of prior information***

$$a_j(t) = \sum_{t=0}^{\infty} f_j^{t} e_j(t - t)$$

- ***Active weights depend on prior alphas***

$$w(t) = hA\left[a(t) + Aa(t-1) + A^2 a(t-2) + \ldots\right]$$

The discount factor is $A \equiv (1 + 2hgV)^{-1}$

- ***Return is correlated with prior innovations***

$$IC_j = \sqrt{1 - f_j^2} \sum_{t=0}^{\infty} T_j(t) f_j^{t}$$

How Large are Active Weights?

- **Memory**

$$w(t) = hA\left[a(t) + Aa(t-1) + A^2 a(t-2) + \ldots\right]$$

- **Discount factor A and signal serial correlations f_j interact, giving**

$$\mathrm{Var}\{w\} = (hA)^2 \sum_j c_j^2 \frac{1 + f_j A}{(1 - f_j A)(1 - A^2)}$$

Solving the Model

Turnover and Tracking Variance

- Turnover depends on active weight size, and so looks similar

$$Turnover = f_R \sum_m \sqrt{(hA_m)^2 \sum_j c_j^2 \frac{4}{p} \frac{1-f_j}{(1-f_j A_m)(1+A_m)}}$$

- For uncorrelated assets, tracking variance is

$$TrackingVariance = \sum_m s_m^2 \mathrm{Var}\{w_m\}$$

Active Return and Trading Costs

- Active return includes multi-period predictive power, as well as the interactions between discount weight and serial correlation

$$ActiveReturn = \sqrt{f_R}\sum_m hA_m s_m \sum_j c_j \sqrt{1-f_j^2}\,\frac{f_j \Im_j\left(f_j\right) - A_m \Im_j\left(A_m\right)}{f_j - A_m}$$

- Portfolio trading cost is

$$TradingCost = tc\ TN$$

where tc is cost per unit of turnover, e.g. 30 bps

Solving the Model

Performance

- *We now have the key performance pieces: risk, return and trading costs*
- *For illustration, we use the after-cost information ratio*

$$IR \equiv \frac{ActiveReturn - TradingCost}{\sqrt{TrackingVariance}}$$

Active Position Size

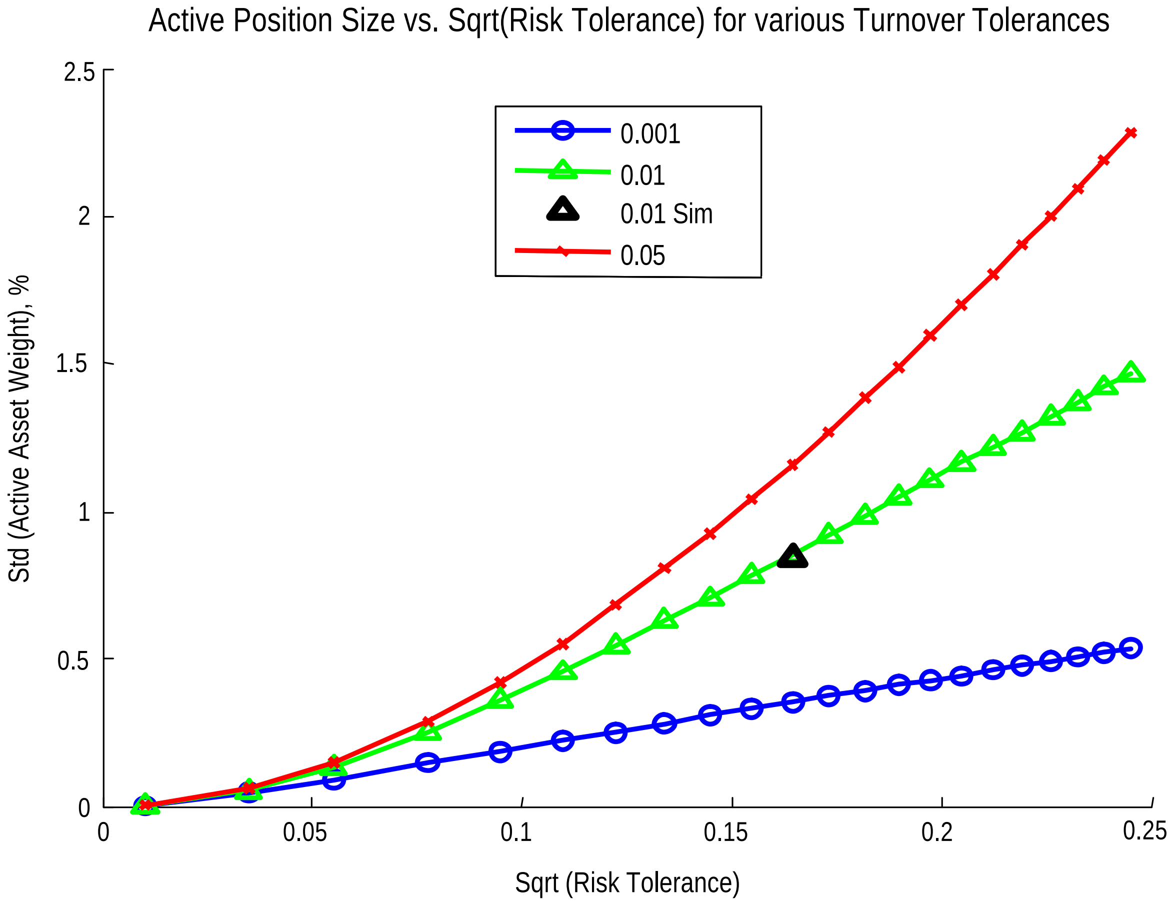

Solving the Model

Turnover

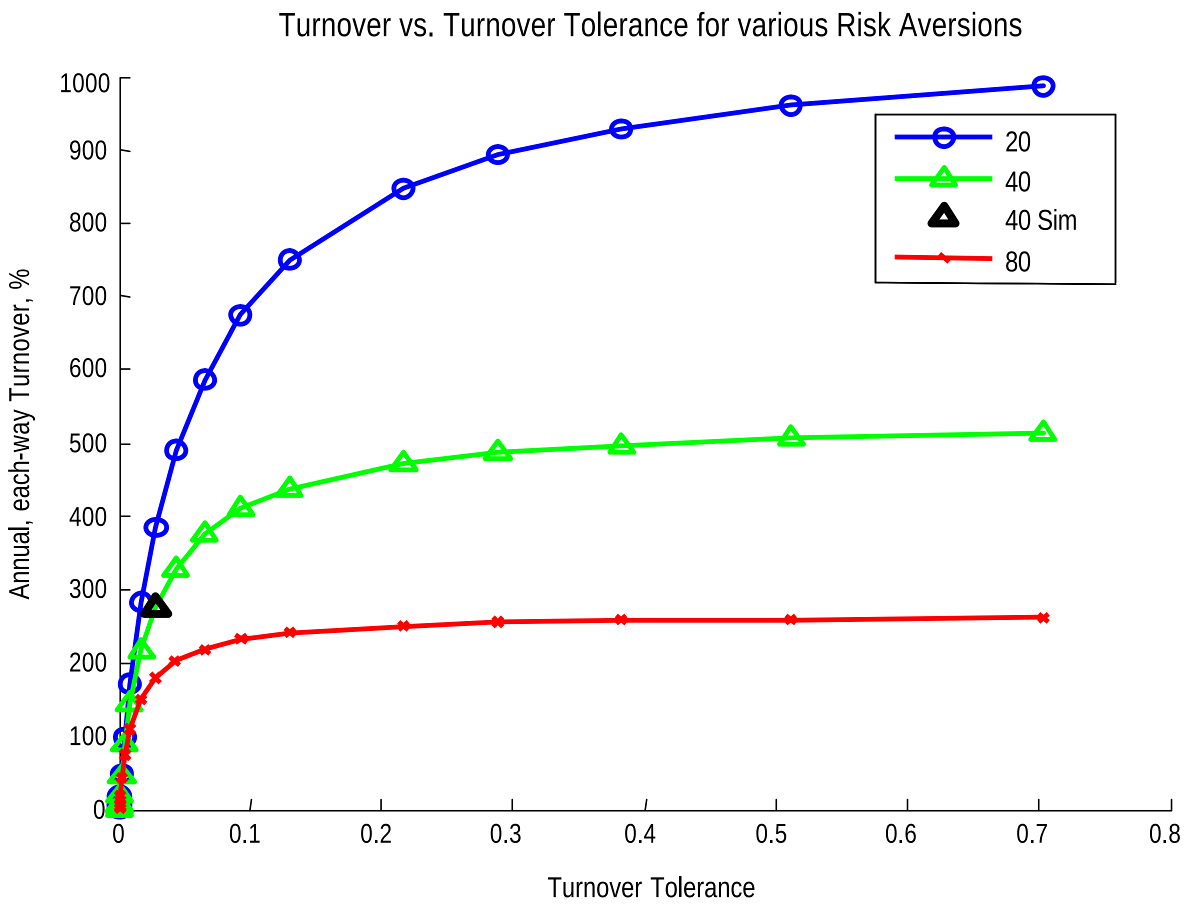

Portfolio Alpha

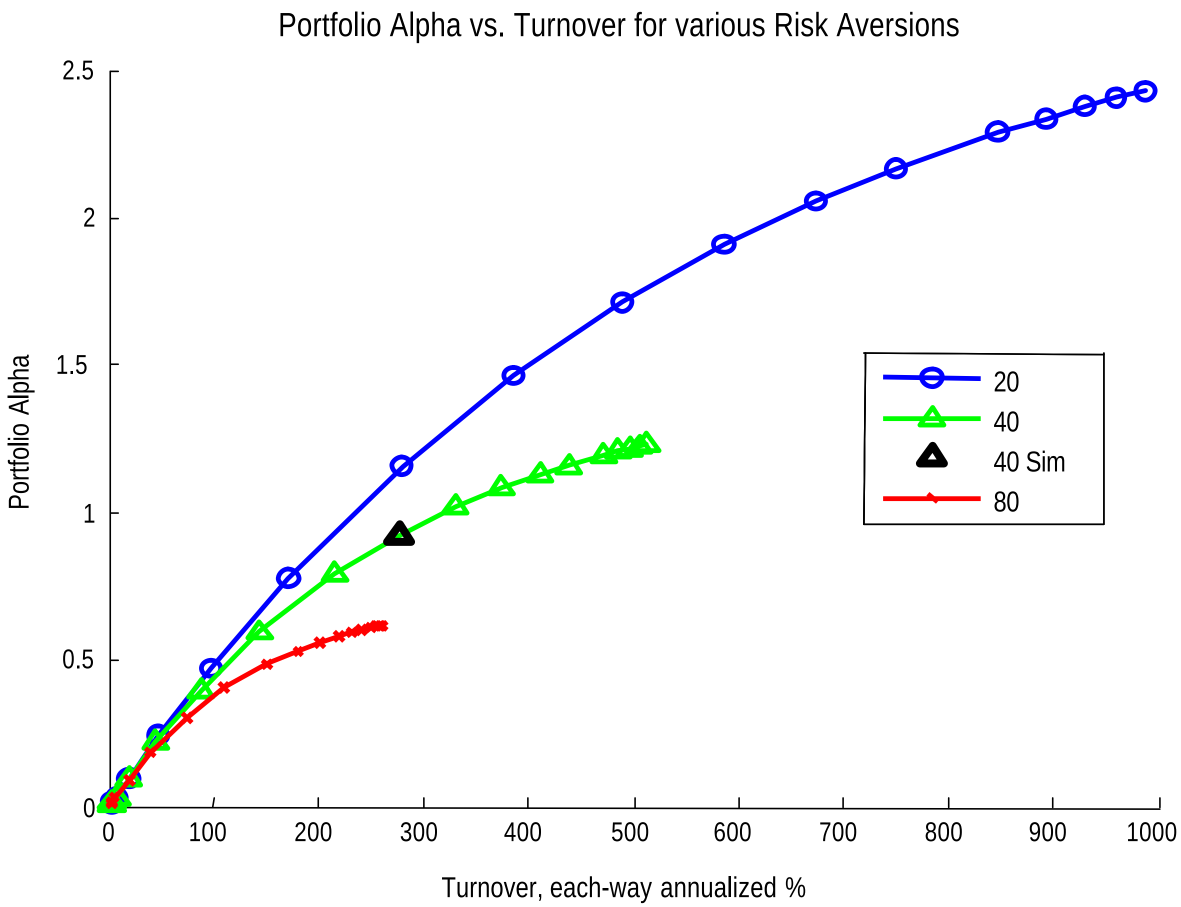

Performance Examples

Immediate and Postponed Predictive Power

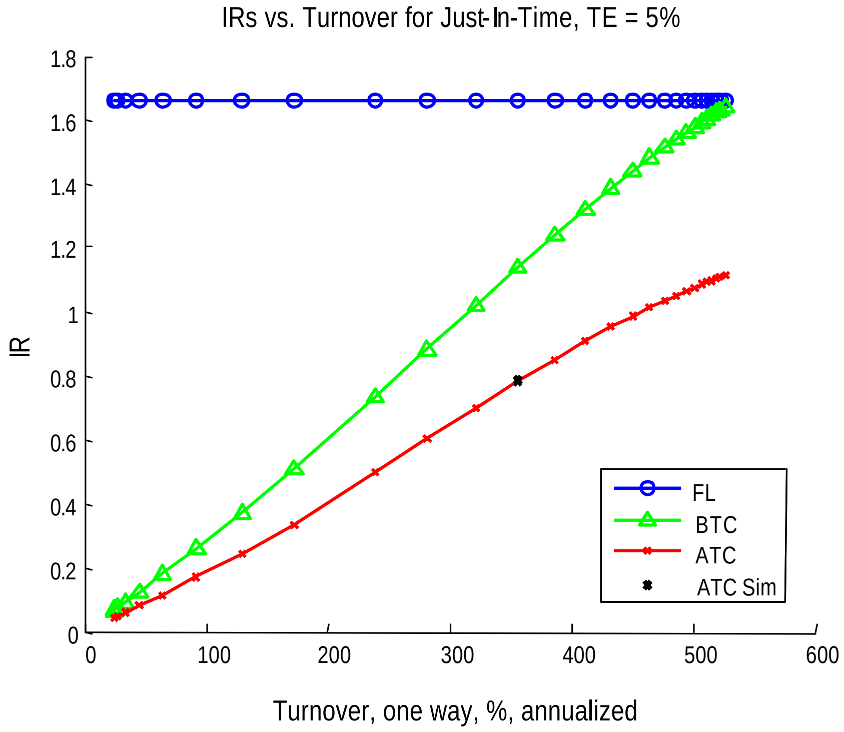

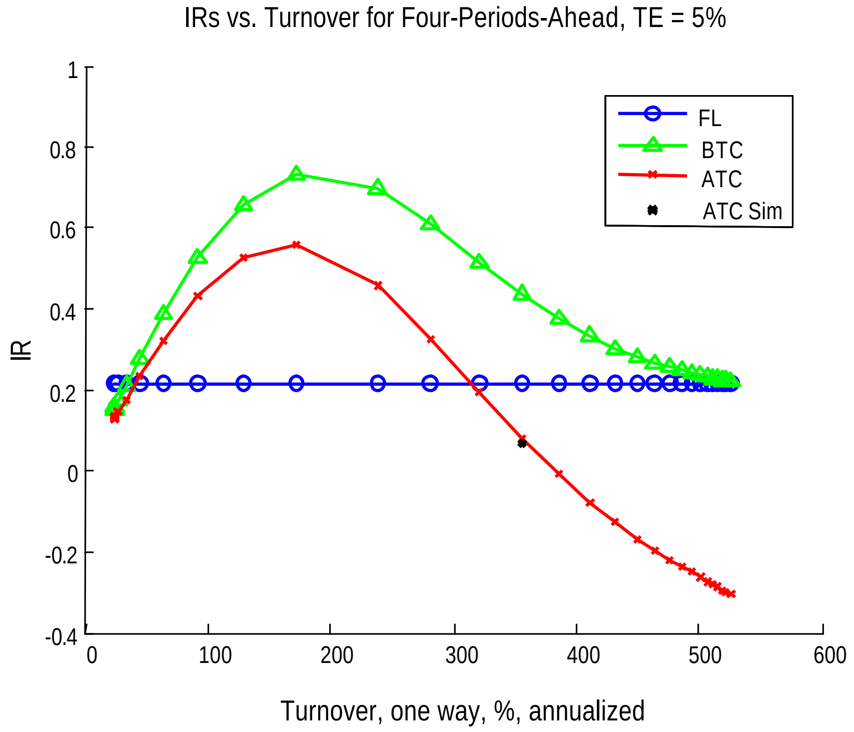

Performance Examples

Mixing the Hare and the Tortoise

The Fundamental Law

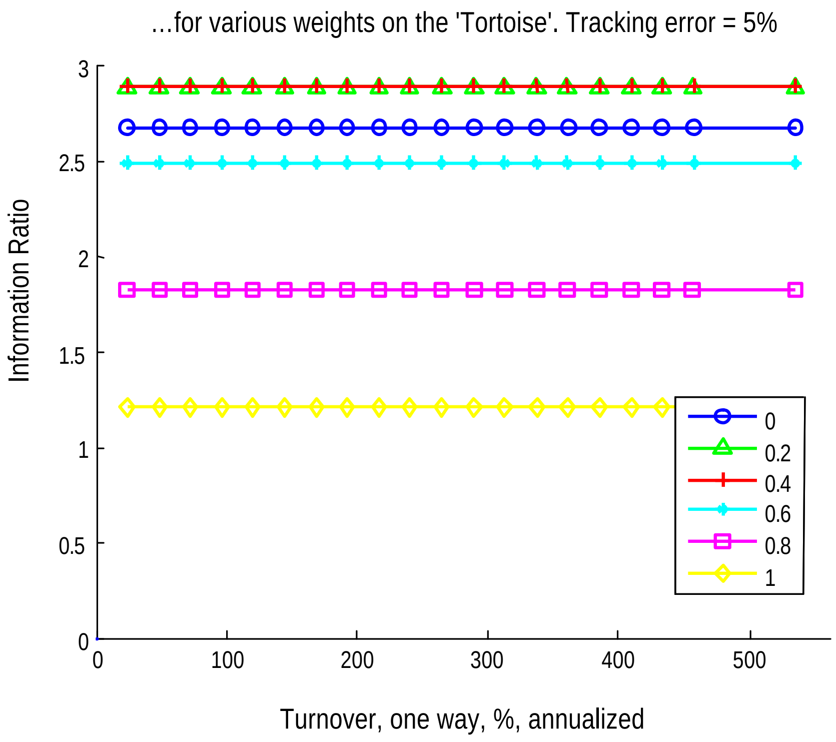

Multi-Period Performance

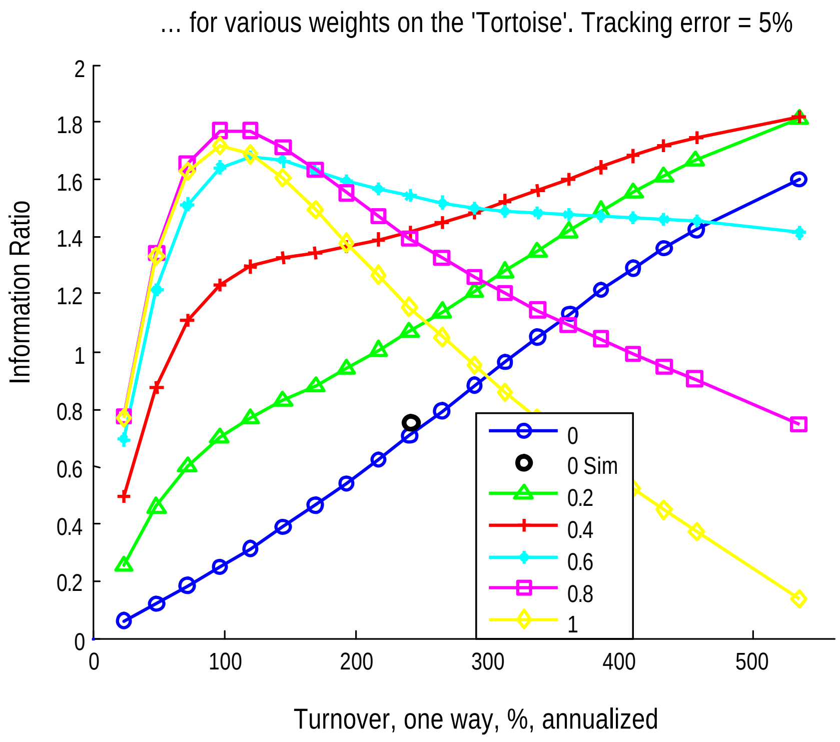

Conclusion

The time dimension of active management is an exploitable opportunity.

Thanks to . . .

***Nardin Baker, Dan DiBartolomeo, Doug Holmes,
Mark Hooker, Angelo Lobosco, John Minahan,
Larry Pohlman,Gita Rao, Easton Ragsdale,
Heydon Traub, and Westpeak colleagues.***