Model-Based Transition Management

A pragmatic approach applied to a multi-asset example.

Jan Bertus Molenkamp

Institutional investors like pension funds spend much time and effort on establishing the optimal strategic benchmark for their assets, given the characteristics of their liabilities. Once the strategic benchmark has been set, the fiduciary manager faces the challenge of an effective transition, reallocating all the assets to the new benchmark weights.

One can set a benchmark several ways. Often it is specified on a particular date, when in a blink of the eye the benchmark changes from (e.g.) 30% equities to 40% equities. For the manager a transition typically increases tracking error on top of the risk of the usual investment processes and raises the transaction costs. Both can have a significant influence on the performance, and possibly on the performance fee, of the fiduciary manager, depending on the specific fee arrangements.

This article researches possible strategies for the implementation of an effective transition. I describe a model to deal with the transition issues that is independent of the trigger of the transition (e.g., a rebalancing strategy or a dynamic allocation strategy). A multi-asset example demonstrates the practical applicability of the model.

ISSUES IN TRANSITION MANAGEMENT

How to execute a transition or portfolio revision has been little researched. If markets are perfect, and there are no transaction costs or liquidity constraints, the transition problem is trivial. In practice, though, markets are imperfect, and there are consequences in terms of return, risk, and costs.

From a strategic point of view the implementation of a transition may not seem particularly interesting, as the return consequences are relatively minor compared

Jan Bertus Molenkamp
is a researcher at the Vrije Universiteit Amsterdam in Amsterdam, The Netherlands, and a senior fiduciary manager at Kempen Capital Management.
jmolenkamp@feweb.vu.nl

to the average asset returns. If returns should fall short of expectation, as has been the case in the new millennium, the focus of pension fund trustees would automatically shift to costs, and transition costs become an issue, as they are not negligible.

Several observations may demonstrate the gap between theory and practice:

- Liquid instruments are not always available to obtain the desired allocation quickly. It takes time to build up exposure to, e.g., private equity.
- If there is high volume involved in transitioning a pension fund compared to the daily trading volume, the transition can distort market prices, which would influence the fund's own active and absolute performance. A market impact could occur if the transition date is a very illiquid trade date (such as December 31 or January 1).
- There can be a mismatch between the instruments used to implement the transition quickly (derivatives) and benchmarks, even in liquid markets. An example is that the futures on a major equity index can diverge from the index for several reasons (basis risk).
- There is often a mismatch between benchmark prices (closing end of day valuation) and implementation prices. That is, for example, the opening price on January 1 could be quite different from the closing price on December 31. If the transition cannot be implemented at the benchmark closing prices of December 31, but rather at the close of the first business day in the new year, one faces (historically) an average absolute return on equities of 0.4% and on bonds of 0.3%.[1]

In practice there are many vehicles for the transition of liquid asset categories at low cost, low risk, and low return impacts, but much in the asset mix of pension funds is illiquid.[2]

Exhibit 1 presents an example of a possible shift in strategic weights. It displays the changes in weight versus the previous benchmark. In this example the institutional investor moves away from pure equity risk and diversifies the asset mix. It is obvious that the transition out of liquid equity markets to more illiquid categories presents challenges other than purely an equity versus Treasuries transition.

So far we have little formal framework to deal with effective transition management. This makes it hard to explicate the way to execute a transition for two reasons. First, the overall consequences in terms of risk, return, and costs are not known and cannot be managed beforehand. So how could a pension fund evaluate the transition capabilities of its fiduciary manager? Second, there can be no clear view on who is bearing the risk involved in the transition and at what price. If a fiduciary manager takes on the transition risk, does it know whether it receives a fair price for doing so?

EXAMPLE ILLUSTRATING THE RISK IMPACT OF A TRANSITION

Let us examine some of the implications of a transition in an example. Say a fiduciary manager for a pension fund has a tracking error versus the pension fund benchmark of 0.5% in its normal investment processes. Let us assume a two-asset world and that 5% of the portfolio has to be reallocated from bonds to equities as a result of some change in the pension fund strategy, or merely as a result of rebalancing to 50% bonds and 50% equities. The benchmark for the manager changes on a specific date (e.g., January 1), and the manager wants to implement the transition during the two months following this transition date. The standard deviations are assumed to be 20% for equities and 8% for bonds. The correlation between equities and bonds and the active strategies of the manager is assumed to be 0.

What is the impact of the transition? For the pension fund, the strategically desired risk level resulting from the new benchmark is 11.2% as opposed to the old risk level of 10.5%. So at the transition date there is a 6% difference from the desired risk level. This risk, however, is taken on by

EXHIBIT 1

Example of Shift in Pension Fund's Strategic Asset Mix

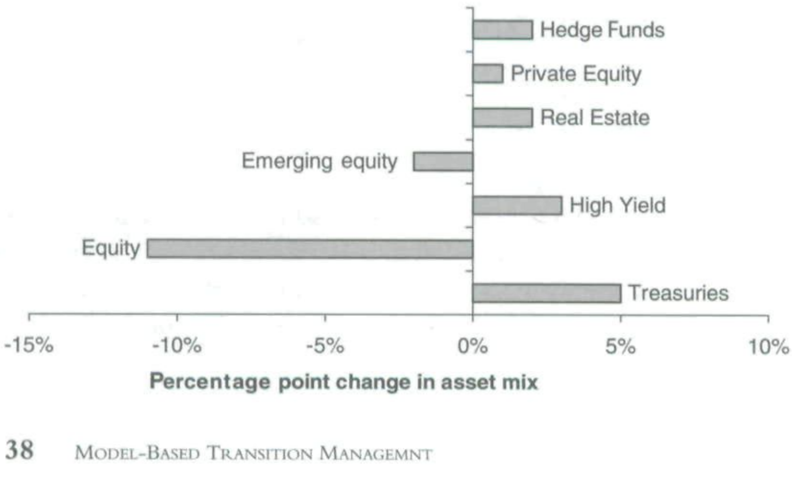

the fiduciary manager and will be reflected in the manager's active performance.

On the transition date the tracking error risk of the manager increases from 0.5% to 1.2%. This represents a 145% divergence from the normal tracking error level for the manager. The extra risk is borne only for two months.

Let us assume the manager implements this transition during the two months in a linear way, so each day assigns an equal part of the reallocation from bonds to equities, resulting in the strategically desired weights for the pension fund. During the two months the manager has on average a 50% higher tracking error. The effect on the average tracking error over one year is more modest, an increase of 10%.

Exhibit 2 displays a possible gradual transition to the benchmark weights, together with a reduction in transition risk. Note that the benchmark weight varies as a no-rebalancing or buy-and-hold strategy is assumed.

If the fund invests only in liquid bonds and equities, with liquid benchmarks, a transition can be executed efficiently in a short period with limited market impact and transaction costs and with low tracking error for the asset manager. If the strategic benchmark involves asset categories such as private equity or direct real estate, however, one cannot easily increase or reduce the weights in these categories in a short period of time. Transaction costs will be higher generally and manager tracking error vis-à-vis the strategic benchmark will increase, the longer the time.

If the benchmarks are illiquid, a proxy hedge could help to avoid an involuntary tracking error increase. If the private equity benchmark is statistically significantly related to a liquid (small-cap) equities benchmark, for example, this can be used to hedge away part of the extra risk.

It may not always be optimal to rebalance to exactly the strategic desirable weights. Mulvey and Simsek [2001], for example, assume it is optimal to rebalance to a zone around the target weights.

Leland [2000] shows that, given proportional costs and aversion to tracking error risk, the optimal strategy is to rebalance to the boundary of a no-trade zone. Leland offers two major observations. First, the effect of illiquidity is to increase the transaction costs. Second, given the no-trade zone, the average weight in illiquid assets tends to be less than the strategic weight.

Although the tracking error concept is difficult to apply to illiquid asset categories on a daily basis, one can use it to make estimates as to the effects of coping with illiquid assets. Leland's approach is to evaluate an optimal cost function representing all future costs (from trading or from risk). This function should satisfy the Hamilton-Jacobi-Bellman equation for all points in the no-trade region and boundary conditions for a finite number of boundary points. It is exact only in the case of one asset. For multiple assets, Leland introduces an approximation.

Donohue and Yip [2003] simulate the Leland approach in multiple-asset examples. They show one can improve on current rebalancing strategies, but the model is complex and not that easy to apply to real-life examples.

One obvious disadvantage of the Leland model is that, although it takes into account all future changes and costs, it keeps transaction costs and the variance of returns constant over time. As this assumption could easily be violated, the incorporation of future costs and tracking error could be unproductive.

Sun et al. [2005] propose a dynamic programming approach that copes with several utility functions, given cost and tracking error. One of their contributions is to analyze the effects of parameter uncertainty. Their optimal dynamic programming solutions compare favorably with other rebalancing approaches under parameter uncertainty in a two-asset case, but for multiple assets the approach is computationally intensive.

To deal with many assets, we propose a simplified pragmatic (one-period) model based on Leland [2000] that is easy to implement in a spreadsheet.

EXHIBIT 2

Example of Transition over Time

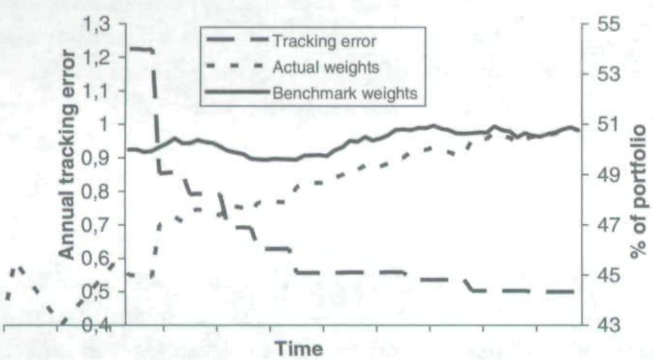

TRANSITION MANAGEMENT MODEL

Let w_i^0 denote the initial asset weight of asset category i, w_i^S the strategic weight of asset i, and w_i^* the optimal weight for asset i, where the weights sum to 100%. Let k_i denote the trading cost per asset i and μ_i denote the expected return on asset i. Let λ be the risk aversion parameter.

I formulate the model as a minimization of the utility function U:

$$
\begin{aligned}
\underset{w_1^*,\ldots,w_n^*}{Min}\; U_{w^*} = {} & \sum_{i=1}^{n} k_i \,|\, w_i^* - w_i^0 \,| \\
& + \lambda \sum_{i=1}^{n} \sum_{j=1;j\neq i}^{n} \{(w_i^S - w_i^*)^2 \sigma_i^2 + \sigma_{ij}(w_i^S - w_i^*)(w_j^S - w_j^*)\} \\
& + \lambda \left\{ \sum_{i=1}^{n} (w_i^S - w_i^*)\mu_i \right\}^2
\end{aligned} \tag{1}
$$

The cost k_i is multiplied by the absolute difference of the optimal weight versus the initial weight of asset i.[3] For risk we assume quadratic utility. Risk is calculated as the variance of the asset returns, given the difference between the strategic weights and the optimal weights plus the squared drift in the expectation of returns.[4]

The last term explicitly copes with the price dynamics. Normally this term will have limited influence, but when there are great differences in the expected returns of the different categories it will have a more significant influence. This is a parametric quadratic programming problem (PQP) that does not have a closed-form solution but can be tackled by several algorithms (even in a spreadsheet).

The model clearly presents a trade-off between the certain out-of-pocket transaction costs and the possible risk/return consequences involved in diverging from the strategic weights. In the analysis I assume immediate execution of the (optimal) transition.[5] The model horizon is explicitly embedded in the variance numbers and the expected returns. It is a specific discretization of the continuous Leland model.[6]

In the Leland model costs are discounted over an infinite horizon, while I focus on one period, not using any discounting, and analyze a quadratic utility function.

EXAMPLE: TWO RISKY ASSETS

Assume we have two risky assets, equities and bonds. The risk-free rate is assumed to be 3% and not volatile. The initial equity weight is 40%, and the strategic weight for next year is 50%. Given an annual standard deviation of 20% for equities and 10% for bonds, and assuming no correlation; an expected return for equities of 7% and for bonds of 4%; 50 basis point trading costs for equity and 20 bp for bonds; and a risk aversion parameter of 10, the optimal weight to rebalance equities to is 49.4%.

It is not optimal to buy equities beyond this percentage, as risk can be reduced at the expense of trading costs only in a less-than-optimal way. If one already has 49.5% in equities, it is optimal not to trade at all, so the 49.4% marks one of the frontiers of the no-trade region.

If the correlation between equities and bonds is not equal to 0 but 0.8 (for example), the optimal rebalancing weight for equities declines to 47.2%. If equities become more volatile, the optimal rebalance weight will increase. Higher costs will naturally reduce the optimal change from the initial weight; a lower risk aversion contributes to the same effect. A higher (positive) difference in expected returns between equities and bonds will also increase the optimal rebalance weight for equities.

The changes in the no-trade region as a result of changes in the parameters are very similar to the results in Donohue and Yip [2003]. Exhibit 3 depicts the no-trade regions of a two-asset case combined with a risk-free asset, with leverage allowed. The solid line represents the no-trade region in the base case. The strategic asset mix is 50% bonds and 50% equities.

In general, the no-trade region will increase (in area or volume) if the correlation shifts away from zero. When the correlation between bonds and equity is positive,

EXHIBIT 3
No-Trade Regions

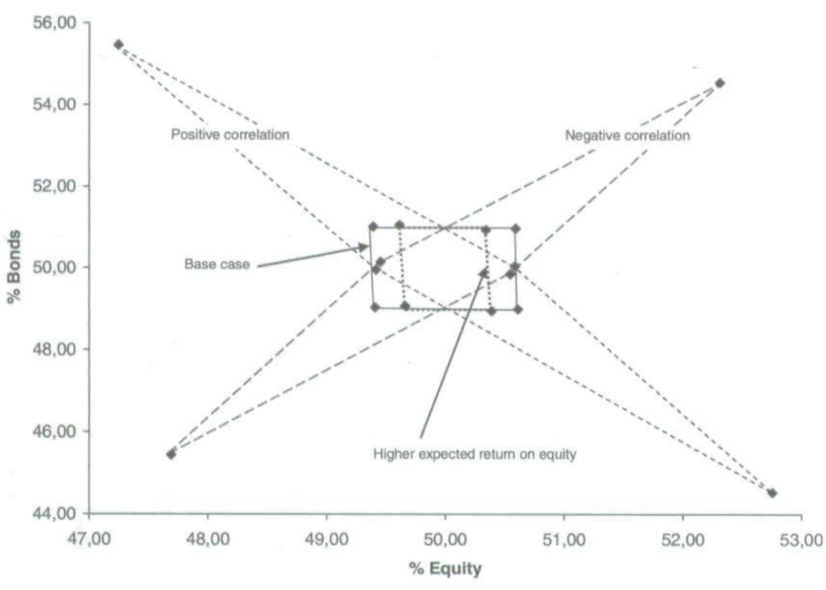

tracking error declines and the no-trade region shifts so that there is a negative slope between the extreme points (left-hand top, right-hand bottom). When the correlation is negative, it is the other way around.

A higher expected return on equity automatically reduces the no-trade region on the equity scale dimension only. Not shown in the graph is the no-trade region for an increase in the volatility of equities. This will reduce the no-trade region for equities further, but will leave the no-trade region for bonds unchanged.

SENSITIVITY ANALYSIS

We can perform some sensitivity analyses based on the two-asset example of the model. The analyses are based on perturbation of one parameter, keeping the other parameters constant. The vertical scales in Exhibits 4–8 are percentage point deviations from the strategic weight of equity and present a boundary of the no-trade zone.

Horizon

The horizontal axis in Exhibit 4 presents the time horizon in months. The extent of the no-trade region is reduced with the length of the time horizon, as the risk will increase. The specific shape of the curve indicates that risk increases with the square root of time.

With short horizons, the no-trade region can be quite substantial due to the dominance of the cost factor. At longer horizons, risk aversion is more dominant.

EXHIBIT 4
Effect of Time Horizon

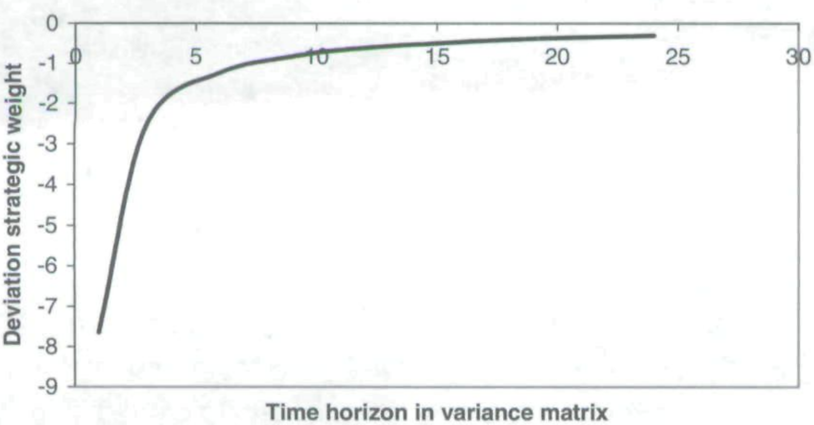

EXHIBIT 5
Effect of Correlation

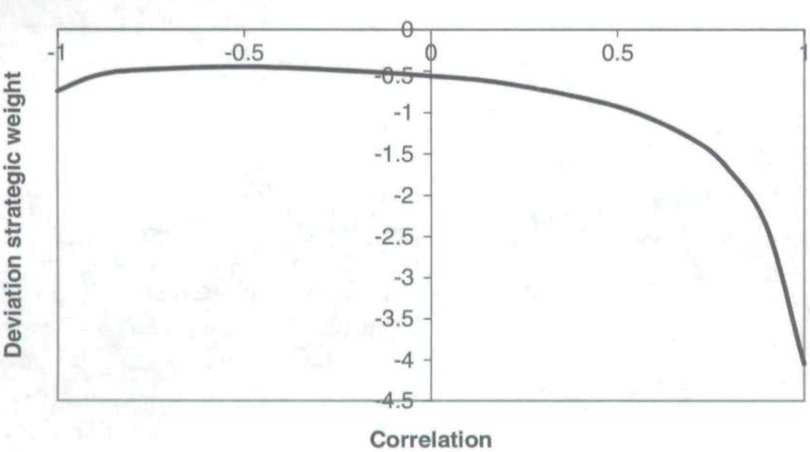

Correlation

The horizontal axis in Exhibit 5 presents the correlation. The no-trade region increases as the correlation between equity and bonds becomes relatively high, keeping the correlation with the risk-free asset constant at 0. Given the assumptions on costs, risk, and returns, the problem at hand is non-symmetrical.

It is important to remember that the volatilities of the asset categories are different. The curvature is influenced by the availability of the risk-free asset and the possibility of leveraging.

Risk Aversion

The deviation from the strategic weights declines with λ, as should be expected. Even at high values of λ, the optimal outcome remains on the boundary of a no-trade region. The horizontal axis in Exhibit 6 presents the risk aversion parameter.

Costs

The deviation from the strategic weight of equity is linear with the difference in cost between equities and bonds. The cost parameters in my examples have been chosen symmetrically (the costs of buying and selling are equal).

According to market microstructure theory, there are several ways to specify the cost function, e.g., correlation between size

EXHIBIT 6
Effect of Risk Aversion

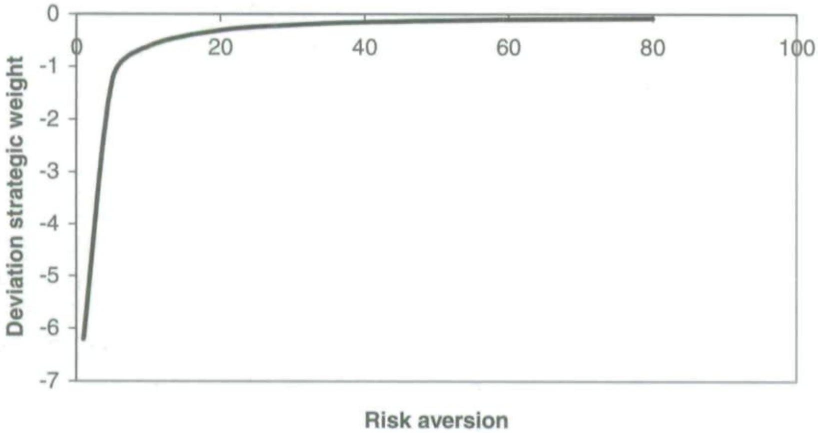

EXHIBIT 7
Effect of Volatility

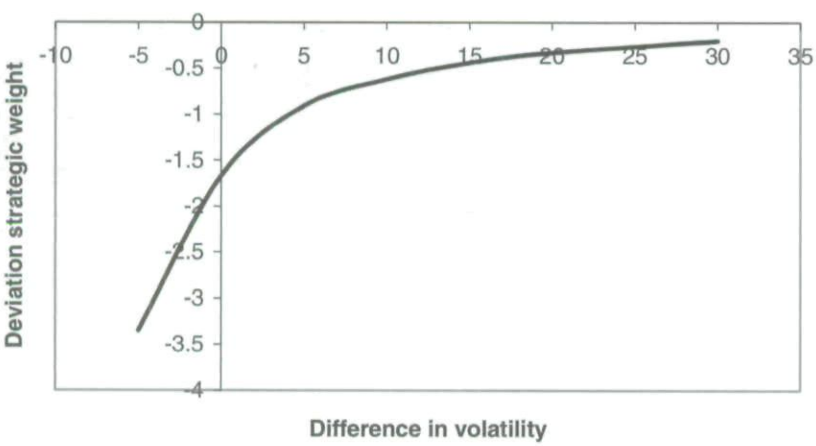

EXHIBIT 8
Effect of Risk Premium

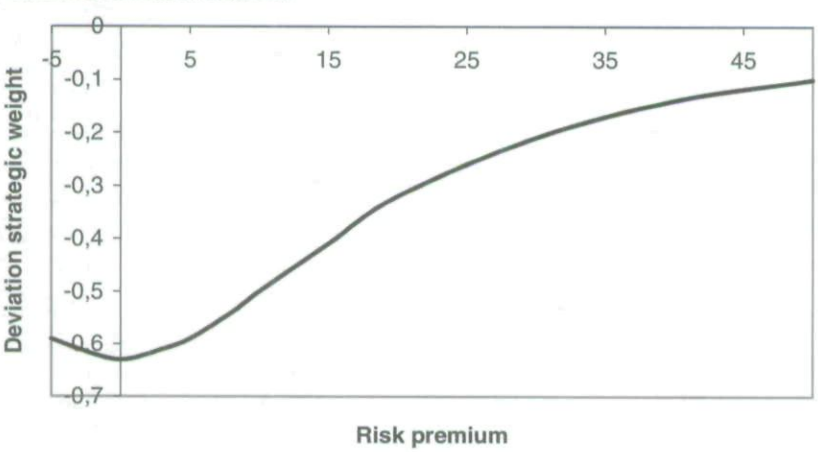

and costs (for an overview see, e.g., Madhavan [2002]). This can affect the form of the no-trade region.

If the cost function is concave, i.e., the marginal costs decline with volume, and if we have a risk-free asset to leverage or deleverage the portfolio, the no-trade region will shrink when the initial asset mix is farther away from the target benchmark. As costs become lower, the trade-off between risks and cost is bent more toward the risk factor. In the next periods, however, the no-trade region will grow as volumes decline and marginal costs are higher. In this case, a multiperiod approach would be necessary to establish an optimal solution.

If the cost function is convex, i.e., the marginal costs increase with volume, and if we allow for a risk-free asset, the no-trade region will grow when the initial asset mix is farther away from the target benchmark. In a later period we will be beyond the no-trade region, as it will shrink. Here also a multiperiod approach would be necessary to establish an optimal solution.

In practice, the cost curves will not be purely concave or convex, and there could be a floor or a cap on marginal costs. In general, depending on the specific parameters, optimal solutions can be obtained in a few periods.

If transaction costs are asymmetrical, i.e., the cost of buying is different from the cost of selling, the optimal solution will shift in the same way as with an increase or a decline in transaction cost per category. The direction of the optimal solution will not change except for extreme differences in the costs of buying and selling.

Volatility

What is the effect of changing the volatility of equities versus bonds, keeping all other assumptions constant? The horizontal axis in Exhibit 7 presents the percentage point difference in volatility of equities and bonds (base case = +10 percentage points). As one would expect, a decline in the volatility of equities reduces the risk of a deviation from the benchmark for equities and widens the no-trade zone for equities (and also for bonds).

The expansion for bonds will be smaller than for equities for two reasons. First, the total risk will be, more than before, dominated by the risk of deviations from the bond benchmark. Second, the risk-free asset will come into play as its cost compares favorably with the cost of the lower risk-inducing equities.

EXHIBIT 9

Initial Asset Mix and Benchmark

	Initial weights	Benchmark weights
Fixed Income	50%	45%
High Yield		5%
Equity Capital Developed	35%	25%
Equity Emerging		5%
Direct Real Estate	15%	10%
Indirect Real Estate		5%
Private Equity		5%
Cash		0%
Total	100%	100%

EXHIBIT 10

Assumptions on Asset Categories

Category	Expected return	Volatility	Trading cost (basis points)	Initial weights	Benchmark weights
Fixed Income	4%	5%	20	50%	45%
High Yield	8%	20%	100	0%	5%
Equity Capital Developed	8%	20%	50	35%	25%
Equity Emerging	10%	34%	150	0%	5%
Direct Real Estate	5%	6%	200	15%	10%
Indirect Real Estate	8%	24%	120	0%	5%
Private Equity	10%	33%	180	0%	5%
Cash	2%	2%	0	0%	0%

Risk Premium/Return Drift

In the transition model the main trade-off is between risk and costs. If there is a great difference in expected returns, this will have an impact on the optimal decision.[7] The effect of differences in the risk premiums is linear within λ, the risk aversion parameter. So if λ doubles, the deviation from the strategic benchmark is cut in half.

The horizontal axis in Exhibit 8 presents the risk premium of equities over bonds. If the expected return of equities increases, the no-trade zone (for equities) will shrink. The effect is non-linear corresponding to the formulation in the model. Note that if the drift term is not included in the formulation, Exhibit 8 would show a straight horizontal line at the level where the risk premium equals 0.

TRANSITION MANAGEMENT MODEL APPLIED TO MANY ASSETS

The interesting feature of the model is its applicability to multiple assets. Let us assume a pension fund with 1 billion USD assets under management. The pension fund wants to diversify away from the classic fixed-income, equities, and real estate portfolio toward categories such as high-yield, emerging equity, indirect real estate, and private equity as shown in Exhibit 9.

Exhibits 10 and 11 show the assumptions and the covariance matrix for all the asset categories.[8] We assume a risk aversion of 30. The optimal solution is in Exhibit 12.

There are some interesting observations we can make in Exhibit 12.

- All costs, correlations, volatilities, and risk premiums are taken into account.
- The optimal weights are on the boundaries of the no-trade zone.
- The no-trade zone is relatively large compared to its size for the two-asset example. The extension to more asset categories (and a broader benchmark) automatically creates a larger no-trade zone. An extra effect is added by the relatively high costs of the extra, more illiquid, asset categories.
- The weight of direct real estate has not changed. This is because of the high transaction costs and because, given the assumptions, direct real estate is more efficient (return/risk profile) than fixed-income.

EXHIBIT 11
Covariance Matrix

Category	Fixed Income	High Yield	Equity developed	Equity Emerging	Direct Real Estate	Indirect Real Estate	Private Equity	Cash
Fixed Income	1	0.44	0.08	-0.17	0.16	0.08	0.12	0.54
High Yield	0.44	1	0.66	0.54	0.62	0.51	0.59	0.25
Equity Developed	0.08	0.66	1	0.61	0.65	0.7	0.77	-0.08
Equity Emerging	-0.17	0.54	0.61	1	0.51	0.51	0.31	-0.06
Direct Real Estate	0.16	0.62	0.65	0.51	1	0.74	0.33	-0.09
Indirect Real Estate	0.08	0.51	0.7	0.51	0.74	1	0.45	0.07
Private Equity	0.12	0.59	0.77	0.31	0.33	0.45	1	0.1
Cash	0.54	0.25	-0.08	-0.06	-0.09	0.07	0.1	1

EXHIBIT 12
Solution

	Optimal weights	Change in weights	Deviation from benchmark weight	Initial weights	Benchmark weights
Fixed Income	41.4%	-8.6%	-3.6%	50.0%	45%
High Yield	4.7%	4.7%	-0.3%	0.0%	5%
Equity Developed	26.0%	-9.0%	1.0%	35.0%	25%
Equity Emerging	4.4%	4.4%	-0.6%	0.0%	5%
Direct Real Estate	15.0%	0.0%	5.0%	15.0%	10%
Indirect Real Estate	4.0%	4.0%	-1.0%	0.0%	5%
Private Equity	4.5%	4.5%	-0.5%	0.0%	5%
Cash	0.0%	0.0%	0.0%	0.0%	0%

- The fixed-income weight falls to below the benchmark weight for the same reason.
- The allocation to private equity is higher than the allocation to indirect real estate. Although the transaction costs for private equity are higher than for indirect real estate, private equity is much more volatile than indirect real estate, producing a larger risk penalty.

Only 70% of the trading costs (44 basis points) are incurred in moving toward the benchmark weights. That implies a trading cost saving of 30% (13 basis points) or U.S.$1.3 million. The resulting (temporary) tracking error is 0.3%. If this is seen to be too high, the risk aversion parameter can be increased. In general, however, if a tracking error of this magnitude is mixed with the tracking error due to active management it does not significantly increase the overall tracking error. Of course, by not trading to the benchmark allocation (allowing for a no-trade zone), the resulting portfolio will entail active bets with an expected return distribution.

The major difference with the two-asset model is that there are far more possible trade-offs among volatilities, correlations, costs, and returns. The essence of the model has not changed, and its great advantage over other approaches is that the results can be calculated easily.

IMPLEMENTATION ISSUES

The model presented here assumes *immediate execution* of the rebalancing and does not take into account all the practical difficulties I have noted. It does, however, point out in agreement with Leland [2000] that it could be unwise to rebalance to exactly the strategic weights.

Optimization of the utility function could present some problems for optimizers, as the absolute value in the cost part of the function is not smooth. This could typically be the case in higher dimensions. This can be remedied by squaring the term and, of course, adjusting the lambda accordingly.[9]

The problem at hand is in itself a dynamic problem. As soon as the optimal solution has been determined, the conditions under which it was derived have changed. The question is whether the initial optimal solution is still in the neighborhood of the "new" optimal solution. I have therefore tested the robustness of the model for shifts in the parameters. If the starting weights change as a result of market movements, the optimal solution (no-trade zones) still holds. If the weights have moved into the no-trade zone, the response is (of course) not to trade.

Overall, with small changes in the parameters (standard deviations, correlations), the initial solution lies near the (new) optimal solution. From Exhibit 5 we know that a measurement error of the correlation can have a significant impact, especially when (absolute) correlation is relatively high.[10] It is assumed that risk aversion and trading costs are constant. As a consequence, these parameters do not contribute to the dynamics.

I have experimented with an extension of the model to make it a multiperiod model. There can be several motivations for this. For example, one could already know the specification of the benchmark in a future period (whether statically or dynamically specified). It is also one way to cope with the illiquidity of some categories, e.g., to restrict trading volume during a period.

In essence, the outcomes of a multiperiod model and a one-period model will be much the same, but a multiperiod model can certainly accommodate extra information and thereby improve on optimality.

CONCLUDING REMARKS

An integrated transition for a pension fund to change its asset mix policy can involve costs, risks, and returns. Given the return and cost effects of transitions, it is wise for pension fund managers to specify how to deal with the execution of a transition in terms of risk. I simplify a model that trades off risks and costs to produce a no-trade zone around the benchmark weights. It is easy to solve even for many assets as we have seen in some examples. The model is easy to implement and can be extended to cope with other constraints or goals.

ENDNOTES

The author thanks Cees Dert, Andre Lucas, Arjen Siegmann, and the reviewer for their comments. The views expressed here are those of the author and should not be taken to represent the views of his employers.

[1]Based on the MSCI World index and the JPM Global Broad bond index over 1989-2002.

[2]Transition managers do not assume the risks, and there is a question as to whether their interests are really aligned with their clients' interests.

[3]At this point I assume symmetric trading costs. Later on I relax this assumption.

[4]To formulate the risk in the utility function, I evaluate the expectation:

$$E\left\{\sum_{i=1}^{n}(w_i^S - w_i^*)r_i)\right\}^2$$

where r_i is the expected return for asset category i. The expectation of a quadratic term can be replaced by its variance plus the square of the expectation, yielding the formula result.

[5]There is literature on execution strategies for trading blocks, where trading volume has a price impact and where dynamic programming approaches have been applied. See, e.g., Bertsimas and Lo [1998].

[6]In a way the one-period model can be viewed as a two-period model if at the end of the period the weights will be changed to the strategic weights at no cost. A more practical interpretation is that at the end of the model horizon a new change in the benchmark will occur.

[7]If one has clear views on asset allocation that are not described in the benchmark, these should be treated separately.

[8]The expected return, volatility, and covariance data are for purposes of illustration only. They have been constructed so to make them as consistent and realistic as possible. The trading costs are estimates and include market impact (see, e.g., Wagner and Edwards [1998]).

[9]My thanks to the reviewer for pointing this out.

[10]This can be concluded from the steepness of the curves. At the same time, this is a more general observation for these kinds of quadratic programming models. The Leland approach also suffers from this parameter uncertainty. Yet even with considerable uncertainty in the parameters, the models still help us understand and analyze the way to approach a transition.

REFERENCES

Bertsimas, D. and A. Lo. "Optimal Control of Execution Costs." *Journal of Financial Markets*, Volume 1 (1998), pp. 1-50.

Donohue, C., and K. Yip. "Optimal Portfolio Rebalancing with Transaction Costs." *The Journal of Portfolio Management*, Summer 2003, pp. 49-63.

Leland, H. E. "Optimal Portfolio Implementation with Transactions Costs and Capital Gains Taxes." Haas School of Business, University of California, Berkeley, December 20, 2000.

Madhavan, A. "Market Microstructure: A Practitioner's Guide." *Financial Analysts Journal*, September/October 2002, pp. 28-42.

Mulvey, J. M., and K. Simsek. "Rebalancing Strategies for Long-Term Investors." EDHEC Risk and Management Research Centre, March 20, 2001.

Sun, W., A. Fan, L. W. Chen, T. Schouwenaars, and M. Albota. "Optimal Rebalancing Strategy Using Dynamic Programming for Institutional Portfolios." Working paper, MIT, 2005.

Wagner, W., and M. Edwards. "Implementing Investment Strategies: The Art and Science of Investing." In F. J. Fabozzi, ed., *Active Equity Portfolio Management*. 1998.