#3 - 2012

NBIM DISCUSSION NOTE

This paper was part of the NBIM memo "On rebalancing" (February 2012)

Empirical Analysis of Rebalancing Strategies

30/03/2012

We review the theoretical foundation for rebalancing regimes and look at the impact of rebalancing on the portfolio's risk and return based on historical return data from 1970 to 2011. We compare both different calendar based rebalancing regimes and different threshold based regimes with the performance of a drifting mix portfolio. Towards the end of the note we focus on the specific design of a rebalancing regime. In particular, we look at a trigger based regime and examine different designs of the threshold level, persistence requirement and implementation rule.

Main findings

- Rebalanced portfolios have had both higher returns and lower risk (volatility and tail risk) than a passive, drifting portfolio. Even if one could argue from a theoretical standpoint that rebalancing involves taking on additional (contrarian) risk, this risk has not materialised in our sample period 1970-2011.

- There are relatively small differences in the risk/return profile between different specifications of the rebalancing regime. The decision to rebalance the portfolio has been far more important than the specifics of the regime.

- The current rebalancing regime for the Government Pension Fund Global is one of the best-performing rebalancing strategies over our sample period. Regional rebalancing has contributed to a higher Sharpe ratio over our sample period. However, the theoretical foundation for regional rebalancing is weak. Rebalancing around any arbitrary dimension such as regions or sectors may add value in a historical analysis if they exhibit mean-reverting tendencies in returns.

- New portfolio theory points out that an investor's portfolio should be constructed by allocating to a wide range of different risk premia, where the optimal exposure to the different risk premia is investor-specific, depending on preferences and the nature of any non-tradeable risks that the investor is exposed to. Constructing a rebalancing regime around the equity share in the portfolio will effectively manage the exposure to the equity risk premium, which is the single most important risk factor.

- We examine a threshold based rebalancing regime focused on keeping the equity share within 3 percentage points of the strategic weight. We also examine the implications of different persistence requirements and compare the impact of different rules for how the index should be adjusted towards the strategic weight after the trigger for rebalancing is met.

Introduction

In the first section of this note we review the theoretical foundation for rebalancing regimes before we look at the impact of rebalancing on the portfolio's risk and return based on historical return data from 1970 to 2011. We compare both different calendar based rebalancing regimes and different threshold based regimes with the performance of a drifting mix portfolio.

In the second part of the note we focus on the specific design of a rebalancing regime conducting both ex ante Monte Carlo simulations and ex post analysis based on historical returns. We look at a trigger based regime and examine different designs of the threshold level, persistence requirement and implementation rule. These design questions balance the concerns of, on the one hand, keeping the deviations from the strategic weights as small as possible and, on the other hand, keeping transaction costs in connection with rebalancing low.

Why rebalance – Traditional portfolio theory

Modern portfolio theory is often thought to have started with the mean-variance analysis of Markowitz (1952). Markowitz showed how investors should pick assets if they care only about the mean and variance of portfolio returns over a single period. A striking implication of this analysis is that all investors will hold the same portfolio of risky assets, representing the unique best mix of stocks and bonds often called the "tangency portfolio". Conservative investors will combine the risky portfolio with cash to achieve a portfolio with lower standard deviation and lower expected returns, whereas more aggressive investors will hold less cash and may even borrow to leverage their holdings of the risky portfolio. The composition of the risky portfolio will, however, be identical for both investors. This result is the mutual fund theorem of Tobin (1958). Whereas Markowitz assumed that investors cared directly about the mean and variance of portfolio returns, later research in portfolio theory derived the same results assuming that the investor had a utility function defined over wealth at the end of one period.

There are several reasons why a long-term investor should hold a different portfolio to a short-term investor. In two seminal papers, Merton (1969) and Samuelson (1969) investigated the conditions under which long-term investors should invest myopically, choosing the same portfolios as short-term investors. They showed that myopic portfolio choice is optimal if asset returns are independent and identically distributed over time or if the utility function takes the log form. Assuming that the returns are independent and identically distributed over time means that the investor faces a constant investment opportunity set. Campbell and Viceira (2002) provide a broader and deeper discussion of the Merton-Samuelson conditions, and the horizon effects on portfolio choice that arise if these conditions are violated. However, due to the difficulty of solving intertemporal portfolio models under more general and less restrictive conditions, the simple models where the Merton-Samuelson conditions hold dominate the practical applications of portfolio theory.

The assumption of constant risk premia implies that optimal portfolios are constant over time for both short-term and long-term investors. Hence, under the standard assumptions mentioned above, an investor who maximises the expected utility of wealth at the end of the investment horizon should hold a portfolio with constant portfolio weights. The rationale for rebalancing a portfolio follows from this: If the optimal portfolio is constant over time, the actual portfolio will have to be rebalanced back to the optimal weights whenever market movements create a discrepancy between the actual and the optimal portfolio.

In a world without transaction costs, the portfolio would be continuously rebalanced and be identical to the optimal portfolio at all times. With transaction costs, however, an investor will weigh the utility loss of deviating from the optimal asset allocation against the cost of trading back to the optimal allocation. Research[1] has pointed out that, in the presence of transaction costs, there will be a no-trade region

1 "Literature focusing on optimal rebalancing in the presence of transaction costs include Constantinides (1986), Davis and Norman (1990), Schroder (1995), Leland (1999), Zakamouline (2002), Donohue and Yip (2003), Kritzman, Myrgren and Page (2007)."

around the target asset allocation. As long as asset weights remain inside this no-trade region, no trading will occur. However, if an asset class weight moves outside the no-trade region, trading will take place in order to bring all asset weights back towards the target allocation. Whether it is optimal to trade all the way back to the target allocation or back to the nearest edge of the no-trade region depends on the investor's cost function. With proportional costs only, the optimal rule would be to trade back to the nearest edge only. If there are non-zero fixed costs in addition to the proportional costs, it will be optimal to trade back to a point inside the no-trade region. The optimal size of the no-trade region and the optimal frequency of rebalancing will depend on transaction costs and tolerance for deviations from target.

Limitations of the constant mix

The design of the fund's current rebalancing regime is anchored in traditional portfolio theory which rests on assumptions that have been challenged in practice. In addition, new portfolio theory developed over the last decade is changing our thinking about dynamic portfolio choice. Below, we examine different limitations of the constant-mix approach.

- **Macro consistency**
  The traditional model for portfolio choice outlined above is solving one investor's problem given an exogenous return-generating process. However, if all investors in the marketplace followed this approach, all investors would buy and sell the same assets at the same time. Clearly, this is not macro-consistent. In general equilibrium, for every investor rebalancing his portfolio there has to be someone else taking the opposite side of the rebalancing trade. The same point has been made by Cochrane (2011), who advocates a portfolio theory where investor heterogeneity leads to different investors holding different portfolios depending on how their characteristics differ from the average investor. Contrarian rebalancers and trend-following investors can then live side by side, ensuring market equilibrium at all times.

- **Time-varying risk premia**
  The assumption of constant risk premia implies that the composition of an optimal portfolio is constant over time for both short-term and long-term investors. However, a large literature finds that expected asset returns seem to vary so that investment opportunities are not constant. The evidence of predictability in various asset classes is surveyed in a separate note. The fact that risk premia are time-varying generates time variation in optimal portfolios: Both short-term and long-term investors should seek to "time the markets", holding more risky assets when the rewards for doing so are high. Campbell and Viceira (1999) show that there are large utility losses from holding a constant-mix portfolio when risk premia are time-varying.

- **Multiple risk factors, multiple agents**
  New portfolio theory expands on the traditional theory by pointing out that an investor's portfolio should be constructed by allocating to a wide range of different risk premia. In the traditional framework, the only source of risk is market risk: All the investor has to do is to allocate his wealth between the risky market portfolio and the riskless asset. Financial research over the last 30 years has, however, documented a wide range of systematic risk premia, including "value", "size", "momentum", "carry" and "volatility", in addition to market risk. In new portfolio theory, therefore, the market is populated with agents who have different preferences and non-tradeable risks, leading to different desired exposures to the various risk factors. Different hedging demands due to investor heterogeneity ensure a market equilibrium where some investors buy "value" or "small cap", whereas others sell exposure to the same risk factors. Hence new portfolio theory has at least two important implications for rebalancing. First, rather than focusing on allocation to traditional asset classes, the focus should be on allocating to various systematic risk premia. Second, the optimal exposure to the different risk premia is investor-specific, depending on preferences and the nature of any non-tradeable risks that the investor is exposed to.

Impact of rebalancing

In this section, we look at the impact of rebalancing on the portfolio's risk and return based on historical return data from 1970 to 2011. We compare the performance and risk of different rebalancing strategies to the performance of a passive portfolio where the asset class mix and regional allocation drift with market developments.

Data

There is a trade-off between data resolution and the availability of data series with a sufficiently long history. In this study, we focus on monthly data for three regions (America, Europe and Asia/Pacific) for both equities and bonds. This level of data granularity puts limits on how much history is available. We have constructed a return history going back to January 1970 using the different data sources described below. Although we could have wished for even more data history, the period from January 1970 covers a number of different market regimes and shocks to the financial markets and the world economy, including the oil price shock in 1973, the stock market crashes in 1987 and 2008, and the IT bubble around the turn of the millennium.

We use FTSE All Cap index data for equity returns in the three different regions back to September 2003. From September 2003 to January 1986 we use the FTSE World index, and from December 1985 and back to January 1970 we use the MSCI World index for the three respective regions.

We use the Barclays Global Aggregate as a total return index for the global fixed-income markets. For the European fixed-income markets we use the Barclays Pan-European Aggregate back to February 1999. From January 1999 and back to January 1970 we construct a total return index based on total return indices for ten-year government bonds in Germany, France and the UK. The three different countries are equally weighted. This approach serves as a good approximation of the return on the Barclays Pan-European Aggregate when we compare the two indexes in the period where we have overlapping observations. The total return indices for ten-year government bonds in the three countries are sourced from Global Financial Data. For America we use the Barclays US Aggregate back to May 1981, and then the total return on ten-year Treasuries from Global Financial Data back to January 1970. For the Asia/Pacific region we use the Barclays Asian-Pacific Aggregate back to July 2000, and then the total return on ten-year Japanese government bonds back to January 1970. The return on ten-year Japanese government bonds gives a very good approximation of the Barclays Asian-Pacific index when we compare the two indices in the period where we have overlapping observations.

Passive portfolio

The natural reference point for an evaluation of different rebalancing strategies will be the performance of a passive portfolio initially invested according to the market capitalisation of the different asset classes. A first challenge arises from estimating a broad market portfolio. Challenges related to estimating the broad market portfolio have been well-known for decades (Roll 1977). In this study, we use the market capitalisation of the FTSE Global All Cap equity index and the Barclays Global Aggregate fixed-income index as our estimate of the market portfolio. Charts 1 and 2 illustrate the development and composition of the equity share and fixed-income share, respectively, in a passive portfolio invested according to the market capitalisation in January 1970.

Chart 1: Equity share in a drifting portfolio

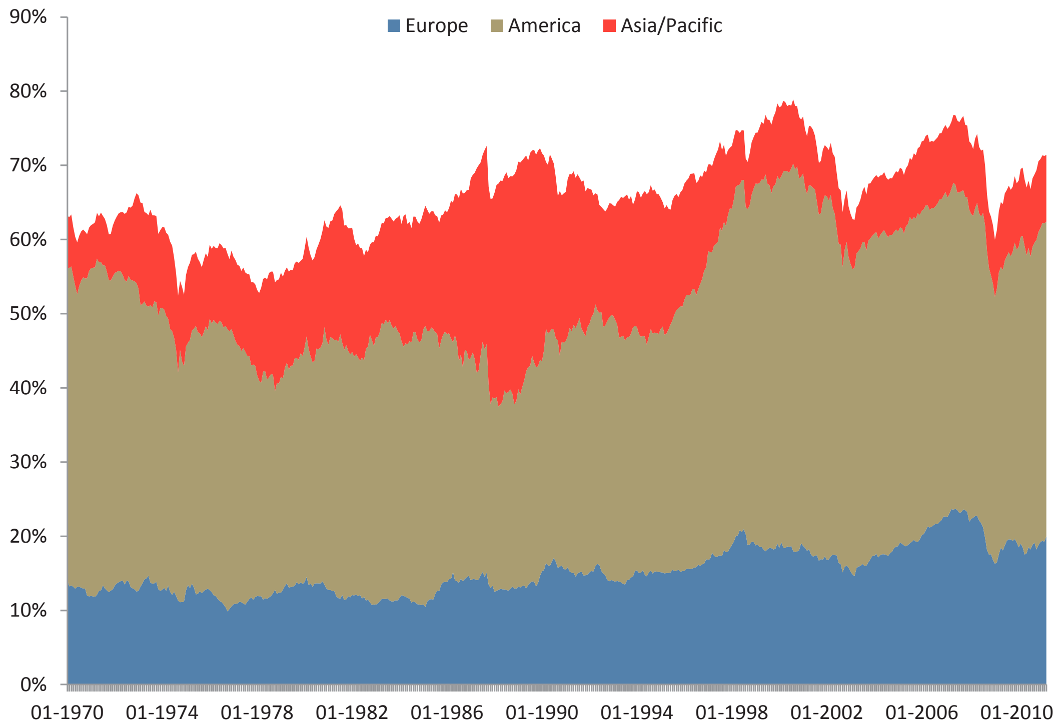

Chart 2: Fixed-income share in a drifting portfolio

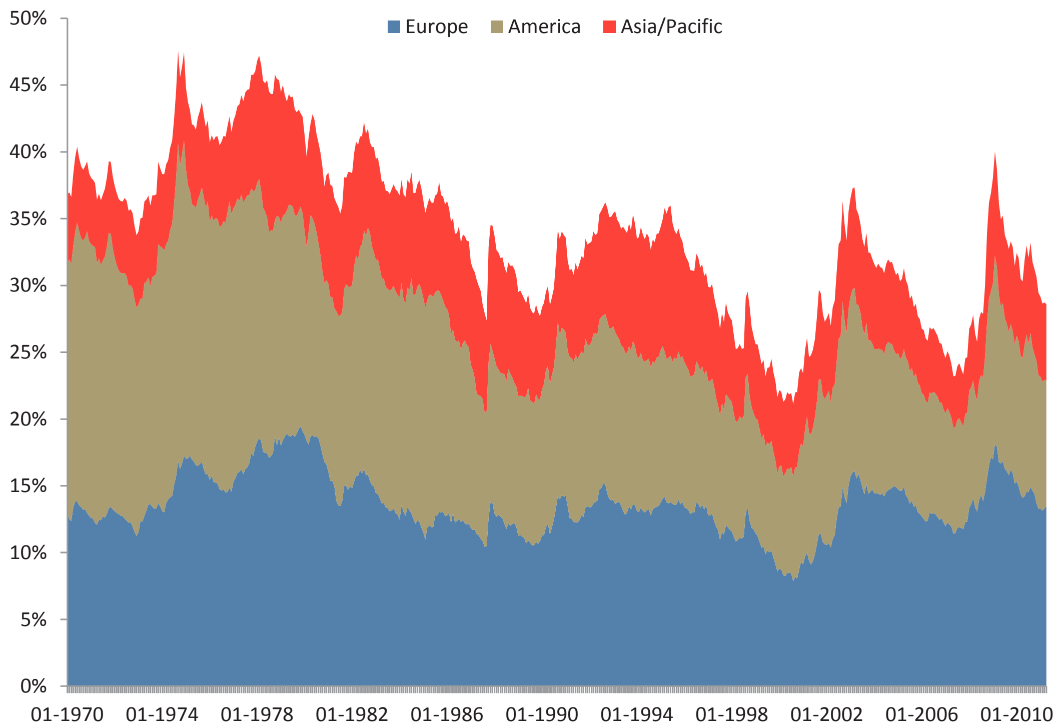

Table 1 summarises the risk and return characteristics of a passive, drifting equity and fixed-income portfolio invested according to market capitalisation weights at the outset of our analysis. We show that the risk and return characteristics differ across regions for both equities and fixed income. Table 1 also examines the risk and return characteristics of an equity portfolio ("Equity with regional factors") and a fixed-income portfolio ("Fixed income with regional factors") initially invested according to some fixed regional weights in contrast to market weights, but allowed to drift in line with market developments since inception. In order to facilitate our later analysis, we have used the same fixed regional weights as currently apply to the fund[2]. As a result of this, the portfolio will be impacted by over- or underweighting the different regions relative to their market capitalisation. We find that the fund's current regional tilts would have added value if held over our sample period.

Table 1 also shows the risk and return characteristics of different drifting portfolios. First, we calculate the performance of a passive, drifting portfolio invested according to market weights (no asset class factors, no regional factors) at the outset of our analysis. Second, we calculate the performance of a passive, drifting portfolio consisting of 60 percent equities and 40 percent bonds to begin with. As the market capitalisation of the two asset classes implied a distribution of 63.2 percent and 36.8 percent in equities and bonds, respectively, in January 1970, the two portfolios exhibit fairly similar risk and return characteristics. Finally, we have also calculated the performance of a portfolio which was initially invested in 60 percent equities and 40 percent bonds with a regional allocation within both asset classes similar to the one described above.

Table 1: Risk and return characteristics of drifting portfolios

	Return[1)]	Volatility	Sharpe[2)]	VaR 95%	CVaR 95%
Equity	8.90	15.30	0.58	-7.06	-9.68
– Europe	9.81	15.73	0.62	-6.89	-10.36
– America	8.91	17.58	0.51	-7.38	-10.44
– Pacific	9.54	19.98	0.48	-8.83	-11.57
Equity w/ regional factors	9.51	15.08	0.63	-6.82	-9.69
Fixed Income	7.90	8.23	0.96	-2.76	-3.87
– Europe	9.00	7.31	1.23	-2.43	-3.76
– America	7.07	11.38	0.62	-4.22	-5.89
– Pacific	9.17	14.05	0.65	-4.88	-6.46
Fixed Inc w/ regional factors	8.48	7.11	1.19	-2.55	-3.43
Portfolio (no asset class factors, no regional factors)	8.89	11.53	0.77	-5.15	-6.76
Portfolio (60/40, no regional factors)	8.86	11.22	0.79	-4.96	-6.52
Portfolio (60/40, with regional factors)	9.15	11.06	0.83	-4.40	-6.47

[1)] Geometric average
[2)] Defined as average return divided by volatility, ie we disregard the riskfree interest rate

A passive, drifting portfolio will, over time, become increasingly concentrated around the riskiest asset class. If we had started out with a 60 percent allocation to equities in January 1970, this share would have been 71.4 percent at the end of April 2011, after peaking in August 2000 at 78.9 percent. Hence, the absence of rebalancing may lead to a portfolio that is risky in an absolute sense, meaning that it may become concentrated in the riskiest asset class. In addition, a drifting portfolio will be suboptimal in a relative sense in that it may drift away from the strategic asset allocation that represents the fund owner's preferences.

Rebalancing asset classes and regions

In the following, we analyse the design of different rebalancing regimes. Our focus is on different designs of threshold-based and calendar-based regimes. These regimes are compared to the risk

2 Equities: Europe 50 percent, Americas 35 percent and Asia/Pacific 15 percent
Fixed income: Europe 60 percent, Americas 35 percent and Asia/Pacific 5 percent

and return characteristics of the passive portfolio with a drifting mix initially invested according to the fund's strategic weights (see above for details).

In the threshold-based rebalancing regime, the portfolio is rebalanced all the way back to the strategic asset allocation[3] if the deviation between the strategic weight and the actual weight is above a certain threshold for two consecutive months. The purpose of this "two-months-in-a-row" rule is to make sure that rebalancing is initiated only when the deviation from the strategic allocation is both significant and persistent. In the calendar-based regime, the rebalancing is carried out at a given frequency, e.g. monthly, quarterly or annually, regardless of how much the actual portfolio deviates from the strategic weights.

Once the design is chosen, one has to define the strategic weights to which these rules should apply. Below, we address an alternative similar to the fund's current strategic index, consisting of two asset classes (equity and fixed income) which are divided in turn into three regions (Europe, Americas and Asia/Pacific). Consequently, we have a total of six asset classes that can trigger a rebalancing of the portfolio. If a rebalancing is triggered, each of the six asset classes is brought back to the strategic allocation. In order to study the characteristics of the different rebalancing regimes in isolation, we disregard possible cash flows in and out of the portfolio.

Table 2 summarises the empirical characteristics of the different rebalancing alternatives. Compared with a drifting mix, the rebalanced portfolios have higher returns and lower volatility. As a result, the Sharpe ratio is also significantly higher. This result holds for all rebalanced portfolios, regardless of the exact specification and threshold level of the rebalancing regime. In addition, both the VaR and the tail risk, measured as Conditional VaR, are lower for the rebalanced portfolios. These results are consistent with the conclusions in Arnott and Lovell (1993) and Plaxco and Arnott (2002).

The lower risk is related to a lower equity share in the rebalanced portfolios, just north of 60 percent on average, compared to 65 percent on average in the drifting mix portfolio. We see from Charts 3 and 4 that the equity share in the drifting-mix portfolio is higher throughout our sample period, with the exception of a period from the mid-1970s to the mid-1980s, and more volatile in the drifting portfolio than the rebalanced portfolios. From a theoretical standpoint one could argue that an investor who rebalances his portfolio, and hence engages in a contrarian investment strategy, takes on additional risk compared to just investing passively in the drifting market portfolio. Based on our sample from 1970-2011, however, this risk has yet to materialise.

The rightmost column in Table 2 shows how much the average equity share deviated from the strategic allocation of 60 percent at the point when rebalancing was triggered. We see that the current rebalancing regime, with a 3 percent threshold level and six strategic weights, has implied an equity share that, on average, has been 4.2 percent off the strategic allocation when rebalancing is triggered.

3 Simulation results for "range-based rebalancing" where the portfolio is rebalanced back to the nearest edge of the no-trade region have also been carried out. These simulations are not significantly different from the threshold-based simulations and are not reported here. Results are available upon request.

Table 2: Absolute risk and return – rebalancing asset classes and regions

	Return[1]	Volatility	Sharpe[2]	VaR 95%	CVaR 95%	Ave Eq Share	End Eq Share	Ave Eq Share Deviation when Rebal
Drifting Mix								
– 60/40 with regional factors	9.15 %	11.1 %	0.83	-4.4 %	-6.5 %	65.0 %	68.9 %	NA
Calender based								
– Monthly	9.61 %	10.2 %	0.94	-4.0 %	-5.8 %	60.0 %	60.0 %	0.0 %
– Quarterly	9.72 %	10.2 %	0.95	-4.0 %	-5.8 %	60.1 %	60.3 %	1.2 %
– Annually	9.89 %	10.2 %	0.97	-3.9 %	-5.7 %	60.4 %	60.9 %	3.1 %
Threshold based								
– 1%	9.73 %	10.2 %	0.96	-4.0 %	-5.8 %	60.1 %	60.6 %	1.7 %
– 2%	9.84 %	10.2 %	0.97	-4.0 %	-5.8 %	60.4 %	60.1 %	2.8 %
– 3%	9.88 %	10.2 %	0.97	-4.0 %	-5.8 %	60.5 %	60.0 %	4.2 %
– 4%	9.81 %	10.2 %	0.96	-4.0 %	-5.8 %	60.6 %	60.1 %	5.0 %
– 5%	9.93 %	10.2 %	0.98	-3.9 %	-5.7 %	60.7 %	60.1 %	6.9 %
– 6%	9.97 %	10.2 %	0.98	-4.0 %	-5.7 %	60.7 %	68.7 %	7.2 %
– 7%	9.72 %	10.1 %	0.97	-4.1 %	-5.7 %	59.5 %	69.6 %	7.7 %
– 8%	10.01 %	10.2 %	0.98	-3.8 %	-5.8 %	61.5 %	68.1 %	11.5 %
– 9%	9.95 %	10.3 %	0.97	-3.8 %	-5.8 %	61.2 %	58.1 %	11.5 %
– 10%	9.60 %	10.2 %	0.94	-3.9 %	-5.7 %	59.1 %	57.8 %	5.2 %

[1] Geometric average
[2] Defined as average return divided by volatility, ie we disregard the riskfree interest rate

Chart 3: Equity share – rebalanced portfolio versus drifting portfolio

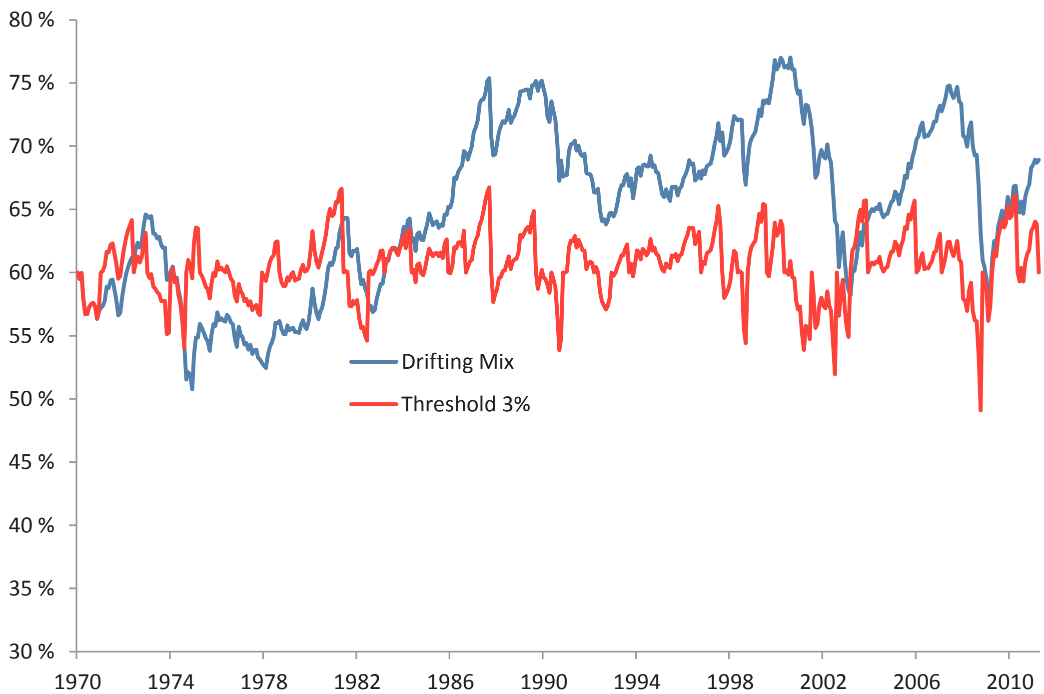

Chart 4: Rolling three-year volatility – rebalanced portfolio and drifting portfolio

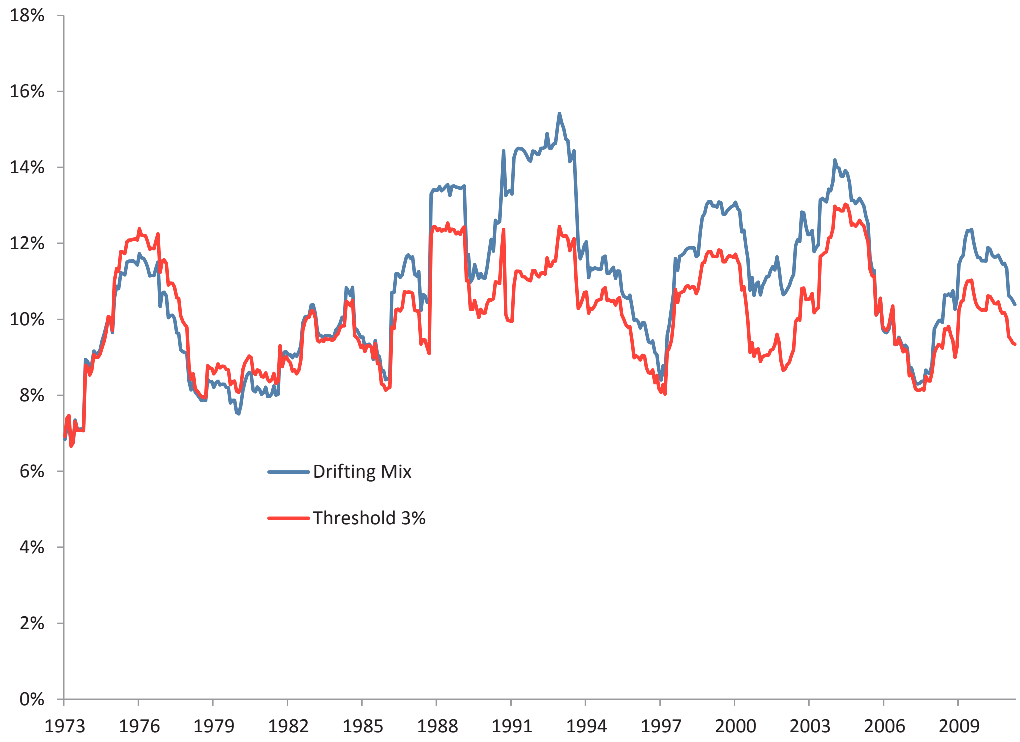

The results in Table 2 are based on a historical simulation over the entire sample period from January 1970 to April 2011. In order to check the robustness of the results, we have conducted the same simulation with different starting points. Chart 5 plots the Sharpe ratios for all rebalancing alternatives as well as the Sharpe ratio for the drifting mix with different starting points. Chart 5 shows that although the Sharpe ratios depend on the starting point for our simulation, the Sharpe ratios of the rebalanced portfolios are always higher than that of the drifting mix.

Chart 5: Sharpe ratios for rebalanced portfolios and drifting portfolio – different starting points

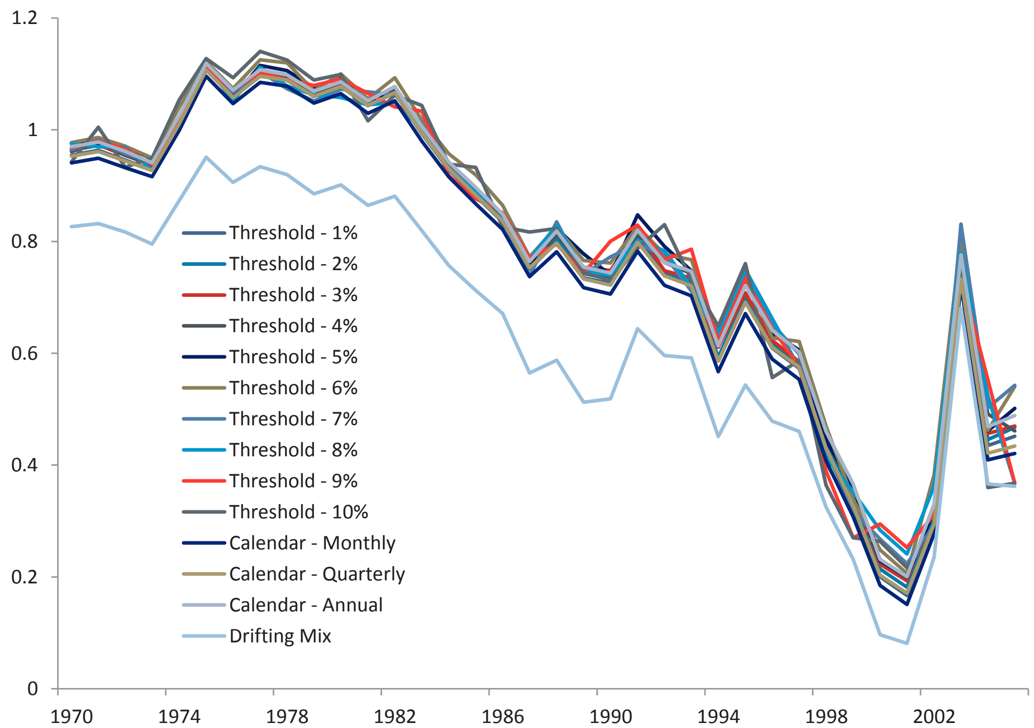

In order to investigate the differences between the different rebalancing rules further, we now look at relative risk/return and also take into account questions related to transaction costs. Table 3 summarises the performance and risk of the different specifications relative to the drifting mix. Transaction costs depend on the turnover in the portfolios, summarised in Table 4.

From Table 3, we can see that all rebalancing regimes are superior to the drifting mix, even after taking transaction costs into account. In order to investigate the robustness of the results, we have used fairly high estimates for transaction costs assuming a flat transaction cost of either 50 bp or 100 bp for both equity and fixed income securities. The assumed transaction costs related to rebalancing are more refined in the second part of this note. We see that there are some specifications with trigger levels higher than 3 percent that have a higher net IR than the fund's current regime. This is partly due to lower transaction costs, but also to the fact that these rules will perform better in trending markets.

Table 3: Relative risk and return – rebalancing asset classes and regions

	Excess Return[1]	Tracking Error[1]	IR	VaR 95%	CVaR 95%	IR after 50bps Tcost	IR after 100bps Tcost
Calender based							
– Monthly	0.34 %	2.7 %	0.13	-1.15 %	-1.80 %	0.08	0.03
– Quarterly	0.44 %	2.7 %	0.16	-1.13 %	-1.77 %	0.13	0.10
– Annually	0.59 %	2.7 %	0.22	-1.09 %	-1.78 %	0.20	0.18
Threshold based							
– 1%	0.44 %	2.7 %	0.16	-1.15 %	-1.78 %	0.14	0.11
– 2%	0.55 %	2.7 %	0.20	-1.14 %	-1.77 %	0.18	0.16
– 3%	0.59 %	2.7 %	0.22	-1.17 %	-1.78 %	0.20	0.18
– 4%	0.52 %	2.7 %	0.19	-1.16 %	-1.82 %	0.18	0.16
– 5%	0.63 %	2.7 %	0.23	-1.13 %	-1.79 %	0.22	0.21
– 6%	0.67 %	2.7 %	0.25	-1.09 %	-1.81 %	0.24	0.23
– 7%	0.42 %	2.7 %	0.16	-1.18 %	-1.85 %	0.15	0.14
– 8%	0.71 %	2.6 %	0.27	-1.13 %	-1.70 %	0.26	0.26
– 9%	0.66 %	2.7 %	0.24	-1.18 %	-1.77 %	0.23	0.23
– 10%	0.32 %	2.8 %	0.12	-1.24 %	-1.85 %	0.11	0.11

[1] Relative to Drifting mix

Table 4: Turnover – rebalancing asset classes and regions

	Ave Turnover Per Year	Ave Turnover Per Rebal	Ave # Rebals Per Decade
Drifting Mix			
– 60/40 with regional factors	0.0 %	0.0 %	–
Calender based			
– Monthly	26.5 %	2.2 %	120
– Quarterly	16.4 %	4.1 %	40
– Annually	9.6 %	9.7 %	10
Threshold based			
– 1%	13.6 %	5.9 %	23.2
– 2%	11.1 %	8.2 %	13.5
– 3%	9.5 %	10.6 %	9.0
– 4%	7.1 %	13.3 %	5.3
– 5%	7.3 %	17.6 %	4.1
– 6%	5.5 %	20.7 %	2.7
– 7%	3.6 %	21.0 %	1.7
– 8%	4.6 %	27.4 %	1.7
– 9%	4.1 %	28.2 %	1.5
– 10%	2.1 %	29.4 %	0.7

Chart 6 compares the performance of the current rebalancing regime to the drifting mix and shows that the fund's current strategy has historically added value fairly consistently relative to the drifting mix. The big dip in the late 1980s is related to the regional allocation to Asian equities. We will revert to this issue in section 1.2.2.

Chart 6: Cumulated excess returns – rebalanced portfolio versus drifting portfolio

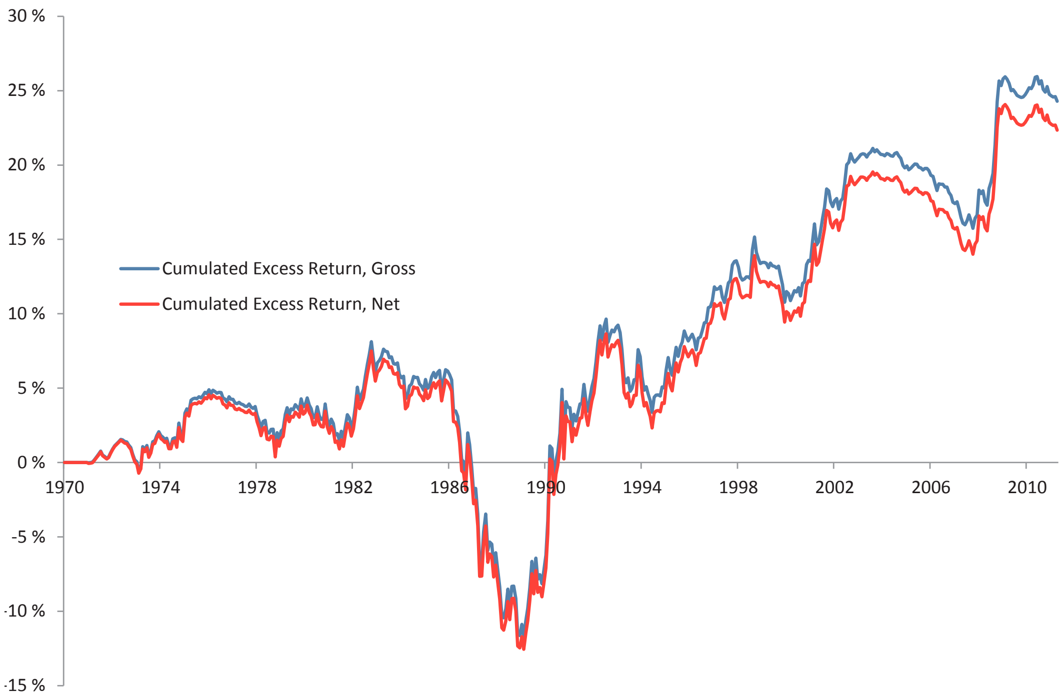

Our main conclusion is that the fund's current rebalancing regime has performed well compared to most alternatives over our sample period. Our empirical examination indicates that a higher threshold of e.g. 5 or 6 percent would have done even better in terms of risk-adjusted performance. We see from Tables 2 and 3 that these alternatives have somewhat higher returns, but with approximately the same absolute volatility and tail risk. The reason for this is not only that these alternatives have lower turnover and transaction costs, but also that these alternatives to a greater extent are exposed to momentum effects.

What to rebalance?

New portfolio theory expands on traditional theory by pointing out that an investor's portfolio should be constructed by allocating to a wide range of different systematic risk premia rather than to traditional asset classes. The most well-known and best documented systematic risk premia include the "equity risk premium", "value", "small cap", "momentum", "carry", "credit" and "volatility". Implementing a framework for dynamically rebalancing between a potentially large set of risk factors would be challenging in practice, especially if one wants to keep the rebalancing process as simple and transparent as possible. Hence, a pragmatic approach is called for.

One would argue that the equity risk premium is the most important risk factor. Furthermore, this equity risk premium is fairly straightforward to define, and the portfolio's strategic exposure to it can easily be controlled by managing the equity share in the portfolio. In this section, we look at the risk and return characteristics of a rebalancing regime constructed to maintain a given strategic weight to equities.

Rebalancing the equity share

The objective of the regime examined in this section is to bring the equity share back to the strategic weight whenever rebalancing is triggered. Rebalancing will be carried out by scaling the equity share up or down without altering the regional distribution. Consequently, the regional distribution for both the equity and the fixed-income portfolios will drift according to market movements.

We start by comparing the absolute risks and returns of the rebalanced portfolios to the drifting mix. These results are summarised in Tables 5 and 6. The passive portfolio with a drifting mix is the same as used in section 2, invested in 60 percent equities and 40 percent fixed income with an initial regional distribution equal to the fund's current regional weights. This drifting portfolio is compared to portfolios where the equity share is rebalanced, either calendar-based or threshold-based. At the outset of our analysis, the rebalanced portfolios are invested according to the fund's current six strategic weights.

In Charts 7 and 8, we compare Sharpe ratios and turnover for this alternative regime where we rebalance the equity share to a regime where we rebalance six assets (two asset classes and three regions).

Table 5: Absolute risk and return – rebalancing the equity share

	Return	Volatility	Sharpe	VaR 95%	CVaR 95%	Ave Eq Share	End Eq Share	Ave Eq Share Deviation when Rebal
Drifting Mix	9.15 %	11.1 %	0.83	-4.4 %	-6.5 %	65.0 %	68.9 %	NA
Calender based								
– Monthly	9.37 %	10.3 %	0.91	-4.1 %	-5.8 %	60.0 %	60.0 %	0.0 %
– Quarterly	9.49 %	10.3 %	0.92	-4.1 %	-5.7 %	60.1 %	60.3 %	1.2 %
– Annually	9.59 %	10.3 %	0.93	-4.1 %	-5.7 %	60.3 %	60.7 %	3.2 %
Threshold based								
– 1%	9.44 %	10.3 %	0.91	-4.1 %	-5.8 %	60.1 %	60.7 %	2.4 %
– 2%	9.48 %	10.3 %	0.92	-4.1 %	-5.8 %	60.2 %	60.0 %	3.2 %
– 3%	9.55 %	10.3 %	0.92	-4.1 %	-5.7 %	60.4 %	60.0 %	4.6 %
– 4%	9.53 %	10.3 %	0.92	-4.1 %	-5.8 %	60.5 %	64.0 %	5.4 %
– 5%	9.59 %	10.4 %	0.93	-4.1 %	-5.7 %	61.2 %	63.8 %	6.4 %
– 6%	9.44 %	10.4 %	0.91	-4.1 %	-5.8 %	61.0 %	63.8 %	7.3 %
– 7%	9.63 %	10.5 %	0.92	-4.2 %	-5.8 %	61.7 %	60.0 %	8.6 %
– 8%	9.56 %	10.4 %	0.92	-4.2 %	-5.8 %	61.3 %	60.3 %	9.0 %
– 9%	9.50 %	10.4 %	0.91	-4.1 %	-5.8 %	60.6 %	60.0 %	10.2 %
– 10%	9.45 %	10.3 %	0.92	-4.1 %	-5.7 %	59.4 %	68.6 %	11.6 %

Table 6: Turnover – rebalancing the equity share

	Ave Turnover Per Year	Ave Turnover Per Rebal	Ave # Rebals Per Decade
Drifting Mix			
- 60/40 with regional factors	0.0 %	0.0 %	-
Calender based			
– Monthly	17.5 %	1.5 %	120
– Quarterly	12.1 %	3.0 %	40
– Annually	7.1 %	7.2 %	10
Threshold based			
– 1%	9.2 %	4.9 %	18.9
– 2%	7.9 %	6.7 %	11.9
– 3%	7.5 %	9.2 %	8.2
– 4%	6.5 %	10.4 %	6.3
– 5%	6.1 %	12.6 %	4.8
– 6%	3.9 %	14.5 %	2.7
– 7%	5.2 %	16.5 %	3.1
– 8%	4.1 %	18.9 %	2.2
– 9%	3.4 %	19.9 %	1.7
– 10%	2.2 %	23.0 %	1.0

Chart 7: Comparing Sharpe ratios

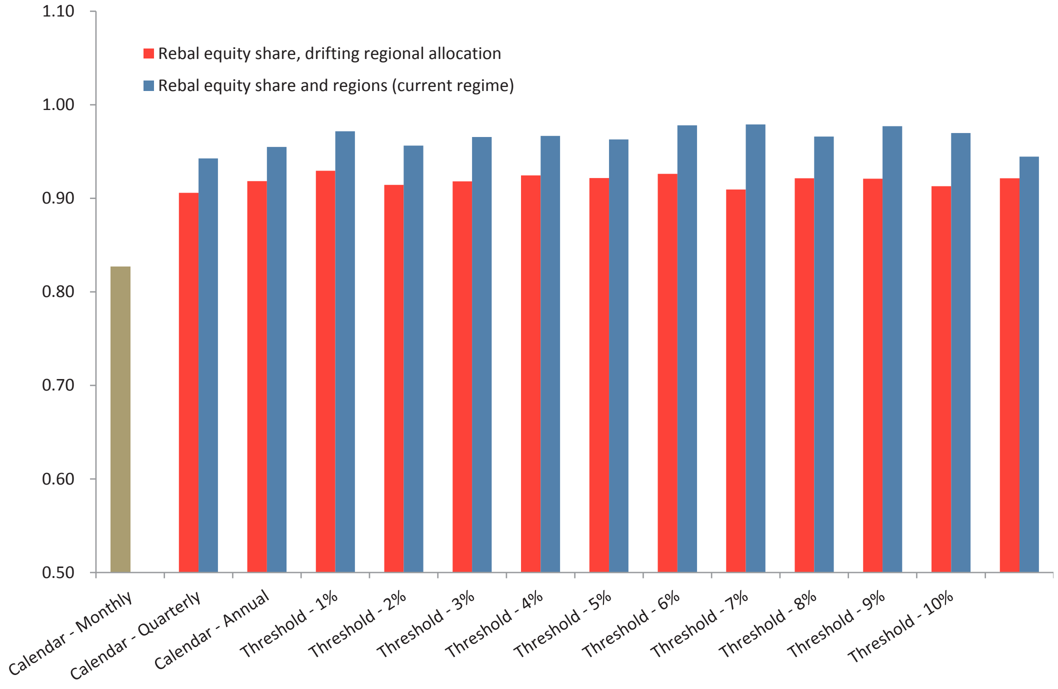

Table 5 tells us that a rebalancing regime focusing on the equity share would have given a higher Sharpe ratio than the drifting mix over our sample period. However, as Chart 7 illustrates, although rebalancing the equity share is the key driver, rebalancing the regional allocation within each asset class yields further improvement in the Sharpe ratio. In other words, rebalancing along both the asset class dimension and the regional dimension contributes to a higher Sharpe ratio than the drifting mix, but the asset class dimension is the most important driver.

In general, a rebalancing regime focusing on many dimensions will most likely incur higher turnover and transaction costs relative to a regime focusing on fewer dimensions. This is confirmed in Chart 8, where the turnover in a regime focusing on the equity share is consistently lower than in a regime rebalancing both the asset class mix and the regional mix within each asset class.

Chart 8: Comparing turnover

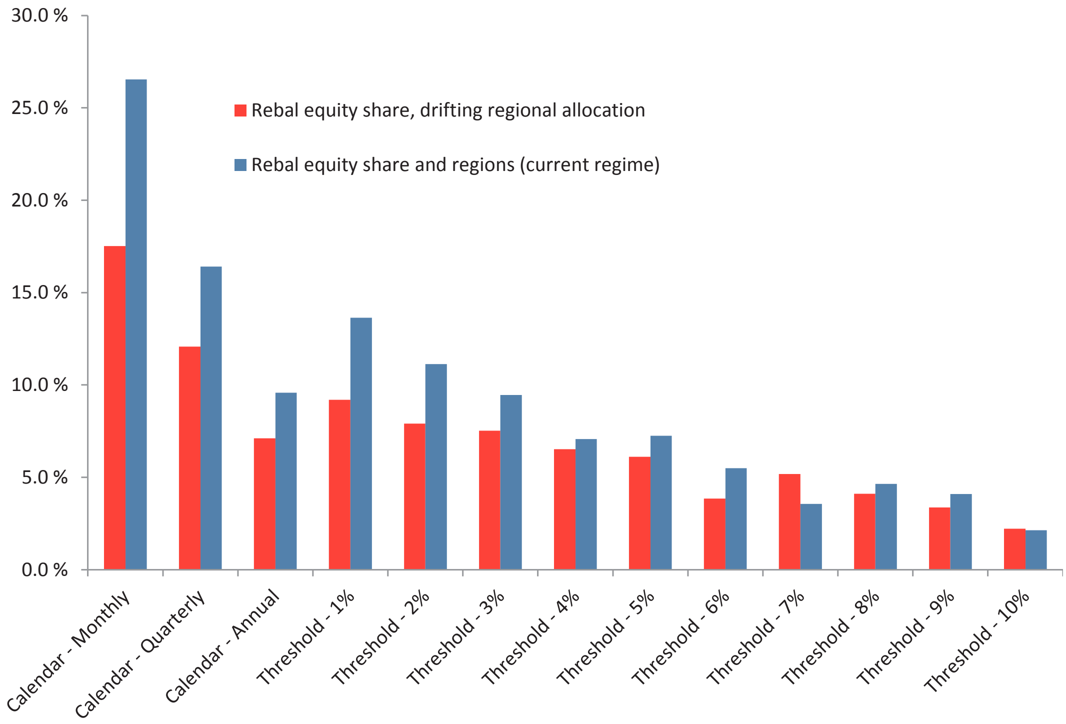

In order to investigate this further, we have to incorporate transaction costs. Table 7 summarises the risk and return of our alternative rebalancing regime (around the strategic equity weight) relative to the drifting mix. The rebalanced portfolios outperform the drifting mix even after taking transaction costs into account. However, the differences in turnover between the two rebalancing regimes illustrated in Chart 8 are not sufficient to alter our conclusion about the risk-adjusted return of the two alternative rebalancing regimes, as can be seen from Chart 9. In our sample, the regional rebalancing has added value despite higher turnover.

Table 7: Relative risk and return – rebalancing the equity share

	Excess Return[1]	Tracking Error[1]	IR	VaR 95%	CVaR 95%	IR after 50bps Tcost	IR after 100bps Tcost
Calender based							
– Monthly	0.13 %	1.30 %	0.10	-0.51 %	-0.71 %	0.03	-0.03
– Quarterly	0.24 %	1.32 %	0.18	-0.46 %	-0.68 %	0.14	0.09
– Annually	0.33 %	1.37 %	0.24	-0.47 %	-0.74 %	0.21	0.19
Threshold based							
– 1%	0.19 %	1.32 %	0.15	-0.46 %	-0.72 %	0.11	0.08
– 2%	0.23 %	1.32 %	0.17	-0.48 %	-0.74 %	0.14	0.11
– 3%	0.29 %	1.32 %	0.22	-0.43 %	-0.74 %	0.19	0.16
– 4%	0.27 %	1.32 %	0.21	-0.43 %	-0.72 %	0.18	0.16
– 5%	0.33 %	1.36 %	0.24	-0.46 %	-0.72 %	0.22	0.20
– 6%	0.20 %	1.29 %	0.16	-0.45 %	-0.71 %	0.14	0.12
– 7%	0.38 %	1.27 %	0.30	-0.42 %	-0.68 %	0.28	0.26
– 8%	0.31 %	1.24 %	0.25	-0.42 %	-0.72 %	0.24	0.22
– 9%	0.25 %	1.13 %	0.22	-0.42 %	-0.63 %	0.21	0.19
– 10%	0.19 %	1.37 %	0.14	-0.58 %	-0.79 %	0.13	0.12

[1] Relative to Drifting mix

Chart 9: Sharpe ratios after transaction costs

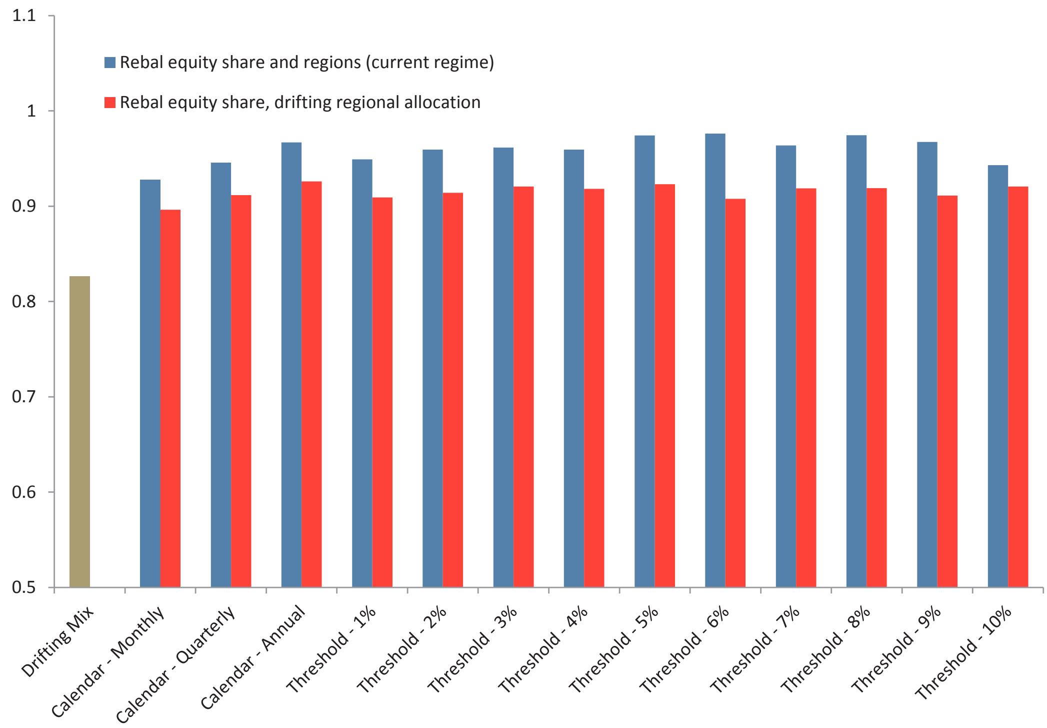

The impact of regional rebalancing

Regional rebalancing has contributed to a higher Sharpe ratio net of transaction costs in our sample period. In this section, we investigate this finding.

In order to isolate the effect of regional rebalancing, we compare the return on two portfolios. The first portfolio is rebalanced around the equity share only, disregarding the regional allocation. We choose a threshold level of 3 percent for the equity share. The second portfolio is also rebalanced around the equity share with a threshold level of 3 percent. However, in this portfolio, we also rebalance the regional weights back to the strategic weights whenever rebalancing is triggered. In other words, regional deviations cannot themselves trigger rebalancing in this portfolio, rather the regional allocation is rebalanced whenever the equity share triggers a rebalancing. Hence, we call this alternative "conditional regional rebalancing".

Regional rebalancing will add value if there is a tendency towards mean reversion in returns at regional level, and detract value if returns are trending. If we look at the regional composition of the drifting portfolio in Charts 1 and 2, we can get an idea of what is driving the gains from regional rebalancing. One striking feature is the increase in the share of Asian equities during the 1980s and the subsequent decrease in late 1980s and the first part of the 1990s.

In Chart 10, we have plotted two graphs. The red graph shows the cumulated return on the portfolio with conditional regional rebalancing in excess of the return on the rebalanced portfolio with drifting regional allocation. Hence, this graph shows the isolated cumulated gain from regional rebalancing over our sample period. The blue graph shows the weight of Asian equities in the portfolio with drifting regional allocation.

The share of Asian equities in both portfolios starts out at the strategic weight in January 1970. However, during the first part of the 1980s, this share increases to almost 40 percent in the portfolio with no regional rebalancing, as the Asian, and particularly the Japanese, equity market outperforms other regions. During the 1990s, this region underperforms, leading the Asian equity share back to around 12 percent. As is evident from the graph, the historical gains from regional rebalancing are strongly related to the rise and fall of the Asian equity market share. A regime with no regional rebalancing will let the share of Asian equities in the portfolio grow as the Japanese equity market rises during the 1980s. As long as the Japanese equity market continues to trend upward, a regime which brings the weight of Asia down to the strategic weight whenever rebalancing takes place will underperform a regime which lets the allocation to the outperforming region increase. However, at the time when the Asian market starts to underperform, Asian equities account for almost 40 percent of the portfolio with no regional rebalancing, whereas the portfolio with regional rebalancing will have a significantly lower share of its equities allocated to Asia. As a result of this, the portfolio with regional rebalancing will outperform as the Asian market starts to underperform. In this case, the net effect over the entire sample period favours regional rebalancing. Our analysis suggest that this result to a great extent can be explained by the build-up and subsequent burst of the bubble in the Japanese equity market.

Chart 10: Exploring the gains from regional rebalancing

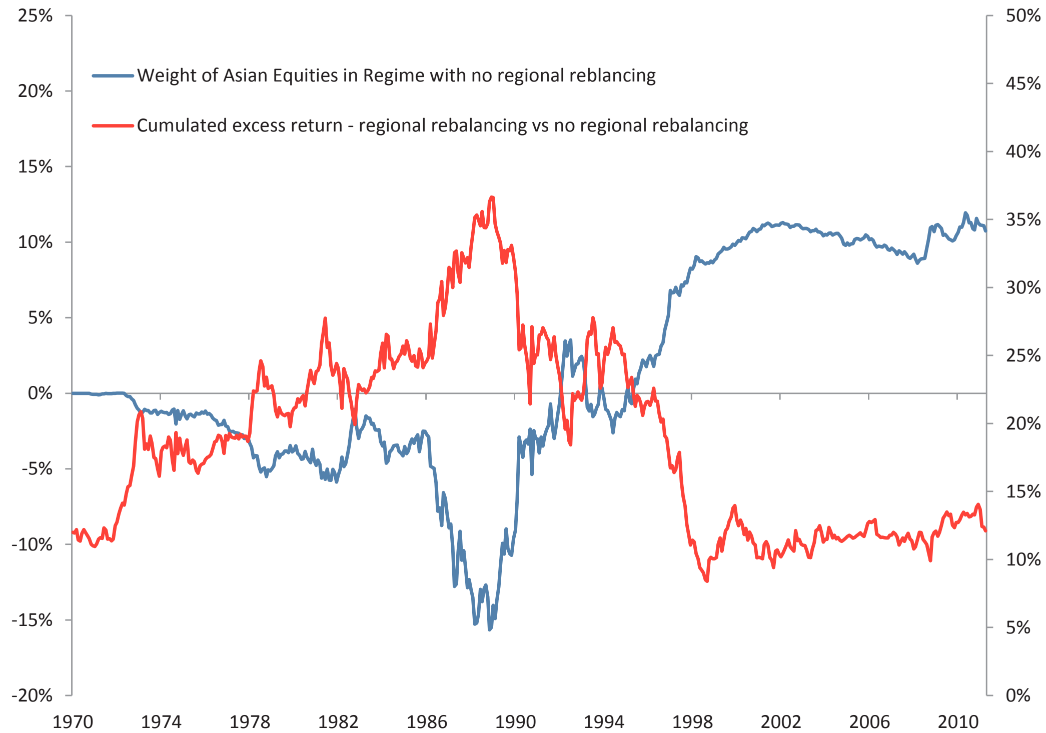

Designing a threshold based rebalancing regime

As discussed in the first part of this note, the rationale for having a rebalancing rule in place is theoretical more than empirical, although empirical analyses of asset returns from 1970 until today show that rebalanced portfolios outperform a drifting-mix portfolio. The rationale for a threshold-based rebalancing regime rather than a calendar-based regime is also theoretically founded. In this note, we examine specific design questions of a trigger-based regime, the deciding level of deviation triggering a full rebalancing, the required persistence of this deviation and, finally, implementation once a full rebalancing is triggered. These design questions balance the concerns of, on the one hand, keeping deviations from the strategic weights as small as possible and, on the other hand, keeping transaction costs in connection with rebalancing low. In these analyses we therefore refine our assumptions on transaction costs compared to the first part of the note. We look at trigger levels of 3, 4 and 5 percent and persistence requirements of one, two and three consecutive months. Finally, we compare implementation of the full rebalancing all-in-one-go versus gradual implementation over three months.

Methodologies

We conduct two types of simulations in order to estimate the costs related to different rebalancing rules.

The ex ante analysis is based on a Monte Carlo simulation, where we make assumptions regarding the expected return and volatility of the asset classes as well as the correlation between the different asset classes. The asset returns are assumed to be normally distributed. Expected returns for the different asset classes are built the following way: For fixed income, we use the current yield on ten-year government bonds. For the other segments in the fixed-income benchmark, we add a risk premium of 0.8 percentage point. Equities in a given region are assumed to give a return that is 2.5 percentage points higher than the expected return on fixed income in the same region. Estimates for volatility and correlation are based on historical observations from the last 20 years. We use the projections for inflows into the fund from the National Budget for 2011.

The ex post analysis is based on historical returns on the different asset classes going back to January 1970. Hence, in this approach we are not making any up-front distributional assumptions for asset class returns. For example, the assumptions made in the Monte Carlo simulation rule out kurtosis (tail events) and any momentum or mean-reversion effects, which may have a significant impact on the performance, risk and turnover of different rebalancing rules. As a result of this, the number and size of rebalancings are expected to be higher in the ex post analysis than in the ex ante analysis. However, we implicitly assume that the historical returns are representative of the future. Moreover, for obvious reasons the historical simulation does not take into account inflows into the fund, and as a result does not capture the effect that cash flows may have on the rebalancing frequency and rebalancing costs.

Our analysis below confirms that the ex ante simulations give a lower number of rebalancings and lower rebalancing turnover. As expected, including projected cash flows in a rebalancing regime with partial rebalancing gives an even lower number of full rebalancings and lower rebalancing turnover. Although the different simulation approaches give somewhat different estimates for rebalancing turnover, the qualitative conclusions drawn in this note are robust to the choice of simulation methodology.

There is significant uncertainty in estimating transaction costs. Transaction costs vary with trade size and other characteristics as well as market conditions. Actual transaction costs will be driven by the trade specifics and market conditions in the relevant rebalancing periods. Large illiquid trades will in general be more expensive to trade than smaller and more liquid orders. In stressed markets where liquidity is low and volatility is high, transaction costs are in general expected to be higher than in more normal market conditions. This uncertainty may be particularly relevant for high-threshold regimes where it may be more likely that large trades are implemented in stressed markets.

In our analyses, we have estimated transaction costs based on the current size of the fund and current market conditions.

To analyse transaction costs for equities, we have looked at the estimated costs of trading a slice of the current benchmark index. An internal NBIM market impact cost model is used to estimate the market impact of the individual securities as a function of spread, volatility, liquidity and the rate of participation. To the estimated market impact from this model, we have added estimated commissions and taxes. The average cost of trading about 1 percent of the fund, or NOK 30 bn based on the current size of the fund, across the equity benchmark index is estimated to be about 30-35 bp. If the trade size increases to about 3 percent of the fund, the estimated average transaction cost increases to about 45-50 bp.

For the fixed-income trade, we have assumed that the average cost of trading a slice of the fixed-income benchmark is about 10 bp. This estimate is based on internal trades over the last two years. However, the estimated trading costs of the different segments of the index vary greatly, and the less liquid and higher-cost segments are most difficult to estimate. Treasuries typically have low trading costs of 0-5 bp, while corporate bonds are expensive to trade. We have assumed an average cost of trading corporate bonds of about 35 bp. Transaction costs for fixed-income securities are also expected to vary with trade size and market conditions. However, in the absence of a good model for estimating this for fixed-income securities, we assume a flat cost across all scenarios.

Rebalancing thresholds

In this section, we compare the impact on rebalancing costs of moving away from the current rebalancing regime with six assets to a rebalancing regime with only two assets, equities and bonds. We also compare different rebalancing thresholds in the regime with only two assets. We look at rebalancing thresholds of three, four and five percent assuming a "two-months-in-a-row" persistence requirement and "all-in-one-go" implementation for all scenarios. The historical simulations show that the average deviation of the equity share in the current regime is about 4 percent when a full rebalancing is triggered.

The simulation results show that rebalancing of six assets implies a higher rebalancing frequency compared to a regime where we only rebalance the equity share, given the same rebalancing threshold of 3 percent. In addition, rebalancing six assets involves more trading once rebalancing is triggered. The reason for this is that, in this case, we also have to bring the regional allocation back to the strategic weights, in addition to bringing the asset class weights back in line with the strategic weights. As a result of both higher rebalancing frequency and higher rebalancing size, the regime with six asset classes involves a higher rebalancing cost relative to a regime with only two assets. Chart 14 summarises the results.

The asset class mix decision is the most important driver of long-term fund returns. In addition to costs related to rebalancing, the owner should therefore also care about deviations from the strategic asset allocation at asset class level. Chart 15 shows that a rebalancing regime with two assets will keep the asset class weights closer to the strategic asset class weights than the current regime with six asset classes. Chart 15 is based on historical data. The tracking error of the equity share relative to the strategic weight of 60 percent is 2.1 percent for the regime with two assets and 2.5 percent for the regime with six assets.

We have also looked at the regime with two assets with higher rebalancing thresholds. As expected, higher thresholds mean a lower number of rebalancings per decade, but the rebalancing size will, of course, also increase. The net effect will, however, mean somewhat lower transaction costs compared to a threshold of 3 percent, as shown in Chart 14. On the other hand, a higher threshold will also give a higher tracking error for the equity share.

Choosing a rebalancing regime involves weighing the desire to keep the tracking error of the equity share low against keeping the rebalancing costs low. In Chart 16, we have plotted the rebalancing costs and the tracking error of the equity share for the different rebalancing rules. We see that the regime with two assets and a threshold of 3 percent gives both lower rebalancing costs and lower tracking error for the equity share than the current rebalancing regime.

A regime with two assets and a rebalancing threshold of 4 percent will give the same tracking error for the equity share as the current regime with six assets and a threshold of 3 percent, but significantly lower rebalancing costs. A higher threshold means that it takes bigger market moves to trigger rebalancing, and also that the average size of the rebalancing will be higher. This introduces more uncertainty regarding the transaction costs, because we will trade more in stressed markets with high volatility and low liquidity. Although the expected rebalancing costs related to a regime with a higher threshold are lower, one should also take into consideration the increased uncertainty about the expected rebalancing costs. This consideration favours a regime with a narrower band and smaller rebalancings in less stressed markets. By setting a lower rebalancing threshold and allowing more frequent but smaller full rebalancings, this will become to a greater extent part of normal fund management, and one can expect more robust procedures for handling full rebalancings.

Chart 11: Number of rebalancings per decade

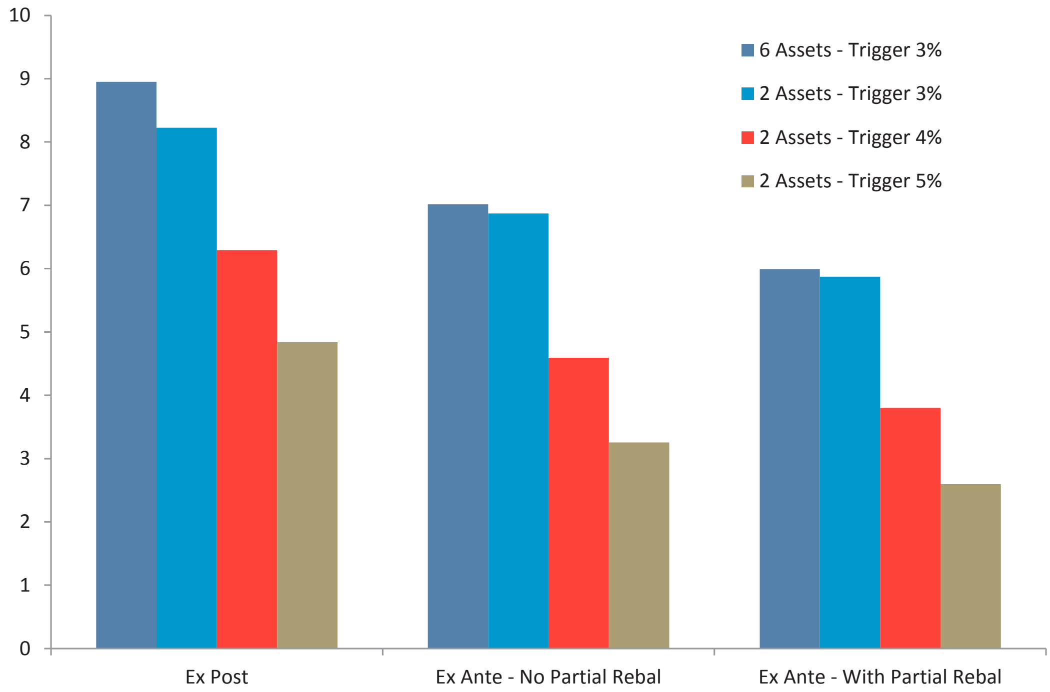

Chart 12: Average size of rebalancings as a percentage of fund size

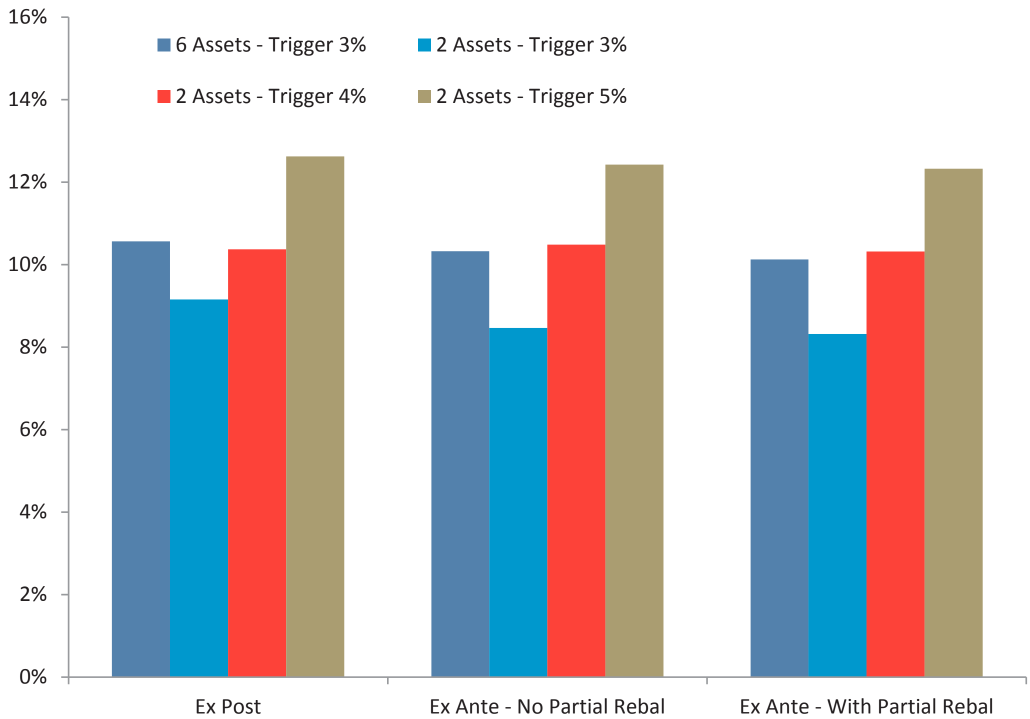

Chart 13: Rebalancing costs per decade as a percentage of fund size

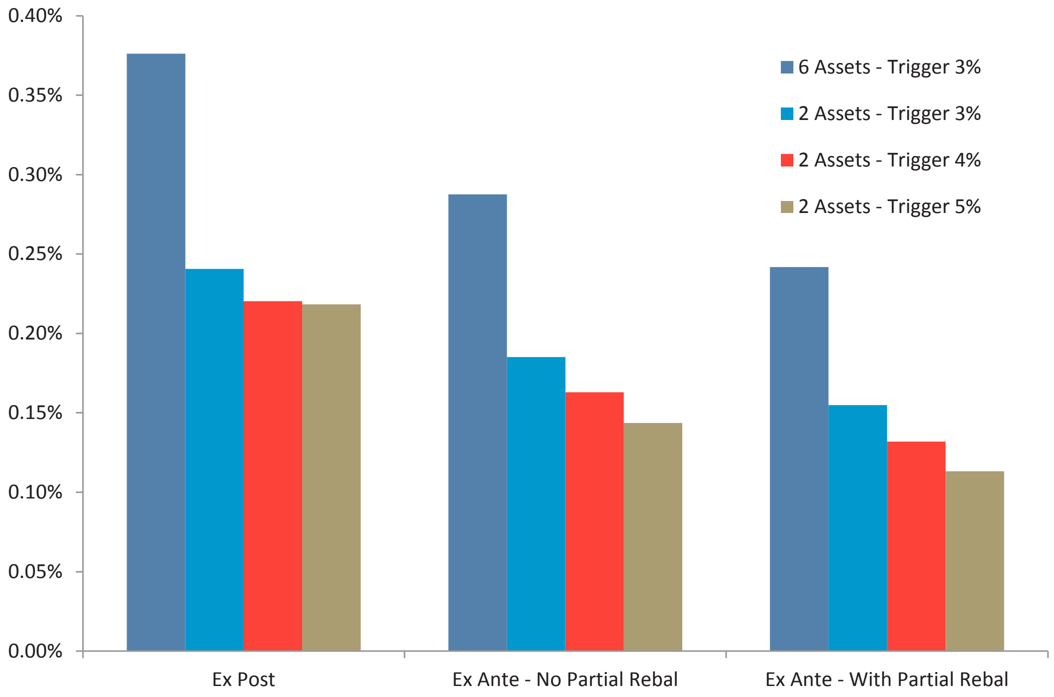

Chart 14: Equity share for two alternative rebalancing regimes

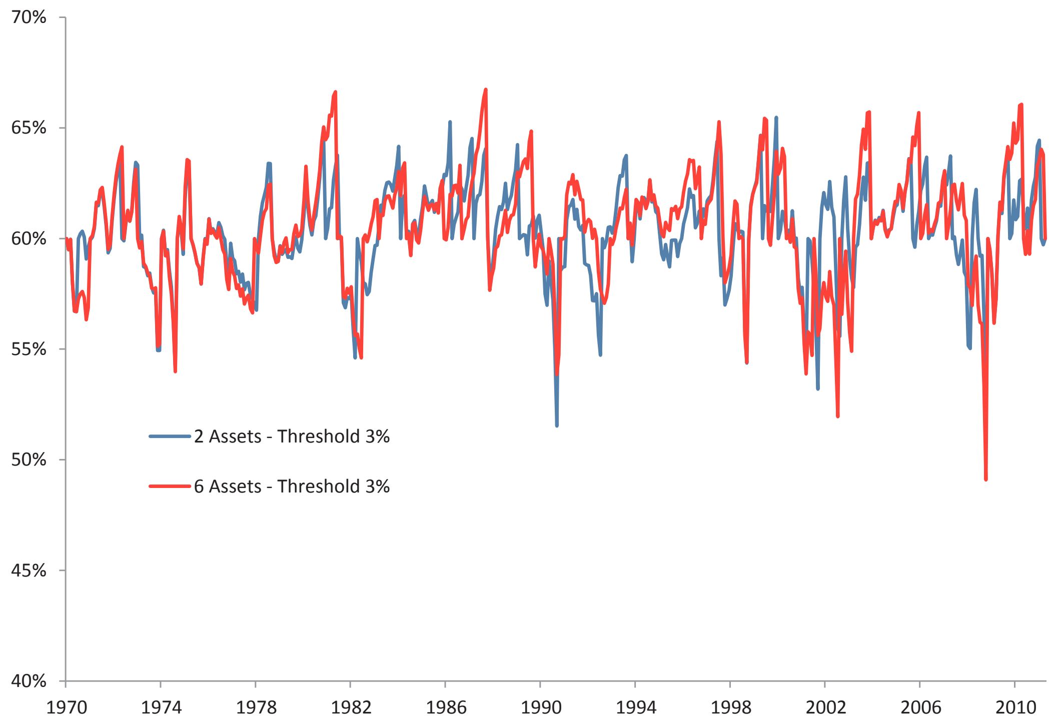

Chart 15: Rebalancing costs and tracking error of equity share

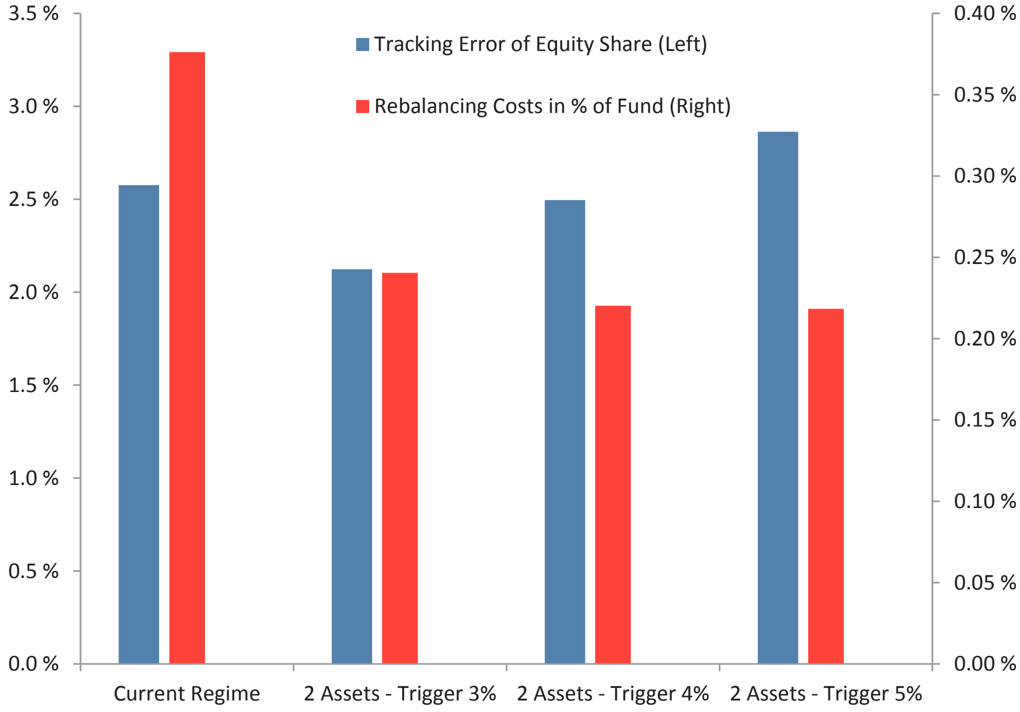

Persistence

In the following we will assume that rebalancing is triggered when the equity share deviate 3 percentage points from the strategic weights. In the current regime, the deviations from the strategic weights have to be outside the no-trade region for two consecutive months for a rebalancing to take place. In this section, we investigate the implications of altering this persistence requirement (assuming "all-in-one-go" implementation for all scenarios).

We expect a negative relationship between the persistence requirement and the average number of rebalancings. The results shown in Chart 17 confirm this: Requiring that the deviations are outside the no-trade region for three consecutive months implies a lower number of rebalancings, whereas a persistence requirement of one month involves a high number of rebalancings. On the other hand, there is a positive relationship between the persistence requirement and average rebalancing size, as shown in Chart 18. This is because the deviation from the strategic asset allocation is larger on average when the deviation is required to be outside the no-trade region for three consecutive months. If just one month outside the no-trade region is required to trigger a full rebalancing, we will get more frequent but smaller rebalancings compared to the alternatives. All in all, however, the rebalancing costs are higher for this alternative, as shown in Chart 19.

The persistence requirement obviously determines how long we let the equity share be outside the no-trade region before we rebalance. As a result, there will be important differences in terms of how closely the portfolio tracks the strategic asset mix. The equity share for the three alternatives is plotted in Chart 20. The tracking error of the equity share relative to the strategic asset class mix of 60 percent is 2.6 percent for the "three-months-in-a-row" alternative, 2.1 percent for the "two-months-in-a-row" alternative and 1.7 percent for the "one-month-in-a-row" alternative.

The question is therefore how much of an increase in transaction costs we are willing to incur in order to track the strategic asset class mix more closely. We have plotted the rebalancing costs and the tracking error of the equity share relative to the strategic asset class mix in Chart 21. Given the significant improvement in the tracking error of the equity share in the "one-month-in-a-row" alternative, the extra transaction costs could be worthwhile.

Chart 16: Number of rebalancings per decade

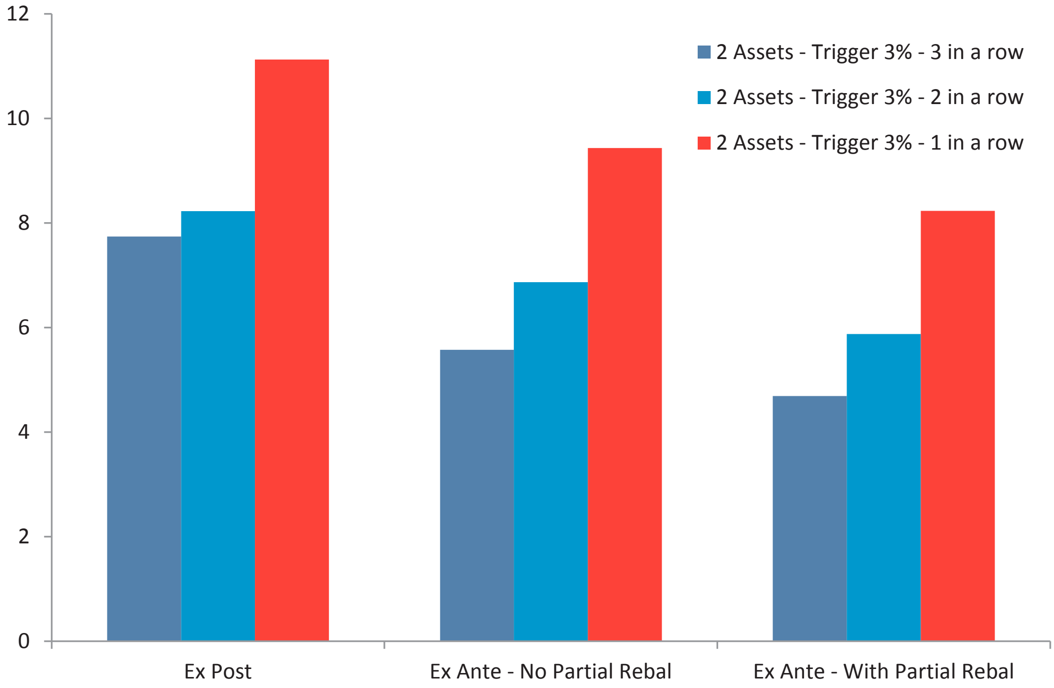

Chart 17: Average size of rebalancings as a percentage of fund size

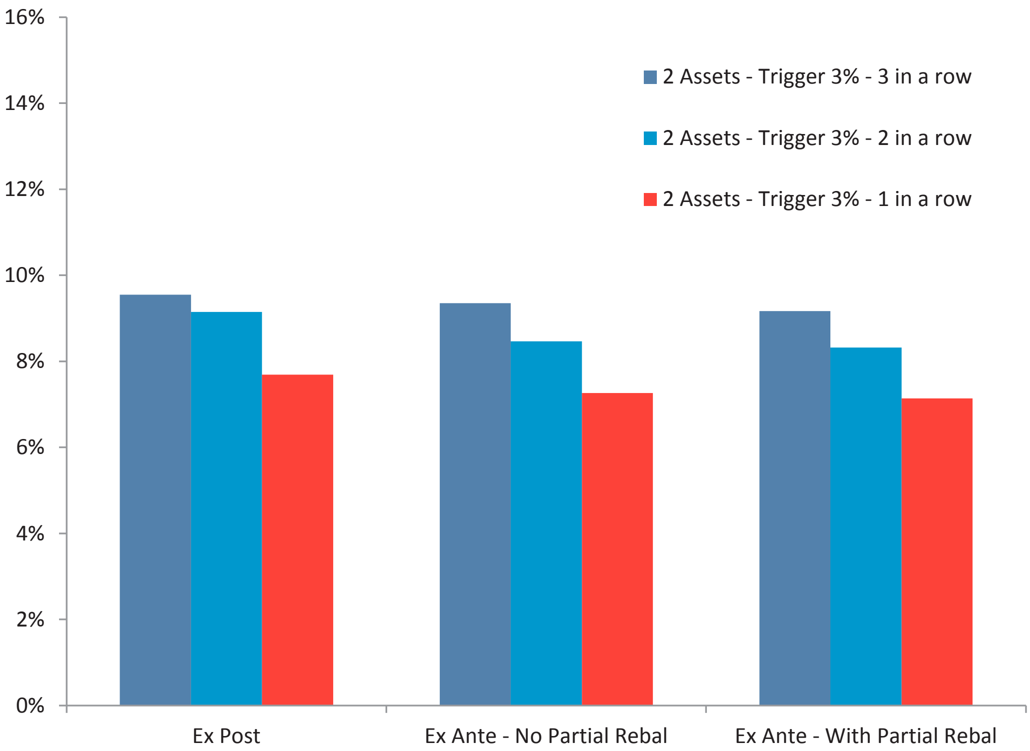

Chart 18: Rebalancing costs per decade as a percentage of fund size

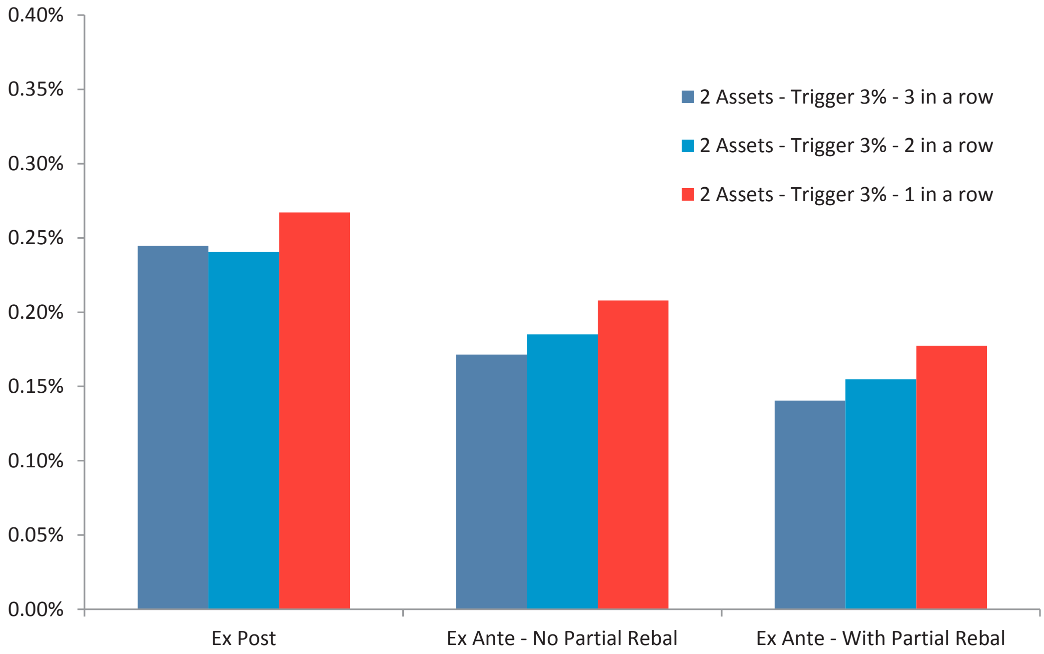

Chart 19: Equity share under three alternative rebalancing regimes

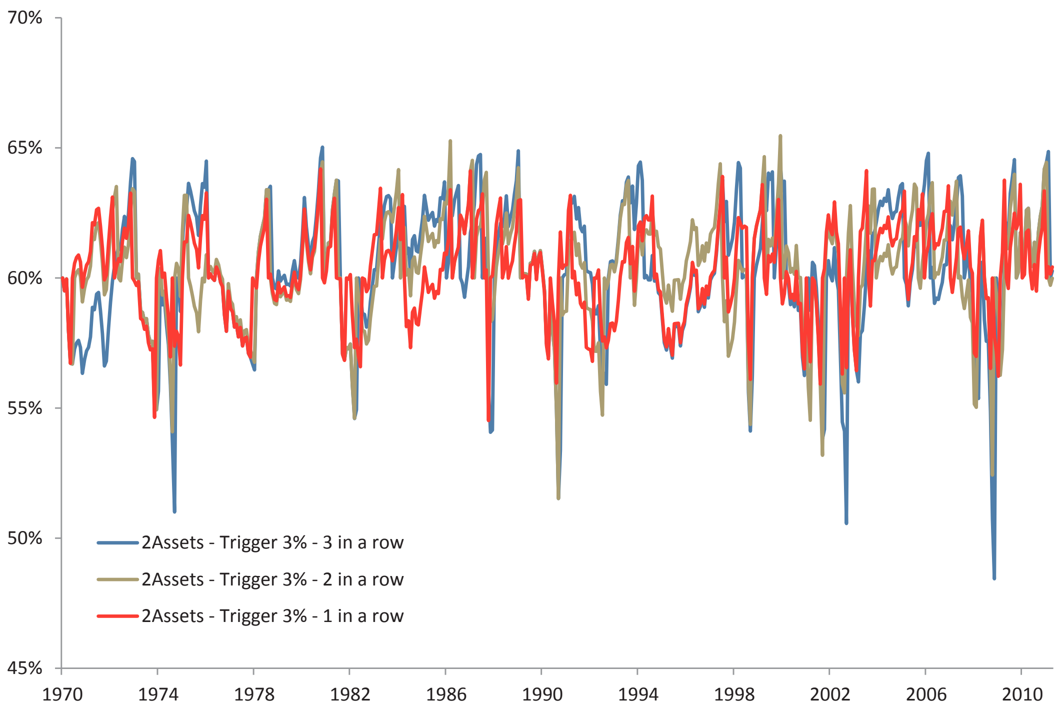

Chart 20: Rebalancing costs and tracking error of equity share

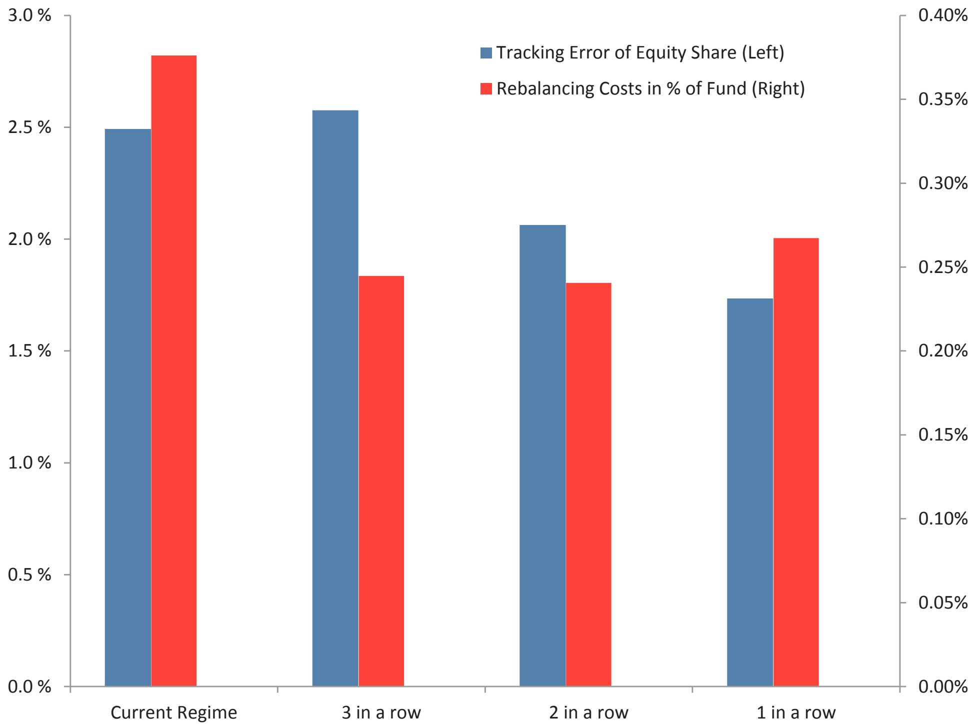

Speed of implementation

Our analysis has so far been based upon an assumption that rebalancing takes place at month-end the month after a full rebalancing has been triggered. We now examine how a more gradual implementation will impact our results, assuming a "one-month-in-a-row" persistence requirement for all scenarios. In the scenario with gradual rebalancing, we assume that one-third of the deviation is eliminated at the first month-end after full rebalancing has been triggered, half of the remaining deviation is eliminated at the second month-end, and the entire (remaining) deviation is eliminated at the third month-end following full rebalancing.

The purpose of a gradual implementation is obviously to smooth the trading over time and reduce the market impact related to rebalancing trading. We have modelled the expected market impact depending on trade size for equities, but our results are based on the assumption of fixed transaction costs of 10 bp for fixed-income trades. Hence the analysis does not fully take into account the fact that the transaction costs due to market impact most likely will be lower if the trading is smoothed out over a longer period of time. As a result of this, the real difference in rebalancing costs in the case of gradual implementation versus all-in-one-go is probably greater than our results predict.

Charts 22 and 23 show that the tracking error of the equity share is the same regardless of the decision to rebalance immediately or gradually. However, the transaction costs are lower for the alternative with gradual implementation.

Chart 21: Average rebalancing frequency and size

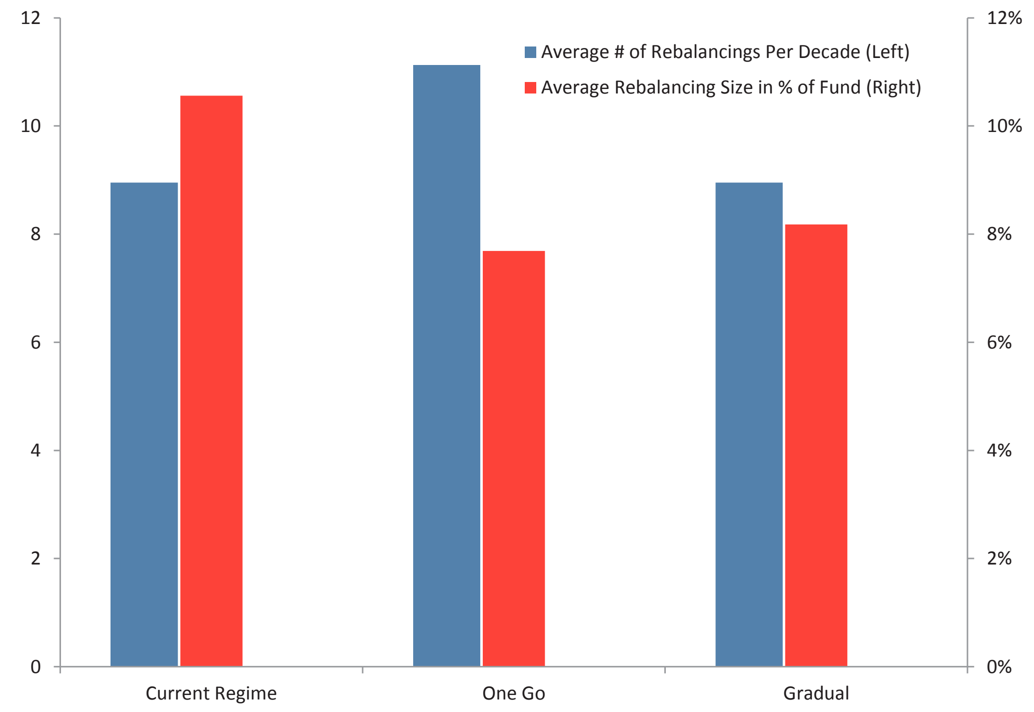

Chart 22: Rebalancing costs and tracking error of equity share

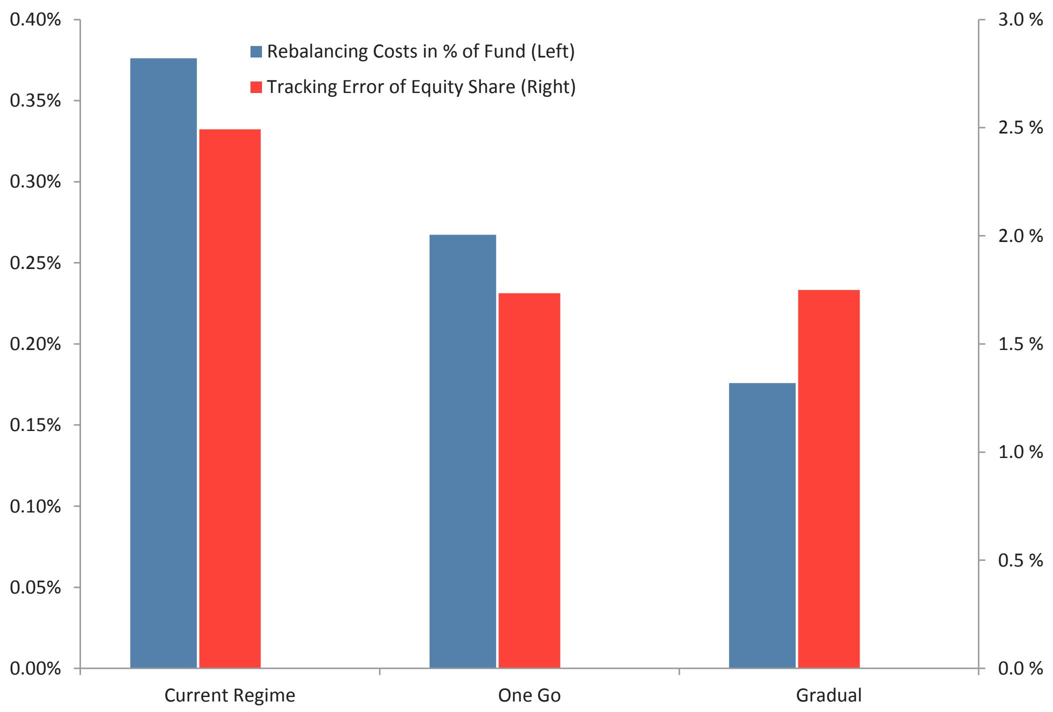

References

Arnott, R. and R.M. Lovell (1993): "Why? When? How often?", *Journal of Investing*, spring 1993, 5-10.

Campbell, J.Y. (2007): "Estimating the Equity Premium", *NBER Working Paper Series,* 13423.

Campbell, J. and L. Viceira (1999): "Consumption and Portfolio Decisions When Expected Returns are Time Varying", *Quarterly Journal of Economics*, 114, 433-495.

Campbell, J. and L. Viceira (2002): "Strategic Asset Allocation: Portfolio choice for long-term investors", Clarendon Lectures in Economics, Oxford University Press.

Cochrane, J. (2011): "Presidential Address: Discount Rates", *Journal of Finance*, 66(4), 1047-1108.

Constantinides, G. (1986): "Capital Market Equilibrium with Transaction Costs", *Journal of Political Economy*, 94, 842-62.

Davis, M.H.A. and A.R. Norman (1990): "Portfolio selection with transactions costs", *Mathematics of Operations Research*, 15, 676-713.

Donohue, C. and K. Yip (2003): "Optimal Portfolio Rebalancing with Transaction Costs", *Journal of Portfolio Management*, summer 2003, 49-63.

Goyal, A. and I. Welch (2007): "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction", *The Review of Financial Studies*, 21, 4.

Kritzman, M., S. Myrgren and S. Page (2009): "Optimal Rebalancing: A Scalable Solution", *Journal of Investment Management*, 7(1), 9-19.

Leibowitz, M. and A. Bova (2011): "Policy Portfolios and Rebalancing Behavior", *Journal of Portfolio Management*, winter 2011, 60-71.

Leland, H.E. (1996): "Optimal asset rebalancing in the presence of transactions costs", working paper, University of California, Berkeley Haas School of Business (March).

Markowitz, H. (1952): "The Utility of Wealth", *Journal of Political Economy*, 60, 151-158.

Lynch, A.W. and P. Balduzzi (2000): "Predictability and Transactions Costs: The Impact on Rebalancing Rules and Behavior", *Journal of Finance*, 55(5), 2285-2309.

Merton, R. (1969): "Lifetime Portfolio Selection under Uncertainty: The continuous time case", *Review of Economics and Statistics*, 51, 247-57.

Plaxco, L.M. and R.D. Arnott (2002): "Rebalancing a global policy benchmark", *Journal of Portfolio Management*, 28(2), 9-22.

Roll, R. (1977): "A critique of the asset pricing theory's tests. Part I: On past and potential testability of the theory", *Journal of Financial Economics*, 4, 129-176

Samuelson, P. (1969): "Lifetime Portfolio Selection by Dynamic Stochastic Programming", *Review of Economics and Statistics*, 51, 239-246.

Schroder, M. (1995): "Optimal portfolio selection with fixed transactions costs", working paper, Northwestern University.

Sharpe, W.F. (2009): "Adaptive Asset Allocation Policies", *Financial Analysts Journal*, May/June 2010.

Zakamouline, V.I. (2002): "Optimal Portfolio Selection with Transaction Costs for a CARA Investor with Finite Horizon", Working Paper, Institute of Finance and Management Science, Norwegian School of Economics and Business Administration.

Norges Bank Investment Management (NBIM)
Bankplassen 2
Postboks 1179 Sentrum
N-0107 Oslo

Tel.: +47 24 07 30 00
www.nbim.no