

Benchmark Rebalancing Calculations

Damien Laker

CompoundingHappens.com

Benchmark calculations for diversified funds involve making assumptions about rebalancing frequency. This paper explores the consequences of choosing different benchmark rebalancing frequencies, and shows how to calculate daily weights and returns for benchmarks that are rebalanced monthly, quarterly, or at any other interval.

This paper appeared in the *Journal of Performance Measurement*, Spring 2003, pp. 8-18. It is reproduced by permission of the *Journal of Performance Measurement*.

Damien Laker is the founder and CEO of CompoundingHappens.com.

Introduction

Modern portfolio theory teaches us the value of diversification in improving the reward/risk ratio of investment portfolios. Accordingly, many funds have benchmarks that combine more than one asset class (*e.g.* 50% bonds, 50% equities). Often, these are called “diversified” or “balanced” funds.

One problem for managers of diversified funds is that the proportion of the fund invested in each asset class might fluctuate away from the benchmark weights. To take an extreme example, “over the 50 years ending in December 1999, a portfolio that began with a 50%/50% stock/bond mix in 1950, with stock dividends reinvested in equities and bond coupons reinvested in bonds, would have drifted to a 98%/2% mix by December 1999”¹. Fund managers are generally not obliged to stick rigidly with the benchmark. Indeed, active managers are paid primarily for deviating from the benchmark in ways they believe will add value. However, if a benchmark is to have any meaning at all, it would clearly not be reasonable for a 50%/50% fund to drift into a 98%/2% allocation due to the long-term out-performance of equities relative to bonds.

Fund managers therefore need to consider the question of how they will rebalance a portfolio so that its asset allocation does not stray too far from the mandated weights. Generally speaking, very frequent rebalancing might lead to unacceptably high transaction costs (and maybe also capital gains tax liabilities), while very infrequent rebalancing might allow a fund to deviate too far from its benchmark weights (which would most likely generate higher levels of tracking error). Many authors have offered guidance on how to manage the trade-offs between risk, return, transaction costs, and taxation liabilities when carrying out portfolio rebalancing. See for example Balvers *et al.* (2000), Bernstein (1996), Buetow *et al.* (2002), Leland (2000), Perold and Sharpe (1998), and Plaxco and Arnott (2002).

One area that has received less attention is rebalancing in benchmark calculations. Just as the portfolio itself can be rebalanced to benchmark weights, the way one calculates the benchmark returns also entails a rebalancing assumption. For example, consider a benchmark that is 50% equities and 50% bonds. A common way to calculate the returns for this benchmark would be to use a weighted-sum of the monthly equity and bond index returns, where the weights were always 50%. This calculation method carries within it an assumption of monthly rebalancing. However, if one started to calculate the benchmark returns as a daily weighted sum (with weights of 50%), this would introduce an assumption of daily rebalancing.

An Unintended Consequence?

There is a powerful industry-wide trend toward daily performance measurement and attribution, due to the greater accuracy this provides. One unintended consequence of this trend might be that fund managers would start using daily benchmark rebalancing rather than monthly benchmark rebalancing. As we demonstrate below, daily benchmark rebalancing can produce results that are significantly different from a monthly rebalanced benchmark. Because of this, anyone who is adopting daily performance calculations might find it very beneficial to plan whether they wish to retain a monthly rebalancing calculation, even after they have adopted daily performance measurement and attribution.

Examples of Extreme Differences

If it is important to give careful consideration to portfolio rebalancing (and the abundance of papers published on this topic suggests that it is), then it must also be important to consider benchmark rebalancing (since the chief purpose of a benchmark is to evaluate manager skill by comparing the benchmark with the portfolio). Two specific examples for an Australian




Figure 1: Cumulative returns for equities and bonds in October 1989

fund whose benchmark is 50% equities and 50% bonds², over the period October 1989 to December 1999 (inclusive) will help to illustrate this point. In October 1989, a daily rebalanced benchmark would have out-performed a monthly rebalanced benchmark by 9 basis points. However, in September 1990, a daily rebalanced benchmark would have under-performed a monthly rebalanced benchmark by 14 basis points. For most fund managers, differences of this magnitude are not negligible. Therefore, the question of which benchmark rebalancing frequency to use is not one that can safely be left to chance.




Figure 2: Weight of equities (rebalanced monthly) in October 1989

The easiest way to understand why a benchmark with daily rebalancing performs differently from a benchmark with monthly rebalancing is to examine the two most extreme cases (October 1989 and September 1990).




Figure 3: Cumulative returns for equities and bonds in September 1990

October 1989

Figure 1 shows cumulative daily returns for bonds and equities during October 1989. It is immediately apparent that equities were dramatically more volatile than bonds during October 1989. Over the first few days of the month, equities quickly gained 2%. However, on October 16, equities dropped by 7.88%. The following day, they whipsawed back up by 3.95%. The resulting graph of cumulative returns gives a vivid impression of the high volatility of equities compared with bonds.

During October 1989, the daily rebalanced benchmark out-performed the monthly rebalanced benchmark by 9 basis points. Figure 2 helps to explain why this was the case. It shows the daily opening weight of equities in the benchmark on the assumption of monthly rebalancing (*i.e.* there was no benchmark rebalancing until the start of the following month).

It is easy to see that the unrebalanced weight of equities fluctuated very much in line with the fluctuations in the cumulative return of equities. On the other hand, a graph of the daily opening weight of equities with daily rebalancing would be very simple indeed: it would be a straight line whose value was 50% on every day. One can understand why the daily rebalanced benchmark outperformed the monthly rebalanced benchmark by considering the daily weights and returns. The monthly rebalanced benchmark was effectively overweight equities (at 50.5%) before the cumulative return for equities zig-zagged its way down to finish below minus 7% on October 16. Right at the low point, the monthly rebalanced benchmark was underweight equities (just below 48%), just before equities went up again.

In short, the daily rebalanced benchmark outperformed the monthly rebalanced benchmark during October 1989 because the market reversed direction so many times. The daily rebalanced benchmark had the same weight in equities on every day, but the monthly rebalanced benchmark was overweight equities when the equity market was at its high point, and underweight equities when the equity market was at its low point.

September 1990

September 1990 presents a different picture. As Figure 3 shows, September 1990 started off badly for equities and only got worse. Accordingly, as Figure 4 shows, the weight of equities in the monthly rebalanced benchmark moved steadily lower over the course of the month. On the other hand, the daily rebalanced benchmark (by definition) maintained a weight of 50% in equities at the start of every day during the month.

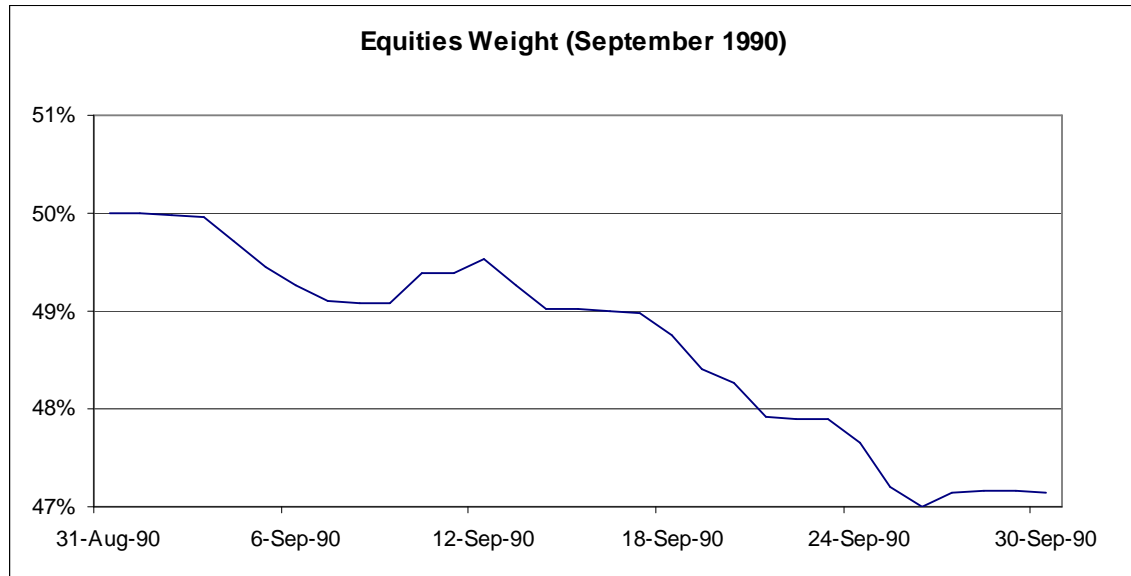


Figure 4: Weight of equities in September 1990

The end result was that the daily rebalanced benchmark under-performed the monthly rebalanced benchmark by 9 basis points during September 1990. The intuition behind this is that the monthly rebalanced benchmark progressively reduced its allocation to equities over the course of the month, while the daily rebalanced benchmark held an allocation of 50% in equities on every day of the month.

Mean Reversion and Trending

October 1989 and September 1990 both produced significantly different results when the rebalancing frequency changed from monthly to daily. An interesting question is whether one might expect the choice of rebalancing frequency to bias the benchmark returns in either direction over the long term. If this proved to be the case, fund managers could perhaps achieve a benchmark that was easier to beat, simply by a judicious choice of benchmark rebalancing frequency.

As it turns out, this question relates very closely to a classic question in finance theory: do market returns exhibit mean reversion? Mean reversion is a very intuitive concept. A mean-reverting market would be one that zig-zagged around a longer-term trend. This corresponds loosely to the popular idea that financial markets tend to overshoot from under-valuation to over-valuation in a cycle of booms and busts. October 1989 is a good example of a period where the equity market seemed to demonstrate mean reversion over the short term. In a mean-reverting market, more frequent rebalancing would tend to bias returns upwards. This is because the process of rebalancing would effectively mean selling out of sectors that have just increased in value (and which are due to revert down), and buying into sectors that have just decreased in value (and which are due to revert upwards). Going back to the example of October 1989, the daily rebalanced benchmark maintained a weight of 50% in equities regardless of market conditions, while the monthly rebalanced benchmark went overweight

equities when they were at their highest point, and went underweight equities when they were at their lowest point.

Conceptually, the opposite of a mean-reverting market would be a trending market. In a trending market, an asset that outperformed in one period would tend also to outperform in subsequent periods. September 1990 was a good example of a period where the equity market seemed to demonstrate trending over the short term.

Over different periods, financial markets sometimes demonstrate trending behavior, and they sometimes demonstrate mean-reverting behavior. Despite the popular perception that financial markets (especially equities) routinely overshoot in a cycle of boom and bust, the literature does not support this very strongly. Kim *et al.* (1999) find only weak evidence for mean reversion. Balvers *et al.* (2000) do find evidence of mean reversion, but they cite a number of other studies that were unable to demonstrate mean reversion. The mean reversion found by Balvers *et al.* had a half-life of three to three and a half years, which would suggest that one would have to rebalance the benchmark over a cycle of several years for this mean reversion to have an effect. This is likely to be unpalatable for other reasons.

Table 1 shows the returns for a benchmark of 50% bonds and 50% periods for a range of rebalancing periods from daily rebalancing right through to annual rebalancing. With the possible exception of annual rebalancing, there is very little difference between these results. Based on these numbers, one would not be able to conclude that the choice of rebalancing frequency (within reasonable limits³) would reliably bias the benchmark returns in either direction. Therefore, this would most likely not be an important criterion in determining which benchmark rebalancing frequency was suitable.

Table 1: Results of Different Rebalancing Frequencies October 1989 to December 1999

	Rebalanced Daily	Rebalanced Monthly	Rebalanced Quarterly	Semi-Annually	Rebalanced Annually
Total Return	220.8%	220.2%	220.4%	221.3%	224.9%
Annualized Return	12.04%	12.02%	12.03%	12.06%	12.18%

Range Rebalancing

The examples so far all fall into the category of *calendar rebalancing* (*i.e.* the benchmark is rebalanced based on months, quarters, *etc.*) Another option would be *range rebalancing*, where the trigger for rebalancing the benchmark would be when the benchmark weights moved more than a specified distance from the strategic weights⁴. For some investment processes, a range rebalanced benchmark may be appropriate. However, a disadvantage of this choice would be that a range rebalanced benchmark is more difficult to calculate, and also harder for clients to understand.

How To Calculate Daily Weights and Returns

It has been traditional for many performance analysts to do most calculations in spreadsheets using monthly weights and returns. As we have already described, this brings with it an assumption of monthly benchmark rebalancing.

When a firm chooses to start using daily performance measurement and attribution, it is not desirable for them to be forced into using a daily benchmark rebalancing assumption purely because they are doing daily calculations. Indeed, for the sake of consistency with past practices, and also for the sake of conformity with industry expectations, it is most likely desirable for them to keep using monthly benchmark rebalancing. This means that it is necessary to calculate daily weights and returns that are consistent with the assumption of

monthly benchmark rebalancing. For the sake of brevity, the following example uses a weekly benchmark rebalancing frequency, even though monthly is the most common case in practice. The only difference between calculations for monthly rebalancing and any other frequency is simply how to determine the first day in any rebalancing period.

For the sake of clarity, we explain the calculation using an example⁵ rather than using algebra.

In this example, there are two sectors in the benchmark (Sector 1 and Sector 2). The mandated benchmark weights are 60% and 40%, with weekly rebalancing. The returns and weights are shown in Table 2.

Table 2: Daily Weights and Returns for a Weekly Rebalancing Example

Week	Day	Benchmark Return			Start of Day Weight		End of Day Weight	
		Sector 1	Sector 2	Total	Sector 1	Sector 2	Sector 1	Sector 2
1	1	-1%	4%	1.00%	60.00%	40.00%	58.81%	41.19%
1	2	4%	-2%	1.53%	58.81%	41.19%	60.24%	39.76%
1	3	-1%	4%	0.99%	60.24%	39.76%	59.06%	40.94%
1	4	4%	-2%	1.54%	59.06%	40.94%	60.49%	39.51%
1	5	-1%	4%	0.98%	60.49%	39.51%	59.30%	40.70%
1	6	4%	-2%	1.56%	59.30%	40.70%	60.73%	39.27%
1	7	-1%	4%	0.96%	60.73%	39.27%	59.55%	40.45%
2	1	3%	1%	2.20%	60.00%	40.00%	60.47%	39.53%

These returns and weights cover a full week in which the markets mean-revert in accordance with a two day cycle, and the first day of the following week⁶.

Start of Day Benchmark Weight

The start of day benchmark weight for a sector on day N is determined by this rule:

- IF N is the first day in the rebalancing period:
- THEN it is the mandated weight (e.g. 60% for Sector 1)
- ELSE it is the end of day benchmark weight for day N-1.

For example, on day 1 of each week, the start of day weights are 60% and 40%, but on the other days, they are the end of day weights for the preceding day.

This rule shows that it is necessary to have a method for calculating the end of day benchmark weight on any day.

End of Day Benchmark Weight

The easiest way to think about end of day weights is to start at the beginning.

On the first day of the rebalancing period, the opening weights will be 60% and 40%. An easy way to do the calculation is to consider the weights as units of currency rather than percentages. For example, if the start of day sector values were \$60 and \$40, and the sector returns were -1% and 4%, then the end of day sector values would be \$59.40 and \$41.60. The total end of day value would be \$101.00, and hence the end of day sector weights for Sector 1 and Sector 2 respectively would be $\$59.40/\$101 = 58.81\%$ and $\$41.60/\$101 = 41.19\%$ respectively.

In practice, you can calculate the end of day benchmark weight for each day by using the start of day benchmark sector weights for each day as well as the benchmark sector returns on that day, or alternatively you can use the mandated weights (in this case, 60% and 40%) along

with the cumulative benchmark sector returns from the start of the rebalancing period until the end of that day. These two calculation methods are equivalent, and the choice is simply one of convenience.

Benchmark Return

The overall benchmark return on any day is simply the weighted sum of:

- the start of day benchmark weights for each sector; and
- the benchmark return for each sector.

For example, on day 1 of week 1, the overall benchmark return is $(60\% \times -1\%) + (40\% \times 4\%) = 1.00\%$.

Turnover and Implied Transaction Costs

In the most extreme case, where you were running an investment strategy that sought to exactly replicate the benchmark, you would be forced to incur portfolio turnover (and hence transaction costs) every time one benchmark rebalancing period finished and a new one began.

For example, in Table 2, the end of day weights on day 7 of week 1 were 59.55% and 40.45%. In accordance with the rebalancing assumption, the start of day weights on the following day must be 60% and 40%. Strictly speaking, in a fully replicated portfolio, one would have to sell 0.45% of the fund out of Sector 2 and buy an equal amount into Sector 1 in order to implement the rebalance. This would entail a “two way” turnover (i.e. buying and selling) of 0.45% of the portfolio. If one assumes that the transaction costs for buying are similar to those of selling, it is more useful to think of “one way” turnover (which would be 0.90% in this case). Supposing that “one way” transaction costs were 100 basis points, this single weekly rebalance would incur 0.9 basis points in transaction costs.

In practice, of course, it would be extremely rare to encounter a fund that attempted to replicate its benchmark in such a mechanical way as this example suggests. Portfolio managers generally don’t mechanically rebalance to benchmark. Also, in most funds the portfolio sector returns do not exactly track the benchmark sector returns.

Furthermore, a fund’s net contributions or withdrawals will in any event generate a need to incur transaction costs. Depending on the size of these contributions or withdrawals, the fund manager can use them to reduce or eliminate “benchmark-driven” turnover, by using the contribution or withdrawal as an opportunity to move the portfolio weights closer to the benchmark weights. In the example discussed above, if the fund received a net inflow equal to 0.9% of its value, the fund manager could rebalance the portfolio to benchmark weight by directing that inflow into Sector 1. Similarly, if there was a net outflow equal to 0.9% of the fund value, the fund manager could meet that outflow by selling assets from Sector 2.

Hence, any calculations that one does about “benchmark driven” portfolio turnover are based on a worst case, where the fund manager chooses to rebalance the portfolio at the end of every benchmark rebalancing period, and where there are no net cashflows to mitigate the situation⁷.

However, some quick calculations can at least give us some idea of the benchmark turnover generated by each benchmark rebalancing assumption. Table 3 shows the annualized “one-way” turnover for each different rebalancing assumption, for the fund that is 50% bonds, 50% equities. Clearly, daily rebalancing is much more likely than any other frequency to generate significant “benchmark-driven” portfolio turnover (and hence transaction costs). There may be some transaction cost benefit in using quarterly (rather than monthly) benchmark rebalancing, but for rebalancing frequencies longer than that, the benefit is very marginal indeed.

These numbers would suggest that caution would be in order before adopting a daily benchmark rebalancing calculation.

Table 3: Annualized “One Way” Benchmark Turnover January 1990 to December 1999⁸

	Rebalanced Daily	Rebalanced Monthly	Rebalanced Quarterly	Semi- Annually	Rebalanced Annually
Annualized Turnover	74.50%	16.61%	8.69%	6.49%	6.47%

Tracking Error of Different Rebalancing Frequencies

Since monthly benchmark rebalancing has been a *de facto* standard, it is interesting to consider how big the impact might be of choosing an alternative rebalancing frequency. A simple way to do this is by calculating the tracking error of (say) a daily rebalanced benchmark compared with a monthly rebalanced benchmark.

Table 4 shows the annualized tracking error (estimated using monthly returns) using the 50% bonds 50% equities data for the period October 1989 to December 1989. Compared with month rebalancing, daily rebalancing introduces about 7 bps of annualized tracking error, while quarterly rebalancing introduces about 11 bps of annualized tracking error. As you would expect, the tracking error gets larger for rebalancing periods longer than quarterly.

These numbers are not completely negligible, but neither are they enormous. Depending on the characteristics of any particular fund, one could make a judgment about how material the impact would be of switching from a monthly rebalanced benchmark to some other rebalancing frequency.

Table 4: Annualized Tracking Error in basis points (compared with Monthly Rebalancing)

	Rebalanced Daily	Rebalanced Monthly	Rebalanced Quarterly	Semi- Annually	Rebalanced Annually
Tracking Error (bps)	7.26	0.00	10.76	19.46	31.12

Investment Mandates

It is customary for managed funds to have an investment mandate, which is a legally-binding document specifying the relationship between the fund manager and the investors. Typically, the mandate will specify the permitted investments, and it will also specify a benchmark against which the fund’s performance can be measured or attributed.

In the author’s experience, when mandates specify a diversified benchmark, they rarely if ever specify a benchmark rebalancing frequency. As discussed above, monthly benchmark rebalancing has (almost by default) been a *de facto* standard in many situations. The personnel who are responsible for specifying investment mandates may like to consider whether it would improve the clarity of these documents if they did indeed specify a benchmark rebalancing frequency. Particularly when there are performance fees or tracking error constraints at stake, the choice of benchmark rebalancing frequency could be of considerable importance.

Attribution

It is almost universally agreed that daily performance measurement and daily attribution are – at least in principle – improvements on monthly calculations. A naïve implementation of daily measurement and attribution might inadvertently fall into the trap of using a daily rebalancing calculation.

There are several reasons why a monthly benchmark rebalancing assumption might very well be the preferred choice even when the calculations are being done daily. These include:

- consistency with traditional methods;
- comparability with independent calculations that use monthly returns in spreadsheets;
- relatively low risk of unacceptable “benchmark-driven” transaction costs; and
- avoiding the excessive drift from strategic benchmark weights that might be involved in (for example) an annual benchmark rebalancing assumption.

In order to do daily attribution on diversified funds with a benchmark rebalancing frequency other than daily, one would need to use the start of day weights and benchmark returns that are illustrated in Table 2. This is an important requirement to consider when building or buying an attribution system.

Conclusion

There is an extensive literature on portfolio rebalancing strategies. However, the literature does not show much interest in benchmark rebalancing calculations. As investment firms make the transition toward daily calculations, this is a topic they should consider, since otherwise they may end inadvertently switching to a daily benchmark rebalancing assumption when they switch to daily performance calculations.

Traditionally, most benchmark calculations have used a monthly or quarterly rebalancing assumption. Taking account of all the issues discussed in this paper, these rebalancing assumptions seem perfectly reasonable. Hence it is likely that many firms who consider this issue will wish to retain a monthly or quarterly benchmark rebalancing assumption, even when they are doing daily calculations. This paper explains how to calculate the daily benchmark weights and returns that are consistent with non-daily benchmark rebalancing assumptions.

References

Balvers, Ronald, Yangru Wu, and Erik Gilliland “Mean Reversion Across National Stock Markets and Parametric Contrarian Investment Strategies”, *Journal of Finance*, **LV**:2, April 2000.

Bernstein, Peter L. “The Rebalancing Bonus: Theory and Practice” *Efficient Frontier*, September 1996.

Buetow, Jr., Gerald W., Elaine Hunt, Ronald Sellers, Donald Trotter, and Willie A. Whipple Jr., “The Benefits of Rebalancing” *Journal of Portfolio Management* **28**:2 Winter 2002, pp. 23–32.

Kim, Chang-Jin, James C. Morley, and Charles R. Nelson, “Time-Varying Mean Reversion in Stock Prices: Evidence and Implications for Market Efficiency”, www.stat.fi/isi99/proceedings/arkisto/varasto/kim_0290.pdf, 1999.

Leland, Hayne E. “Optimal Portfolio Implementation with Transactions Costs and Capital Gains Taxes”, www.haas.berkeley.edu/faculty/pdf/leland2000.pdf, 2000.

Perold, Andre F. and William F. Sharpe “Dynamic Strategies for Asset Allocation” *Financial Analysts Journal*, January/February 1998, pp. 16-27.

Plaxco, Lisa M. and Arnott, Robert D. “Rebalancing a Global Policy Benchmark” *Journal of Portfolio Management*, **28**:2 Winter 2002, pp. 9-22.

¹ Plaxco and Arnott (2002) p. 10.

² The benchmark is 50% ASX All Industrials (a broad index of Australian equities) and 50% UBSW All Maturities (a broad index of Australian bonds). There is no reason to suppose that Australia would be sufficiently different from other countries that the conclusions drawn in this paper would be materially different if the example was based on data from another country.

³ As Plaxco and Arnott's 50 year example demonstrates, a very long benchmark rebalancing interval (measured in years) would tend to bias the benchmark returns upwards since the weight of high risk, high return asset classes (such as equities) would drift upwards. However, for practical purposes, it would probably not be palatable to clients to calculate a benchmark using a rebalancing frequency longer than one year.

⁴ See Plaxco and Arnott (2002) p. 10 for a comprehensive list of rebalancing alternatives.

⁵ A spreadsheet containing the detailed calculations for this example is available at <http://www.CompoundingHappens.com/benchmarks.htm>. The spreadsheet is called "Simple Example of Benchmark Rebalancing.xls".

⁶ The example spreadsheet includes a comparison of these weekly-rebalanced benchmark returns with some daily-rebalanced benchmark returns.

⁷ In practice, portfolio managers also have other tools at their disposal (such as optimizers and the use of derivatives) to mitigate turnover and hence transaction costs.

⁸ These turnover comparisons exclude the 3 months October 1989 to December 1989, because the only fair basis to compare annual rebalancing with other rebalancing frequencies is over an integral number of years.