Original Article

How much value should you expect to gain or lose by replacing your investment manager?

Received (in revised form): 10th April 2012

Robin Penfold

is a Senior Investment Consultant at Towers Watson Limited, UK. In this role, he writes papers on a wide range of strategic investment subjects. Amongst its services, Towers Watson provides investment advice to pension funds, foundations and other large institutional funds.

Correspondence: Robin Penfold, Towers Watson Limited, Watson House, London Road, Reigate, Surrey RH2 9PQ, UK
E-mail: robin.penfold@towerswatson.com

ABSTRACT Goyal and Wahal showed that plan sponsors typically lose value when firing an underperforming manager and hiring its replacement. Experience suggests that most investors appreciate this research but then ignore it when making decisions. To redress this problem, we build a model that shows an investor its expected loss from hiring and firing its managers for performance reasons. By knowing this likely loss, we hope that investors will make fewer performance-related decisions and therefore improve their returns. Another useful part of the model is that it can also project an investor's likely gain from 'taking profits' from a manager with exceptional outperformance.

Journal of Asset Management (2012) **13,** 243–252. doi:10.1057/jam.2012.10;
published online 5 July 2012

Keywords: portfolio management; portfolio transitions; mean reversion

INTRODUCTION

Goyal and Wahal (2008) show that plan sponsors tend to lose value when firing managers that have performed poorly. The authors consider the hiring and firing decisions of 412 'round-trip' transitions by US plan sponsors between 1996 and 2003. They study the performance of the incumbent manager in the 3 years before it was terminated and the performance of its replacement in the following 3 years. Where poor performance led to the manager being fired, they find that plan sponsors lost 0.79 per cent of cumulative potential value in the 3 years after the transition (before accounting for transition costs).

As investment consultants to institutional investors, these findings trouble us. Moreover, the reasons for these performance-related decisions are irrelevant – we expect a loss, whether these decisions are inadvertent, obfuscated or deliberate. (Later in this article, we show how trustees can diagnose whether their decisions have been performance related.[1]) Even worse, our experience is that many investors still hire and fire managers for performance reasons, even in the knowledge that they are expected to lose value.

Why could this be? To us, a large contributor is the inconsistency of the information that investors get when they

make hiring and firing decisions. Think about three key questions that are often asked in these instances:

1. How should our manager perform in the future (compared with other managers that we may hire)?
2. How has our manager performed for us (and is this performance what we were led to expect)?
3. How much value will we likely lose by switching managers for performance reasons?

In the first two questions, investors can get a *quantified* answer for *their experience* with the manager. With the third question, they can only get information about the average value that US plan sponsors lost a decade ago when they changed managers for performance reasons.

As the third answer is less tangible to investors than the first two answers, we believe that they often de-emphasise it in their hiring and firing decisions. From a practical perspective, that response is understandable. Yet if the answer to the third question could be made more tangible to investors – and quantify the expected costs to *them* of switching *their* manager for performance reasons – then they might take it into account when deciding a manager's fate. Doing so would make them less likely to lose value from poor decision-making.

The key issue for an investor therefore becomes: How can we quantify the value that *you* should expect to gain or lose by switching *your manager* for performance reasons?

It turns out that quantifying this expected gain or loss of value is not so difficult.

CALCULATING THE GAIN OR LOSS THAT YOU SHOULD EXPECT FROM FIRING YOUR INVESTMENT MANAGER

Our model emulates the structure of the Goyal and Wahal study, but looks forwards rather than backwards. We provide more detail on the model's assumptions in a later example, but can summarise its process in four steps:

1. *Build return scenarios*. Build many random, but collectively sensible, scenarios of a manager's excess returns over two consecutive 3-year periods (pre-transition and post-transition).
2. *Transition the portfolio*. In each scenario, transition from the incumbent manager to the replacement manager after 3 years.
3. *Measure performance*. In each scenario, record the post-transition 'performance impact' and the pre-transition performance of the incumbent manager.
4. *Forecast your likely performance impact*. Use the strong linear pattern between these two recorded statistics to forecast the likely performance impact of replacing your manager, given its recent performance.

Stage 1 – Build return scenarios

We begin by studying many random scenarios of the manager's excess returns. Some excess returns are excellent, some are awful, but most are in between. Put together, they resemble the manager's projected return distribution: the average is the manager's expected excess return and will vary according to the manager's tracking error. To make our projection as realistic as possible, we also assume that the manager's excess returns mean-revert a little over time. This means that an above-average return in one period will typically lead to a below-average return in the following period.

As we explain later, we do not just study the excess returns of the initial manager – we also project scenarios for three managers that could replace it. In the simplest case, we assume that these possible replacements have the same *ex ante* return distribution as the initial manager, albeit with different investment styles.[2]

For the purposes of this article, we consider 10 000 of these possible scenarios. In each scenario, and for each of the *four* managers, we project the manager's excess returns over 6 consecutive years. We then assume that the investor fires the initial manager after 3 years. This leaves us with two discrete 3-year periods of performance: 'pre-transition' and 'post-transition'.

Stage 2 – Transition the portfolio

At the point of transition, we assume that the investor

- searches for a replacement manager from the three potential candidates; (We assume a shortlist of three candidates to reflect our typical experience with institutional investment boards.)
- excludes a fixed number of these candidates, based upon their pre-transition performance; (For example, the investor may choose to exclude the manager with the worst pre-transition performance of the three candidates.)
- randomly hires the replacement manager from the remaining candidates; and
- incurs costs for transitioning the portfolio of the initial manager to the replacement manager.

Stage 3 – Measure performance

We then measure the post-transition performance of the replacement manager, reflecting transition costs. Clearly, values above zero represent gains against the benchmark. Stated alone, however, they only convey part of the picture. What we also need to know is what performance would have been had the investor retained, rather than fired, its initial manager. For that reason, we compare the post-transition performance of the replacement manager (including transition costs) with that of the former manager (excluding transition costs). This relative return shows the performance impact of changing the manager.

We then record two values: this performance impact of changing the manager and the pre-transition performance of the former manager. We then repeat these calculations for all 10 000 scenarios. This provides us with 10 000 values of these two recorded statistics. From these values, we can see if there is a pattern between the pre-transition performance of the former manager and the performance impact incurred.

Stage 4 – Forecast your likely performance impact

We find a statistically significant linear pattern between these two recorded statistics, and assume that this pattern will persist. This enables us to calculate the expected performance impact of firing a manager that has the excess return distribution modelled in Stage 1, so long as we know its achieved excess return over the last 3 years.

MODEL SUMMARY

With five simple ingredients, we can calculate the value that we would expect an investor to gain or lose if it replaced its investment manager today. Those five ingredients are:

1. The distribution of projected excess returns for the investor's existing manager.
2. The distribution of projected excess returns for the three candidate replacement managers.
3. The cost of transitioning the portfolio from the existing to the replacement manager.
4. The extent to which the investor considers past performance when hiring replacement managers.
5. The latest 3-year excess return of the investor's existing manager.

As we show later, the results from our model tend to support the lessons of Goyal and Wahal. For example, we find that investors

should take great care to avoid replacing managers for performance reasons alone, particularly when they expect their manager to have a high information ratio and mean-reverting excess returns.

We now illustrate our model with an example, featuring a hypothetical investor and manager.

SHOULD THE TRUSTEES FIRE ABC INVESTORS?

ABC Investors (ABC) has managed global equity assets for the XYZ Pension Trust for some time. In the early days of its appointment, its performance was fine. Then performance dipped. Over the last 3 years, ABC underperformed its benchmark by −2.0 per cent a year.

Imagine that we act as investment consultant to the Trustees of the XYZ Pension Trust.

Despite ABC's dip in performance, we still rate it highly. After accounting for our own fallibility, we expect it to outperform by 1.2 per cent a year (after fees). Of course, we would not expect outperformance in every year, and therefore assign ABC a tracking error of 6 per cent a year.

Even with our assurances, ABC's appointment is in jeopardy. For that reason, the Trustees have asked us to quantify how much value they could gain or lose by replacing ABC today. To do that, we assume that ABC and its potential replacements have:

- an expected net excess return of 1.2 per cent a year relative to their benchmark;
- a standard deviation of excess returns of 6 per cent a year;[3]
- excess returns that follow a lognormal distribution;[4]
- excess returns that exhibit a modest level of mean-reversion[5] (specifically, an annual excess return autocorrelation of −0.2); and
- different investment styles (and thus low correlations between the managers' excess returns).

We also assume that the Trustees:

- assess 3-year performance when making hiring and firing decisions;
- replace managers by considering shortlists of three candidate managers (but hardly ever hire the worst-performing candidate). As such, we assume that the replacement manager is chosen at random from the two best-performing candidate managers over the prior 3 years; and
- spend 0.5 per cent to transition their portfolio from the existing manager to its replacement.

We then model 10 000 scenarios of the future excess returns of ABC over two consecutive 3-year periods. In each case, we record two values: the performance impact of the transition and the pre-transition performance of ABC. We show these two values for every scenario in Figure 1.

Two points are clear. First, no clear pattern emerges immediately, as the chart contains a lot of noise. However, as one looks closer, it becomes clear that there is an upwards slope to the dots in the chart. It does not explain too much – there is still a lot of noise – but it is pronounced.

This understanding can be confirmed statistically. We perform a regression on the data in the chart and find that there is a highly significant linear pattern between the performance impact of the transition (I) and ABC's pre-transition performance (P).[6] We express this pattern in its functional form in equation (1), showing the relevant t-statistics and measures of statistical significance:[7]

$$\begin{aligned} I &= -0.003 + 0.08\ P + \varepsilon \\ &\quad (-8.0) \qquad (5.7) \\ R^2 &= 0.3 \text{ per cent} \end{aligned} \tag{1}$$

(P-values for the regression coefficients are significant at the 0.1 per cent level)

To consider the implications of this regression, we now use Figure 2 to zoom in on the previous chart and remove the noise

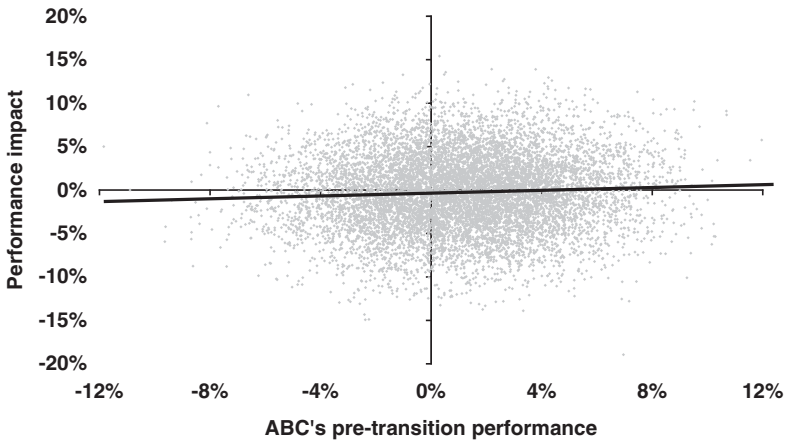

Figure 1: Performance impact associated with ABC's pre-transition performance (results from 10 000 scenarios).

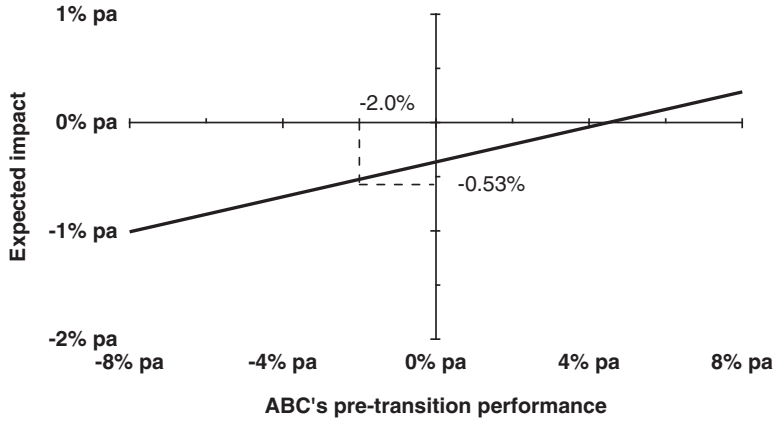

Figure 2: Expected performance impact associated with ABC's pre-transition performance.

of the individual scenarios. By doing so, we are only considering the *typical* – or likely – performance impact.[8]

We see that the likely performance impact of replacing ABC rises if its pre-transition performance worsens. Much of that worsening is due to our assumption about the mean-reversion of ABC's excess returns – a trait that is usually apparent in the returns of many real-world mandates.

Getting back to the example, we know that the XYZ Pension Trust's investment with ABC underperformed by −2.0 per cent a year over the last 3 years. A quick glimpse at the line in Figure 2 shows the implications for the Trust if the Trustees were to fire ABC today. That is, we would expect this decision to lose the Trust −0.53 per cent a year on this mandate over the next 3 years (or −1.57 per cent, when compounded). Of course, we doubt that the actual loss or gain will be −1.57 per cent over this time, but it is just as likely to be above this value as it is to be below it.

The precise value of this expected loss would be useful information to the Trustees. Armed with it, we would like to think that they would be less likely to fire ABC now – and therefore less likely to lose value.

ANOTHER INSIGHT OF THE MODEL: WHEN TO TAKE PROFITS

Figure 2 shows us more than just the expected loss arising from the Trust replacing ABC after a bout of underperformance. Looking at the right-hand side of the chart, it also shows us the value that the Trust can expect to gain from replacing – or rebalancing – a manager that has had exceptional outperformance. In other words, it helps us to understand when the Trust should 'take profits' from an investment with ABC. In this instance, that would be when ABC has outperformed by more than about 4 per cent a year for 3 years.

Investors that are wary of the perils of hubris would be well advised to consider such an analysis on their outperforming managers. For each of their managers, for example, they could use this model to determine the level of excess returns over 3 years that would suggest that assets be added or removed from the manager. In times of exceptional manager performance, the investor should then be more able to take decisive and dispassionate action (which would, by definition, be contrarian in nature).

WHAT ASPECTS OF A TRANSITION MOST INFLUENCE ITS LIKELY PERFORMANCE IMPACT?

Just like asset allocation models, the value of this model extends beyond its ability to project an expected outcome. For example, by repeatedly changing the value of one of the model's inputs, we can see whether that input strongly influences the model's results (the likely performance impact of replacing a manager for performance reasons).

Having conducted these 'sensitivity tests', we find that these three inputs have a strong influence on the model's results. They are the:

- *expected excess return* of the managers;
- number of possible *replacements excluded from consideration on performance grounds*; and
- *mean-reversion of excess returns* of the managers.

We now discuss our sensitivity tests for each of these cases, comparing them to the 'base case' of the ABC example above.

The expected excess return of the managers

In Figure 3, we plot the best-fit lines from three different runs of the model. As before,

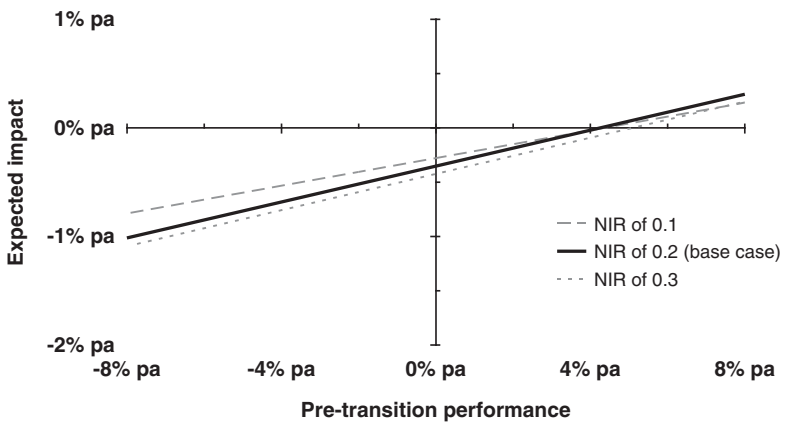

Figure 3: Expected performance impact for a given level of pre-transition performance (for different levels of manager NIR).

we continue to set equal expected excess returns for the initial manager and the potential replacement managers. In each case, however, we alter the expected excess returns. We consider net expected excess returns from 0.6 per cent to 1.2 per cent to 1.8 per cent a year, representing a net information ratio (NIR) of 0.1, 0.2 and 0.3.

We find that the expected excess return of the managers strongly influences the model's results. As one might imagine, the higher the expected excess return, the worse the performance impact for firing the manager at a given level of underperformance. After all, a given level of underperformance in these cases will be further from the manager's expected return, implying that any future reversion to the mean will be more pronounced.

The number of possible replacements excluded from consideration on performance grounds

How much do the model's results depend upon the way that the investor hires replacement managers? We examine this question in Figure 4. In doing so, we assume that three possible managers are available, with ever increasing numbers of managers ruled out for performance reasons. In the extreme case, the investor merely chooses the best performer from the previous 3 years. In the base case, the investor randomly chooses the replacement manager from the two candidates with the best recent performance.

We already know that investors can lose value by firing managers that have recently underperformed. With Figure 4, we now also know that they can lose value by hiring managers for performance reasons, particularly if they usually hire the best-performing candidate manager. (Experience suggests that, for whatever reason, investors often choose the best performer from a shortlist of candidates.)

Figure 4 tells us more, though. We can find the impact of hiring solely for performance reasons from the difference between the performance impacts of 'choosing from the top 1' and 'choosing from the top 3'. We can also find the impact of firing managers for performance reasons by taking the annualised performance impact of 'choosing from top 3' and subtracting a third of the transition cost. By doing this, we tend to find that the impact of firing an underperforming manager exceeds that of hiring a manager for performance reasons, although not by much.

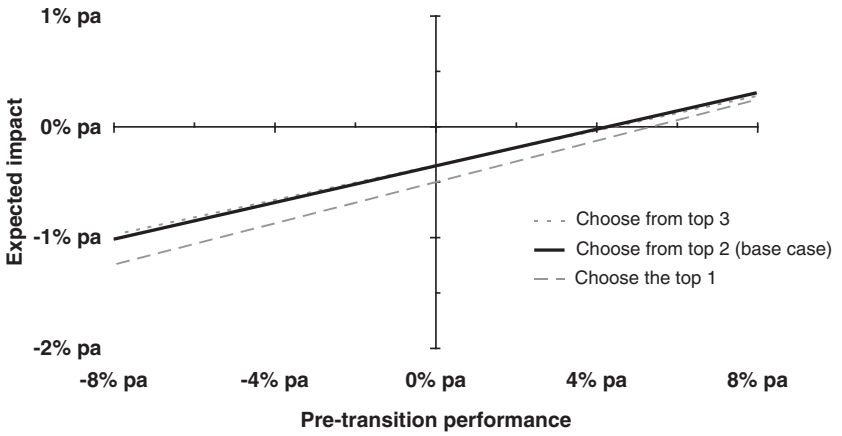

Figure 4: Expected performance impact for a given level of pre-transition performance (for different ways of finding replacement managers).

The mean-reversion of managers' excess returns

When replacing underperforming managers, we can attribute an investor's expected loss to three factors:

- the one-off cost of transitioning the portfolio from the initial to the replacement manager;
- the opportunity loss from firing an underperforming manager (whose excess returns are expected to 'bounce back' soon – that is, revert to their mean); and
- the anticipated loss arising from hiring a manager with relatively good performance (whose excess returns are expected to 'fall back' soon – that is, revert to their mean).

Given the latter two factors, we would expect the model's results to be sensitive to a change in our assumption for the autocorrelation of the managers' excess returns. We confirm that this is the case in Figure 5. (Recall that, in our base case, we follow the implications of Goyal and Wahal's results and assume a slight mean-reversion of excess returns. Specifically, we set the annual autocorrelation of excess returns to be −0.2.)

To confirm our intuition about the three factors, we ran the model with parameters that did not induce mean-reversion in the excess returns of any managers. We did so by assuming that the initial manager's past performance was as expected – there was no autocorrelation of manager excess returns – and that replacement managers were chosen randomly from the three available candidates. In this instance, we found that the model expects a loss of −0.18 per cent a year, which is essentially one-third of the cost of the transition.

IGNORANCE IS NO LONGER A CREDIBLE DEFENCE

We began this article by noting that investors tend to lose value when they replace managers for performance reasons. We went on to say that many investors understand this point but still tend to replace managers in this way, inadvertently or otherwise. Perhaps we were being kind, but our reasoning for this irrationality was that investors could not quantify their likely loss from these actions. Now armed with our model, investors will no longer be able to rely on this reasoning for such irrationality.

Those who do not have access to our model can also learn from its results, however. Whilst seemingly obvious, investors can now be more confident about being wary of replacing managers for performance

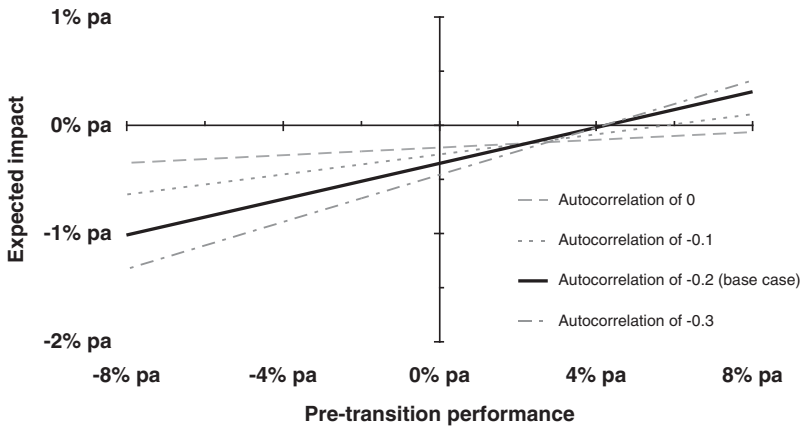

Figure 5: Expected performance impact for a given level of pre-transition performance (for different levels of mean-reversion in the managers' excess returns).

reasons, especially if their excess returns have a high expected value or are likely to mean-revert. Caution is also recommended to avoid the related error of hiring managers for performance reasons.

Either way, if you replace managers for performance reasons, ignorance is no longer a credible defence.

ACKNOWLEDGEMENTS

I would like to thank George Aristides, Robert Brown, Stephen Miles and Charlie Ullman for their help with this article. All errors are my own.

NOTES

1. Many investors find it hard to acknowledge that their manager replacement decisions are performance related. After all, the interaction between a manager's past performance and its appeal to an investor is sometimes subtle. For example, one would expect managers to make better presentations to investors after good performance (as they have more confidence and can use better examples). Moreover, all managers have some weaknesses in their process; they just tend to come under more focus when the manager has recently underperformed. Given these difficulties, we encourage all investors to analyse the past performance of their recent manager appointments and terminations. For each appointment in the last 3 years, for example, investors could rank the performance of the shortlisted managers in the 3 years before the hiring decision, noting the position of the manager that was actually hired. If the investor typically chose the 'best performer of the three', then the investor's preference for past performance is unlikely to be accidental. In a similar vein, investors could study the managers that they recently terminated. They could calculate the average information ratio of these managers in the 3 years before their termination. A negative average would again suggest that the investor's preference for past performance is more than accidental.
2. To aid simple exposition, we assume that these possible replacements have the same return distribution as the incumbent manager. However, this need not be the case.
3. Taking this assumption alongside the one directly before it implies a net information ratio of 0.2. We consider this appropriate for the base case, given our experience of client portfolios that contain our highly rated managers. Given Penfold (2004), we form our tracking error estimate through a qualitative risk assessment of the manager's portfolio, rather than relying exclusively on the results of risk models.
4. The model works on any distribution of excess returns. We have tested the Variance-Gamma distribution, as it provides more realism through its fat tails. In our base case, however, we opted for the lognormal distribution, given its broad appeal and the ease with which its parameters can be interpreted.
5. We assume a slight level of mean-reversion in our base case. We do so because our experience – and that of Goyal and Wahal – shows that investors tend to lose value when they replace managers for performance reasons. Of course, if the excess returns of most managers typically trended – rather than mean-reverted – then these performance-following investors would tend to gain value. (After all, they would buy future winners and sell future losers.) As no gain in value is generally observed, we conclude that the excess returns of most managers exhibit some level of mean-reversion. Support for this view is also generally available from empirical studies that remove the effects of survivor bias. In one such study by Busse *et al* (2010), the authors find that better manager excess returns from the previous year tend to lead to worse excess returns in the following 3 years.
6. Furthermore, polynomial forms barely improved upon the explanatory power of the linear regression. In addition, the *P*-values of all the regression coefficients in this article are significant at the 0.1 per cent level, with the exception of the case with no autocorrelation in Figure 5. The *P*-value of this regression's slope coefficient is significant at the 20 per cent level.
7. As we noted, the 'best-fit' line only explains a small amount of the overall variation in this chart. However, it can lead to a change in expected value that is comparable to a manager's expected excess return, and so is certainly worth considering.
8. We accept that the explanatory power of the model's regression is low for any particular manager. Using the model across many of the investor's managers, however, improves the chance that the investor gains good performance by following the model's results.

REFERENCES

Busse, J., Goyal, A. and Wahal, S. (2010) Performance and persistence in institutional investment management. *The Journal of Finance* 65(2): 765–790.

Goyal, A. and Wahal, S. (2008) The selection and termination of investment management firms by plan sponsors. *The Journal of Finance* 63(4): 1805–1847.

Penfold, R. (2004) *How Well Do Three Risk Models Forecast Risk?* Russell Research Commentary, Russell Investment Group.

intended by Towers Watson to be construed as the provision of investment, legal, accounting, tax or other professional advice or recommendations of any kind, or to form the basis of any decision to do or to refrain from doing anything. As such, this document should not be relied upon for investment or other financial decisions and no such decisions should be taken on the basis of its contents without seeking specific advice.

This document is based on information available to Towers Watson at the date of issue, and takes no account of subsequent developments after that date. In addition, past performance is not necessarily indicative of future results. In producing this document Towers Watson has relied upon the accuracy and completeness of certain data and information obtained from third parties.

This document may not be reproduced or distributed to any other party, whether in whole or in part, without Towers Watson's prior written permission, except as may be required by law. In the absence of its express written permission to the contrary, Towers Watson and its affiliates and their respective directors, officers and employees accept no responsibility and will not be liable for any consequences howsoever arising from any use of or reliance on the contents of this document including any opinions expressed herein.