Perspectives

Evaluating an investment manager

in an uncertain world

TOWERS WATSON

Evaluating an investment manager in an uncertain world

When an investment board hires a manager, it has beliefs about the future performance that the manager will generate. Typically, the manager is given three years to do its job. The investor then reviews it in light of the performance that it generated. In this technical paper, we formulate a method that enables an investor to perform such a review. In doing so, we reveal some profound issues with many of the prevailing conventions of performance evaluation.

Setting the scene

This paper reflects upon a familiar scenario. Consider three parties: an investment board ('the investor'), a fund manager and an investment consultant. In this simplified world, assume that there are only two types of fund manager: stars and flops[1]. Managers within each of these groups have the same expected excess return, with stars expected to outperform flops[2].

Let's also assume that this investor holds US equities and wants them to outperform the S&P 500 Index. As the investor feels unable to manage this portfolio directly, it agrees to hire an external provider of fund management services. To find the fund manager that is most appropriate for its needs, the investor has to decide amongst the available options. With about 4,000 such options available, any decent research effort would be costly and time consuming. The investor therefore agrees to delegate this task and hire an investment consultant, in order to access its research efforts in this area.

By paying the consultant's fee, the investor reveals a belief that its chance of finding a star fund manager via the consultant materially exceeds what it would have been otherwise[3]. Yet this chance is unlikely to be 100%. After all, the investor cannot be certain that the consultant (who we will call Stephen) will find it a star manager.

This lack of certainty complicates the investor's evaluation of its fund manager. After all, if the investor had total confidence that Stephen would find it a star manager then, all other things being equal, any subsequent underperformance from that manager would only arise from bad luck, and so should not require any action from the investor.

In the real world, however, investors often react to experienced underperformance by firing their fund manager. To some consultants and commentators, these actions are evidence of investor irrationality. They talk of 'chasing returns' and claim that much value is lost.

We take a different view. As we discuss in the rest of this paper, these actions show that investors are rational in their response to the underperformance that they experienced. Doing so requires us to respond like a Bayesian statistician. The main difference between the two camps then becomes about the choice of parameters.

What is Bayesian statistics?

Of course, this paper begs the question: what is Bayesian statistics? As The Economist puts it:

"The essence of the Bayesian approach is to provide a mathematical rule explaining how you should change your existing beliefs in the light of new evidence. In other words, it allows scientists to combine new data with their existing knowledge or expertise.

… Imagine that a precocious newborn observes his first sunset, and wonders whether the sun will rise again or not. He assigns equal prior probabilities to both possible outcomes, and represents this by placing one white and one black marble into a bag. The following day, when the sun rises, the child places another white marble in the bag. The probability that a marble plucked randomly from the bag will be white (that is, the child's degree of belief in future sunrises) has thus gone from a half to two-thirds. After sunrise the next day, the child adds another white marble, and the probability (and thus the degree of belief) goes from two-thirds to three-quarters. And so on. Gradually, the initial belief that the sun is just as likely as not to rise each morning is modified to become a near-certainty that the sun will always rise."

In similar vein to the Economist example above, imagine that the consultant, Stephen, picks a ball from two buckets that lie behind a curtain. (We are using the balls in this example to represent fund managers.) In putting his hand between a gap in the curtains, Stephen does not know the bucket from which he is picking. It just so happens that one bucket (with a width of seven units) contains stars, whilst the other (with a width of three units) contains flops. In that way, when Stephen picks the ball from the bucket, the investor considers him to have a 70% chance of picking a star.

The investor then immediately appoints the manager representing the chosen ball[4]. At this time, the investor expects that both manager types have the same tracking error (and so have the same volatility of excess returns relative to the chosen benchmark). It also expects that stars and flops have equal and opposite excess returns. (We make this final assumption to simplify our example, before calibrating it further in the coming pages. When we do, we find that equal and opposite expected excess returns are reasonable in this specific case.)

Now, fast-forward three years, to a time when the chosen manager is underperforming. Using this new evidence, the investor can use a Bayesian approach to update its belief about whether Stephen picked a star manager from behind the curtain. This answer will then update the investor's view about the manager's expected return, and guide it on whether it should retain the manager.

In the following pages, we provide a mathematical formulation of this Bayesian approach. Doing so provides some clear insights about how investors should evaluate fund managers. To make this mathematical model more realistic, we then extend it in a few simple ways. We end with advice on how to use – and not misuse – the model.

A formulation of this familiar scenario

To consider this model, we introduce the following notation:

- t is the time elapsed, in years, between the manager being appointed and the manager's performance being evaluated.
- p_t is the investor's view, at time t, of the probability that Stephen picked a star.
- p_0 therefore represents the investor's initial view of this probability.
- $E(\alpha_{star})$ and $E(\alpha_{flop})$, respectively, represent the investor's unconditional expected excess return (or outperformance) of the star and flop upon appointment. The tracking errors of the respective managers are expressed as ω_{star} and ω_{flop}. We assume that these excess returns are net of fees and independently and identically distributed lognormally. We also assume that these unconditional values remain the same when the investor evaluates the manager in the future.
- r is the annualised excess return experienced by the investor's fund manager in the t years since its appointment.

We then calculate the cumulative t-year expected excess return and tracking error of the star and flop manager. We do so by using the following standard approach for moving from one-year arithmetic returns to multi-year geometric returns. To make the notation clearer, we also remove the subscript for the star and flop manager.

$$\boldsymbol{E}(\alpha_t) = tA$$

$\omega_t^2 = tB^2$, where

$$A = \ln\{1 + \boldsymbol{E}(\alpha)\} - \frac{B^2}{2}$$

$$B^2 = \ln\left\{1 + \frac{\omega^2}{[1 + \boldsymbol{E}(\alpha)]^2}\right\}$$

The cumulative value of the manager's experienced return, r, over t years is:

$$r_t = t\ln(1 + r)$$

We can then calculate the investor's initial expectation for its excess return over t years, $E_0(\alpha_{i,t})$. In doing so, we consider the similar expectations for the star, s, and flop, f.

$$\boldsymbol{E}_0(\alpha_{i,t}) = p_0\boldsymbol{E}(\alpha_{s,t}) + (1 - p_0)\boldsymbol{E}(\alpha_{f,t}) \quad (1)$$

Next, we assume that the manager is a star and calculate the likelihood that its cumulative experienced return over the t years would lie close to r_t. (By close, we mean that this return lies within a certain distance, ε.) We denote this likelihood as $L_{s,t}$. We then repeat the corresponding calculation on the assumption that the manager is a flop and get $L_{f,t}$. In this context, Φ represents the standard normal cumulative distribution.

$$L_{s,t} = \boldsymbol{\Phi}\left\{\frac{r_t + \varepsilon - \boldsymbol{E}(\alpha_{s,t})}{\omega_{s,t}}\right\} - \boldsymbol{\Phi}\left\{\frac{r_t - \varepsilon - \boldsymbol{E}(\alpha_{s,t})}{\omega_{s,t}}\right\}$$

Using Bayes theorem[5], we can now update our probability that Stephen picked a star from behind the curtain. (That is, we can update p_0 to become p_t.)

$$p_t = \frac{L_{s,t}p_0}{L_{s,t}p_0 + L_{f,t}(1 - p_0)} \quad (2)$$

Substituting p_t (rather than of p_0) into equation (1), we obtain the investor's revised expectation for its excess return, $E(\alpha_{i,t})$:

$$\boldsymbol{E}(\alpha_{i,t}) = p_t\boldsymbol{E}(\alpha_{s,t}) + (1 - p_t)\boldsymbol{E}(\alpha_{f,t}) \quad (3)$$

Making sensible assumptions for the expected returns of stars and flops

In the simple example that follows, the client expects star and flop managers to generate benchmark-relative performance of 1% and -1% a year, respectively. Whilst these values do not seem outlandish, it is important that they accord with practice and theory.

In particular, we need to ensure that:

- The expected outperformance (or information ratio) for the skilled manager is reasonable. Much empirical work has been done on establishing reasonable values in this context.
- Given Sharpe, the total outperformance before costs of all managers should be zero.

Following this latter point, we calibrate the expected outperformance of stars and flops as follows:

$$\pi_s\{\boldsymbol{E}(\alpha_s) + c\} + (1 - \pi_s)\{\boldsymbol{E}(\alpha_f) + c\} = 0$$

This calibration leads us to define:

$$\boldsymbol{E}(\alpha_f) = -\left[\frac{c + \pi_s \boldsymbol{E}(\alpha_s)}{(1 - \pi_s)}\right]$$

In the simple example above, we use the following parameters:

- We use the data from French on the cost of active investing. Specifically, we assume a value for c of 0.67% a year.
- Re-casting the data from Barras et al into a model with only two types of investor, we assume a value for π_s of 15%.

We began with an expected outperformance for a star of 1% a year, net of costs. On that basis, and using the calibration and parameters above, we arrive at an expected net outperformance for a flop of -1% a year.

A simple example

The best way to illustrate this model is to consider a simple example. We imagine that:

- The client hires the manager, after following the recommendation of its consultant. At that moment, the client believes that the consultant has a 70% chance of picking a skilled manager and a 30% chance of picking a flop manager[6].
- The client also expects stars to outperform by 1% a year, with a tracking error of 3% a year. It expects flops to have the same tracking error but to underperform by -1% a year.
- It also restates these expectations as compound returns over three years. The client expects stars and flops to generate 2.9% and -3.2% of benchmark-relative performance, respectively. It expects the tracking errors of both managers to be 5.2% over three years.
- The client gives the manager three years to generate its outperformance. Once this time has passed, the client will review the manager's appointment.
- Three years then pass. In this time, the manager underperforms its benchmark by -1% a year (or around -3.0% compounded).

This performance is obviously disappointing and puts the appointment of the manager (and perhaps the consultant) into question. However, to respond appropriately, the client needs a way to revise its former belief about the skill of the manager (much in the way that the newborn in our earlier example needed to update his belief that the sun will rise tomorrow). Clearly, there is now a smaller chance than first thought that the consultant originally picked a skilled manager, as the manager's performance has fallen below expectations since its appointment three years ago. The question is: how much less than that original 70% chance?

We can help to answer this question by calculating the chance that a star manager would have generated the manager's actual performance. To do so, we actually calculate the chance that a star manager got a return close to -3.0% over three years (that is, within 0.1% of -3.0%)[7]. We then do the same on the assumption that the manager was actually a flop. When we do so, we learn that these respective probabilities are 0.8% and 1.5%.

To some extent, these results tell us what we knew before – that the manager's past performance is more in keeping with a flop than a star. Yet the results also tell us the **level** to which these outcomes are more likely. With this information, we can use equation (2) from before to update the client's probability that the consultant originally picked a star manager.

$$p(star) = \frac{(70\% * 0.8\%)}{(70\% * 0.8\%) + (30\% * 1.5\%)} = 0.55$$

In other words, the client now has a 55% belief, not a 70% belief, that it hired a skilled manager. As such, it now expects the manager to outperform by:

$$\text{E (manager)} = (0.55 * 2.9\%) + ((1\text{-}0.55) * \text{-}3.2\%) = 0.2\%$$

As this expected outperformance is positive, and transition costs are material, the client decides to retain the manager.

In **Figure 01**, we show how this expectation varies when using different values for two key factors. First, the investor's initial belief that Stephen picked a star (the columns). Second, the excess return generated by this manager since its appointment three years ago (the rows).

We then colour the resulting expectations according to their sign and magnitude. Negative outcomes are in red, whilst positive outcomes are in teal. Outcomes of greater magnitude are displayed in darker colours than outcomes of lesser magnitude. In the absence of transition costs, one could argue that the investor should terminate its relationship with the manager if the expected excess return becomes negative (and thus red)[8].

For reference, we also circle the outcome from the example on the previous page.

One trait of this model is that, barring two extreme beliefs, investors will give managers some latitude to underperform (and the amount of this latitude depends upon their initial probability that Stephen picked a star). The patience of all investors will eventually be exhausted, though, implying that investors should fire all managers that perform sufficiently poorly.

Extreme beliefs lead to different outcomes, though. It seems intuitive that, if one has absolute faith in Stephen's ability to pick a star, then any level of underperformance from the manager in the following three years will not shake your absolute faith. (The same holds for absolute faith in Stephen's inability to pick a star.)

Figure 01. The investor's expectation about its cumulative excess return in the next three years, given different values of two key factors

Rows – the excess return generated by the manager, since its appointment three years ago
Columns – the investor's initial probability that the consultant picked a star

	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
-6% pa	-3.2%	-3.1%	-3.1%	-3.1%	-3.1%	-3.1%	-3.0%	-3.0%	-2.9%	-2.5%	2.9%
-5% pa	-3.2%	-3.1%	-3.1%	-3.1%	-3.0%	-3.0%	-2.9%	-2.8%	-2.5%	-1.9%	2.9%
-4% pa	-3.2%	-3.1%	-3.1%	-3.0%	-2.9%	-2.8%	-2.6%	-2.4%	-2.0%	-1.0%	2.9%
-3% pa	-3.2%	-3.1%	-3.0%	-2.8%	-2.7%	-2.5%	-2.2%	-1.8%	-1.1%	0.1%	2.9%
-2% pa	-3.2%	-3.0%	-2.8%	-2.5%	-2.3%	-1.9%	-1.4%	-0.9%	-0.1%	1.1%	2.9%
-1% pa	-3.2%	-2.8%	-2.4%	-2.0%	-1.6%	-1.1%	-0.5%	0.2%	0.9%	1.8%	2.9%
0% pa	-3.2%	-2.5%	-1.9%	-1.3%	-0.7%	-0.1%	0.5%	1.1%	1.7%	2.3%	2.9%
1% pa	-3.2%	-2.1%	-1.1%	-0.4%	0.3%	0.9%	1.4%	1.8%	2.2%	2.5%	2.9%
2% pa	-3.2%	-1.4%	-0.2%	0.6%	1.2%	1.6%	2.0%	2.2%	2.5%	2.7%	2.9%
3% pa	-3.2%	-0.5%	0.7%	1.4%	1.8%	2.1%	2.3%	2.5%	2.6%	2.8%	2.9%
4% pa	-3.2%	0.4%	1.4%	1.9%	2.2%	2.4%	2.6%	2.7%	2.7%	2.8%	2.9%
5% pa	-3.2%	1.2%	2.0%	2.3%	2.5%	2.6%	2.7%	2.7%	2.8%	2.8%	2.9%
6% pa	-3.2%	1.8%	2.3%	2.5%	2.6%	2.7%	2.8%	2.8%	2.8%	2.8%	2.9%

Reflecting the potential for a manager's excess returns to mean-revert

As sensible as the earlier model may seem from a Bayesian perspective, following it can run counter to received wisdom. Studies show that investors tend to lose value when they hire and fire managers for performance reasons[9]. One popular reason for this loss is that investors tend to fire managers that have underperformed, only to see the performance from these managers rebound in the following years.

In order to resolve this seeming contradiction, we need to incorporate this behaviour and reflect the tendency of the manager's excess returns to mean-revert over time. We assign a parameter (λ) for this tendency and call it the one-period force of trending to the unconditional mean, $\overline{E}(\alpha_i)$. Formally, but simplifying the notation, we use λ to update our expected excess return for the manager in the next period[10]:

$$E_t(\alpha_i) = \lambda r_{t-i} + (1-\lambda)\overline{E}(\alpha_i)$$

Akin to a correlation statistic:

- A **positive** λ implies that the manager's excess returns will typically trend. That is, we expect good performance in the last period (a good r_{t-1}) to lead to above-average performance this period (and vice-versa).
- A **negative** λ reflects mean-reversion. Bad performance from last period will therefore tend to precede above-average performance this period (and vice-versa).
- A **zero** λ implies no expected relationship between the returns of both periods.

Revised results

We now revise our earlier example to reflect the tendency of the manager's excess returns to mean-revert. We show the impact of changing this λ parameter from 0 to -0.2 across consecutive triennium[11]. In doing so, we still arrive at the previous example's conclusion that the client now has a 55% belief, not a 70% belief, that it hired a skilled manager.

However, the manager's expected excess return now differs from before. Specifically:

- The expected outperformance of a star is now 4.0%, not 2.9% (as it performed worse than expected in the last period and so should expect better returns than normal for this period).
- The expected underperformance of a flop remains at -3.2% (as it performed in line with its unconditional expectation in the previous three years and so would not be expected to revert from its unconditional expectation in the coming three years).

By calculating the weighted average of these two returns, the expected outperformance of the manager is now 0.8%, not 0.2%. This improvement in expected return comes from the manager's poor performance last period and the client's assumption that excess returns across consecutive periods tend to mean-revert[12].

In order to be consistent with our analysis without mean-reversion, we chart the outcomes with slight mean-reversion in a similar style to those in **Figure 01**. These outcomes are shown below in **Figure 02**.

Figure 02. The investor's expectation about its cumulative excess return in the next three years, given different values of two key factors (with mean-reversion)

Rows – the excess return generated by the manager, since its appointment three years ago
Columns – the investor's initial probability that the consultant picked a star

	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
-6% pa	-0.1%	-0.1%	0.0%	0.0%	0.0%	0.0%	0.1%	0.1%	0.3%	0.7%	7.1%
-5% pa	-0.7%	-0.7%	-0.7%	-0.6%	-0.6%	-0.5%	-0.4%	-0.3%	0.0%	0.8%	6.5%
-4% pa	-1.3%	-1.3%	-1.2%	-1.1%	-1.0%	-0.9%	-0.7%	-0.4%	0.1%	1.2%	5.9%
-3% pa	-2.0%	-1.9%	-1.7%	-1.6%	-1.4%	-1.1%	-0.8%	-0.3%	0.5%	1.9%	5.3%
-2% pa	-2.6%	-2.4%	-2.1%	-1.8%	-1.5%	-1.1%	-0.5%	0.2%	1.1%	2.5%	4.6%
-1% pa	-3.2%	-2.8%	-2.3%	-1.8%	-1.3%	-0.7%	0.0%	0.8%	1.7%	2.8%	4.0%
0% pa	-3.8%	-3.0%	-2.3%	-1.6%	-0.8%	-0.1%	0.6%	1.3%	2.0%	2.7%	3.4%
1% pa	-4.4%	-3.1%	-2.0%	-1.0%	-0.2%	0.4%	1.0%	1.6%	2.0%	2.5%	2.8%
2% pa	-5.0%	-2.8%	-1.4%	-0.5%	0.2%	0.7%	1.2%	1.5%	1.8%	2.0%	2.2%
3% pa	-5.6%	-2.4%	-0.9%	-0.1%	0.4%	0.8%	1.0%	1.2%	1.4%	1.5%	1.6%
4% pa	-6.1%	-1.9%	-0.6%	0.0%	0.3%	0.6%	0.7%	0.8%	0.9%	1.0%	1.1%
5% pa	-6.7%	-1.5%	-0.6%	-0.2%	0.1%	0.2%	0.3%	0.4%	0.4%	0.5%	0.5%
6% pa	-7.3%	-1.4%	-0.7%	-0.5%	-0.3%	-0.2%	-0.2%	-0.1%	-0.1%	-0.1%	-0.1%

As we can see from the difference between **Figure 01** and **Figure 02**, the assumption of mean-reversion has a large impact on these results. Before, there was only one effect at play: the worse a manager performed, the lower the client's confidence that it was a star, which reduced its future expected return. We now have another effect, which runs counter to the first effect: the worse a manager performed, the better the expected performance of the star (and often the flop). As extreme performance amplifies any reversion to the mean, we find the greatest difference between these two charts at their extremes. For instance:

- **In the last three years, the manager underperformed by -6% a year**. In this scenario, the manager strongly underperformed what was expected of both the star and the flop. Whilst the client now has almost no confidence that the manager is a star, the performance of the flop is expected to rebound to such an extent that it will be positive for the next three years. (The performance of the star will be even higher.) For that reason, and absent any other factors, the client should retain the manager. In fact, the phrase 'a rising tide floats all boats' resonates with the circumstances of this instance.
- **In the last three years, the manager outperformed by 6% a year**. The forces in the previous instance also apply here, but in the opposite direction. The client should therefore 'take profits' from the manager after such exceptional outperformance.

For reference, we also include **Figure 03**, which shows how the client's revised return expectation for the manager varies when using different values for two key factors. First, the investor's initial belief that Stephen picked a star (the columns). Second, the mean-reversion of the manager's excess returns (the rows). In doing so, we assume that the manager underperformed its benchmark by -1% a year in the three years between selection and evaluation.

The two values circled in **Figure 03** are those circled in **Figures 01** and **02**.

Figure 03. The investor's expectation about its cumulative excess return in the next three years, given different values of two key factors (and underperformance of -1% a year)

Rows – the mean-reversion of the manager's excess returns
Columns – the investor's initial probability that the consultant picked a star

	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
-0.5	-3.2%	-2.7%	-2.2%	-1.5%	-0.9%	-0.1%	0.8%	1.8%	2.9%	4.2%	5.8%
-0.4	-3.2%	-2.7%	-2.2%	-1.6%	-1.0%	-0.3%	0.5%	1.5%	2.5%	3.7%	5.2%
-0.3	-3.2%	-2.8%	-2.3%	-1.7%	-1.1%	-0.5%	0.3%	1.1%	2.1%	3.3%	4.6%
-0.2	-3.2%	-2.8%	-2.3%	-1.8%	-1.3%	-0.7%	0.0%	0.8%	1.7%	2.8%	4.0%
-0.1	-3.2%	-2.8%	-2.4%	-1.9%	-1.4%	-0.9%	-0.2%	0.5%	1.3%	2.3%	3.4%
0.0	-3.2%	-2.8%	-2.4%	-2.0%	-1.6%	-1.1%	-0.5%	0.2%	0.9%	1.8%	2.9%
0.1	-3.1%	-2.8%	-2.5%	-2.1%	-1.7%	-1.3%	-0.7%	-0.1%	0.5%	1.3%	2.3%
0.2	-3.1%	-2.9%	-2.6%	-2.2%	-1.9%	-1.5%	-1.0%	-0.5%	0.1%	0.8%	1.7%
0.3	-3.1%	-2.9%	-2.6%	-2.3%	-2.0%	-1.7%	-1.2%	-0.8%	-0.3%	0.4%	1.1%
0.4	-3.1%	-2.9%	-2.7%	-2.4%	-2.2%	-1.8%	-1.5%	-1.1%	-0.6%	-0.1%	0.5%
0.5	-3.1%	-2.9%	-2.7%	-2.5%	-2.3%	-2.0%	-1.8%	-1.4%	-1.0%	-0.6%	-0.1%

Conclusion

In our earlier examples, we modelled how investors could best alter their expectations for a manager's future performance in light of its past performance. By doing so, we hope to give investors a clearer framework with which to evaluate their managers.

Of course, we are not suggesting that investors can get perfect clarity on whether to retain a manager based on just one model that considers two factors. Rather, users of models should be cautious, as models simplify the real world in order to aid our understanding. Yet this simplification can lead models to miss the complexity of real life. Investors should therefore use this model diligently and alongside existing qualitative ways of evaluating managers.

However, the model's use goes further than just providing a manager evaluation framework. The thinking behind it also draws some of our industry's misunderstandings about manager evaluation into sharper focus. We now consider two of these misunderstandings.

Do not always expect the same answer if you are asking two different questions

Investors often ask their consultants two different questions about the performance expected from their manager, and are surprised when they get two different answers. First, they want to know the outperformance that a skilful manager can attain. Second, they ask for the expected outperformance of this manager's portfolio. The first question is contingent on the fact that the manager is skilled. Its answer therefore exceeds the answer to question 2, which also needs to reflect the chance that the manager is a flop. This question may seem arcane, but we often hear talk of consultants or managers "over-promising but under-delivering". Some of this may be happening, but another explanation is also plausible. That is, managers say what they can achieve, on the assumption that they will succeed, and the investor assumes this value to be its expected outperformance. In doing so, Equation (1) tells us that the investor is implying that the investment consultant is perfect at selecting managers[13].

The importance of the mean-reversion assumption for the manager's excess returns

For the sake of brevity, this paper has not considered in detail the sensitivity of the model's results to this assumption. As the difference between **Figure 01** and **Figure 02** attests, however, the model's results are highly sensitive to a change in this assumption. To us, this sensitivity does not limit the model's use. Rather, it reveals the importance of the level of mean-reversion in decision-making. It also shows that investors should take time to consider how much the excess returns of their manager will likely mean-revert. Investors often do this already, albeit implicitly, saying "we think the manager will have a rebound" or "there is some intrinsic value in its portfolio that has yet to be realised." Instead of making these qualitative statements, we encourage investors to consider quantifying this level of mean-reversion (or trending) in the manager's excess returns. They can then test the robustness of their manager retention decision by considering the sensitivity of their conclusions to small changes in the value of the manager's mean-reversion of excess returns.

We also feel that a lack of appreciation for mean-reversion generates a misunderstanding in performance evaluation. Consultants often feel frustrated by investors' decisions to react to performance. As we have shown, though, this is a reasonable Bayesian response from the investors. However, the consultants tend to see something that the clients miss: mean-reversion in manager excess returns. If investors also incorporate this factor appropriately into decision-making, then they will be less likely to lose value when replacing managers, but can still respond as they (and the Bayesians) see fit.

Further reading

Barras, Lauren, O Scaillet, and R Wermers. *False discoveries in mutual fund performance: measuring luck in estimated alphas.* The Journal of Finance, February 2010.

Busse, Jeffrey, Amit Goyal, and Sunil Wahal. *Performance and Persistence in Institutional Investment Management.* The Journal of Finance LXV, no. 2 (April 2010).

French, Kenneth R. *The cost of active investing.* The Journal of Finance LXIII, no. 4 (August 2008): 1537-1573.

Goyal, Amit, and Sunil Wahal. *The Selection and Termination of Investment Management Firms by Plan Sponsors.* The Journal of Finance LXIII, no. 4 (August 2008): 1805-1847.

Penfold, Robin. *How much value will you likely gain or lose by replacing your manager for performance reasons*? Towers Watson Limited, 2011.

Sharpe, William. *The Arithmetic of Active Management*. Financial Analysts Journal, January/ February 1991.

The Economist. *In praise of Bayes.* 2000, 30 September.

Further information

For further information, please contact your Towers Watson consultant or:

Robin Penfold
+44 1737 274462
robin.penfold@towerswatson.com

Endnotes

1 One area for further investigation is whether the presence of more than two types of manager alters the model's results. For that reason, we plan to build a model with three types of manager: skilled, mediocre and bad. This approach also has intuitive appeal, as it mirrors that used by Barras et al in their excellent study of manager performance.

2 Of course, this approach can apply just as much to portfolios of managers. In fact, by comparing the results of the model when used on the portfolio and its underlying managers, we can begin to understand the cost (or value) of the mental accounting by most investors in this area. (That is, where they ignore the portfolio effect and fire the worst-performing manager, even though there will likely always be a poor-performing manager in any large portfolio of managers.)

3 This statement must be true – if the investor does not hold this view, then it should not have hired the consultant and paid it a fee. This conclusion follows directly from the assumption that there is no difference in expected return amongst the stars (or amongst the flops). As an aside, we conjecture that this conclusion should help to reveal the appropriate size of the consultant's fee for manager research, given the investor's perceived level of marginal benefit.

4 In doing so, we assume that the investor does not see the excess returns of the manager before its appointment. In reality, as Goyal and Wahal show, many investors tend to screen managers based upon their past performance, often rejecting those that have underperformed. We have built a more complex model that reflects this realism, although do not report its results here for the sake of brevity and simplicity.

5 Strictly speaking, we are working backwards from posterior to prior.

6 One way that we can use to calibrate this parameter is to consider the relative frequency of outperformance for the managers that we selected in our model portfolios.

7 Our results are not sensitive to sensible changes in the width of this return range.

8 One implication of this statement is that the proceeds of any disinvestment from the manager are invested passively, until other actively managed opportunities arise.

9 See Goyal and Wahal, and Penfold for more details.

10 We note that the tracking error of the manager will not change, as this is a one-period autoregression.

11 We assume a slight level of mean-reversion in our example. 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 (Busse, Goyal and Wahal), the authors find that better manager excess returns from the previous year tend to lead to worse excess returns in the following three years. One possible reason for such a dynamic is that outperforming managers attract assets, which hampers future outperformance.

12 One area of ongoing research is what happens to the following period's likelihood that the consultant picked a star manager, once the manager has underperformed significantly. To do this, we are using a two-stage Monte-Carlo analysis, and paying particular attention to the impact on false negatives.

13 We also find that consultants often apply a rule of thumb (such as halving the manager's view of its expected outperformance) to arrive at likely portfolio outperformance. Rather than doing this, it would be far better for the consultant to use equation (1).

I would like to thank many colleagues for their help with this paper, particularly Alasdair Macdonald (with whom I had the initial discussion on this subject), Zacky Choo, Martin Dietz and Stephen Miles.

About Towers Watson

Towers Watson is a leading global professional services company that helps organisations improve performance through effective people, risk and financial management. With 14,000 associates around the world, we offer solutions in the areas of employee benefits, talent management, rewards, and risk and capital management.

towerswatson.com

TOWERS WATSON