
Paying (Only) for Skill (Alpha)— A Practical Approach

Donald M. Raymond, CFA
Senior Vice President
Canada Pension Plan Investment Board
Toronto, Ontario

Agency problems can cause asset managers to act in a way that is not necessarily in the best interests of their clients. To align the interests of asset managers with their clients, what is needed is a straightforward fee structure (and associated terms) that seeks to distinguish skill from luck and that has a continuous incentive structure.

Founded in 1997, the Canada Pension Plan Investment Board is the investment arm of Canada's national social security system, known as the Canada Pension Plan (CPP). All working Canadians contribute to the plan, and only those who contribute draw a pension from it. The board operates independently of the government and has a mandate to maximize return without undue risk.

Historically, the assets under management were passively managed externally in globally developed market equities. But in 2003, the board began to focus on building active capabilities. As a large and growing pension fund about to embark on hiring its first set of external active managers, it had the good fortune to be able to draw on investment professionals with experience on both the sell side and the buy side of the industry. As a result, the board decided to develop its own fee structure for paying managers. This presentation is the story of that fee structure.

Steven Landsburg (1993) talks about the power of incentives. One of my favorite quotes is "Most of economics can be summarized in four words: 'People respond to incentives.' The rest is commentary" (p. 3). He also describes research conducted at Texas A&M University in which rats and pigeons could "purchase" food and drink by pushing levers. Each item had a "price" equal to a certain number of pushes, and the animals had an "income" of so many total pushes a day.

This presentation comes from the Asset and Risk Allocation: Alpha and Beta Investment Strategies conference held in Philadelphia on 13–14 November 2007.

The researchers found that rats and pigeons responded differently to changes in price, changes in income, and changes in wage rates. When the price of root beer, for example, went up, they bought less root beer. When wage rates went up, they worked harder unless their income was already high, in which case they chose to enjoy more leisure.

The responses observed by the researchers are the same responses economists would expect to observe in humans. I am not suggesting that active managers are anything close to being pigeons or rats, but I am suggesting that when the client signs an investment management agreement, it is conveying a certain set of incentives to the external manager. Consequently, the client should be clear and deliberate about what the incentives are, which is effectively what the board set out to do with its fee structure.

But before I explain the fee structure, I will talk about some foundational elements of fee structures. First, I will discuss manager incentives and how, in general, they relate to two common fee structures found in the industry today—the standard fixed-fee and performance-fee structures. Second, I will consider manager skill and, in particular, how at the Canada Pension Plan Investment Board we quantify manager skill and distinguish skill from luck.

Manager Incentives

Typical fee structures provide different manager incentives. The first common fee structure is the fixed-fee structure. It pays managers regardless of their skill and whether they add value or not. This

structure obviously creates an incentive to gather assets. One could argue that there is also an incentive for performance in order for the manager to retain assets over longer periods of time, but I would argue that is a fairly weak incentive compared with what the manager can get with a performance-fee structure, particularly a well-designed performance-fee structure.

Under the standard performance-fee structure, a manager is paid a base fee plus a performance fee subject to a high-water-mark provision. Often, the base fee is a large proportion of the total expected fee. In my experience, this structure tends to be common in fixed income because these strategies tend to run at a lower tracking error with a very high information ratio assumption. In a typical performance-fee structure, managers are paid for skill and luck as long as they add value. So, there is no emphasis on quality of returns; as long as managers add value, they get paid.

Moral hazard issues also exist with this fee structure. A moral hazard can be said to exist when investors or property owners are in some sense protected from the downside risks of poor investment decisions or outcomes, as is the case for active managers who have no downside in their performance fee. Another moral hazard issue arises when managers maximize the performance-fee “option” by taking higher active risk, which requires that they be monitored on both an *ex post* and an *ex ante* basis. Finally, managers have an incentive to gather assets beyond the capacity of the strategy, which is even more true in the case where the base fee is a

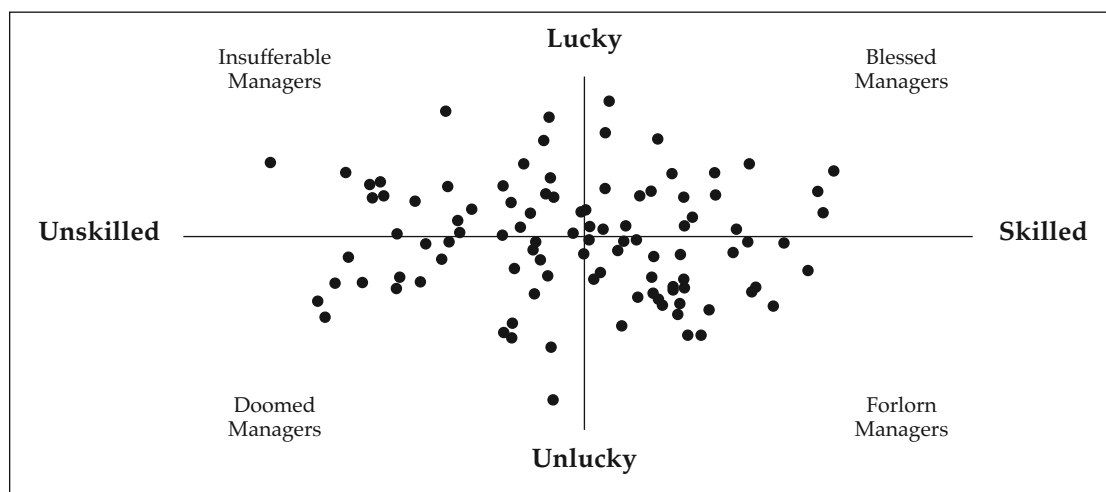
significant portion of the total fee because it is essentially the same situation as the fixed-fee structure.

Manager Skill

Paying managers for skill requires a way to measure skill. **Figure 1** is based on a diagram from Grinold and Kahn (1999). It is a stylized diagram that breaks a manager’s realized return into two components—skill and luck—and plots it along the two dimensions. Each dot represents a manager. The skill component is plotted along the horizontal axis; on the right side are the skilled managers, and on the left side are the unskilled managers. The component resulting from luck is plotted along the vertical axis; managers above the horizontal line are the lucky managers, and the ones on the bottom are the unlucky managers.

This diagram divides the manager universe into four quadrants. In the bottom left quadrant are the doomed managers. They are managers without skill who additionally have had a bad draw from the return distribution and are destined to go out of business. The top left quadrant represents the insufferable managers. These are managers without skill who happen to have had good returns, so they have been lucky. In the top right quadrant are the blessed managers. They have skill and have also been lucky, so they are generally very happy. In contrast, the managers on the lower right are forlorn; they have skill but have had a bad draw and have either not added value or added less value than they otherwise should have based on their skill alone.

Figure 1. Four Stylized Manager Types



Source: Based on Grinold and Kahn (1999).

Objectives and Definitions

As we began to think about designing our own fee structure, the main questions were whether we could construct a better negotiation process and whether we could come up with a fee structure that aligned interests. In terms of process, our approach was to define the fee structure up front using a few key design principles—have few parameters, keep it simple, offer continuous incentives, reduce moral hazards, and pay a manager for skill in adding value—and then negotiate the parameters.

Discontinuities in a fee structure can introduce moral hazard, thus making a simple structure imperative. The more complicated a fee structure is, the more likely it is that moral hazard issues will be introduced inadvertently. Paying managers for skill in adding value involves defining skill and introduces a notion of quality to the return stream that is linked to how the managers are paid. Understandably, moral hazard cannot be completely removed, but the goal is to try to minimize it and align the manager's interests with the client's interests. The fee structure the board designed was not an attempt to reduce the amount of fees paid but an attempt to design a fee structure that would best align interests. Therefore, the board included termination events in the contract, tried to impose minimal parameters, and allowed for negotiation of the parameters once the structure of the performance fee was clear.

In designing the fee structure, we had to answer the question of how to define manager skill. We defined the probability (Pr) that a manager has skill as 1 minus the probability that the manager is lucky. If one assumes that active returns are normally distributed, this probability is the cumulative normal distribution, $F(\bullet)$, of the information ratio, IR, at time N (number of years) multiplied by the square root of time. We define Q_N , the quality factor at year N , as the probability that the manager has skill:

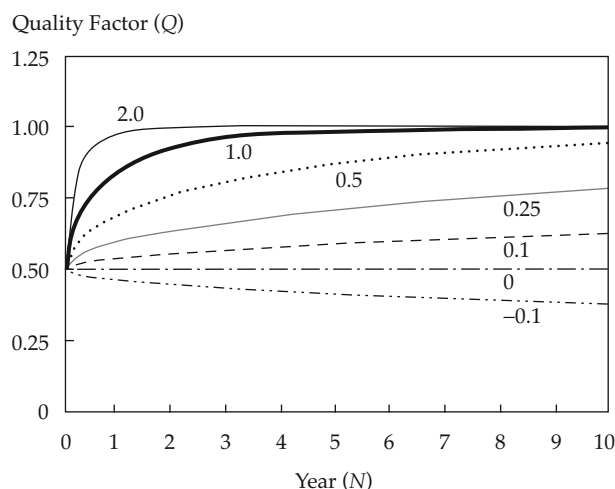
$$\begin{aligned} Pr\{\text{skill}\} &= 1 - Pr\{\text{luck}\} \\ &= 1 - Pr\{\text{realized returns were due to chance}\} \\ &= F(IR_N \times N^{1/2}) \\ &= Q_N, \end{aligned}$$

where Q_N is bounded between zero and one.

Figure 2 is a visual representation of Q , or the quality factor. Each line is a function of time for managers with different IRs. They all start off at time zero with the quality factor equal to 0.5, meaning that without any performance information, the managers start off with a 50 percent probability that they have skill. As they start to generate active returns, the quality factor can be measured, and as shown, for a manager with a high IR (such as 2), the

quality factor goes up to 1 very quickly. After one year, it is at about 0.95, meaning there is a 95 percent chance they have skill, and after two years, it is essentially 1. A manager with an IR of 1 after one year has a quality factor of about 0.8. It takes five or six years before it finally goes up to 1, indicating that the value added is most likely a result of skill. For a low IR, such as 0.1, it takes a long time to get the quality factor up to 1. In fact, even after 10 years, it is approximately 0.62.

Figure 2. Quality Factor over Time for Various IRs



So, Q starts at 0.5 and rises to 1 faster for managers with high IRs. Note also that Q is a continuous function because, as I mentioned earlier, discontinuities create moral hazard.

The next step we faced was to build the Q factor into the fee structure.

Proposed Fee Structure

After setting design objectives and devising a way to define manager skill, we were ready to begin designing the proposed fee structure. In this section, I will look at the two key parts of the fee structure, the performance fee and the base fee, and then I will talk about putting it together and negotiating the parameters.

Performance Fee. The cumulative performance fee, F , that is paid at year N is equal to a participation rate, P , which is a negotiated parameter equal to the proportion of total value added that the manager earns, multiplied by the quality factor (Q) at year N multiplied by X , the cumulative value added to year N (in dollars):

$$F_N = P \times Q_N \times X_N.$$

For a manager with skill, an IR greater than zero, and a quality factor that goes to 1, the upper limit of the cumulative fees earned in year N is a proportion of the total value added:

$$\lim_{N \rightarrow \infty} F_N = P \times X_N.$$

The Q in the first formulation serves to hold back some of the performance fee until the manager demonstrates that the value added is a result of skill. The fee is effectively deferred and gets released as Q goes toward 1 over time. As a simple example, consider the line in Figure 2 for the manager with an IR equal to 1. (Obviously, this is a very simplified example because even a manager with an IR of 1 would have some variability over time.) By Year 6, the Q factor reaches 1. Applying the fee formula shows that by Year 6, the total fees paid out to the manager are a fixed percentage of the total value added, based on the agreed-upon participation rate.

Using the same example, now assume the manager has added \$20 million a year in value and that the negotiated participation rate is 25 percent—meaning over the long term, the manager will be

paid 25 percent of the total value added. In **Figure 3**, the gray bars represent the cumulative performance fee earned, which is the participation rate multiplied by the total value added ($P \times X_N$), and it goes up by \$5 million a year (25 percent of the \$20 million). The white bars are the cumulative performance fee paid (F_N), which is the cumulative performance fee multiplied by Q . By the sixth year, the cumulative fee paid nearly equals the cumulative fee earned. The difference is represented by the small black bars [$F_N - (P \times X_N)$]. A certain amount of the fee earned is held back, which is the result of Q starting at 0.5 and going to 1. Note in Year 1 that the fee earned is \$5 million; the fee paid is around \$4 million, which is $Q = 0.8$ multiplied by the fee earned; and \$1 million is held back.

In **Figure 4**, the annual fees earned and paid are shown.¹ Again, the annual performance fee the manager earned ($P \times x_N$) is \$5 million, the annual performance fee paid is approximately \$4 million

¹In the notation, uppercase letters refer to cumulative quantities and lowercase letters refer to annual quantities.

Figure 3. Cumulative Fees Earned, Paid, and Held Back over Time

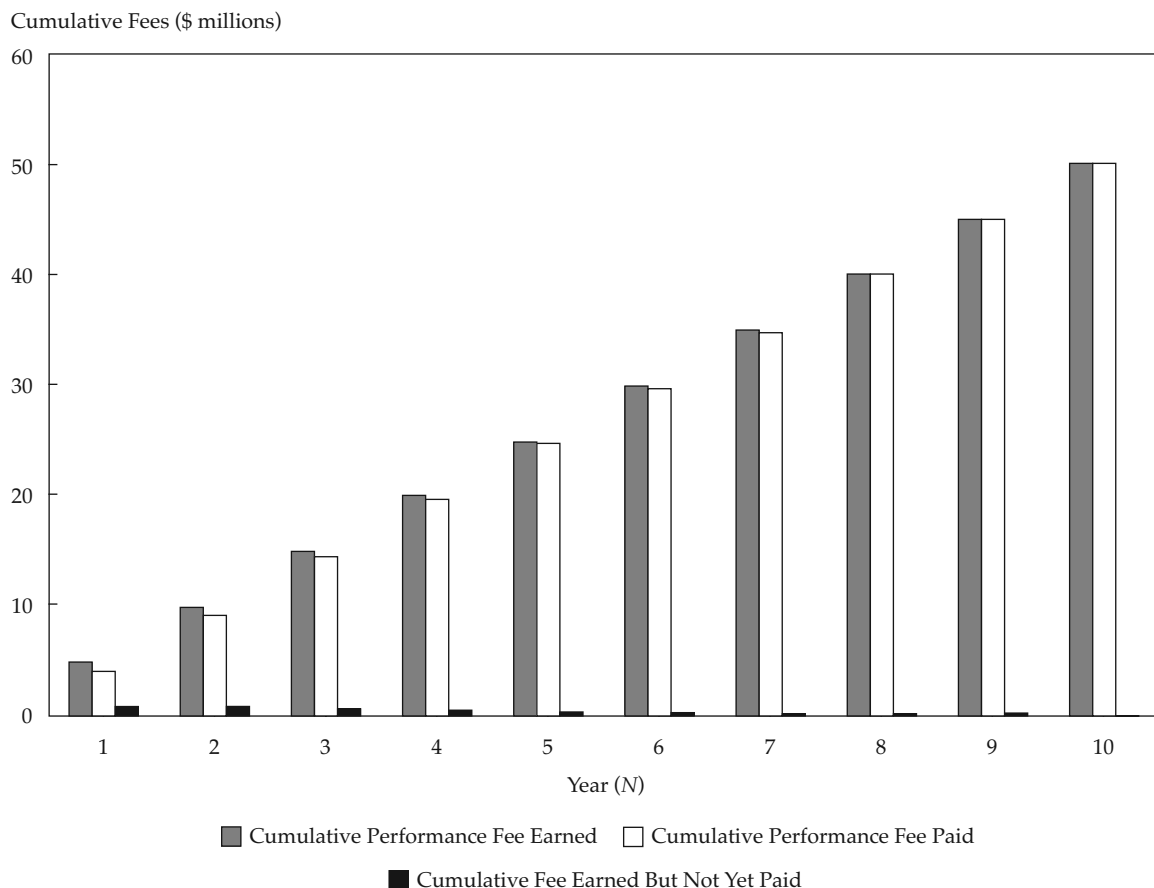
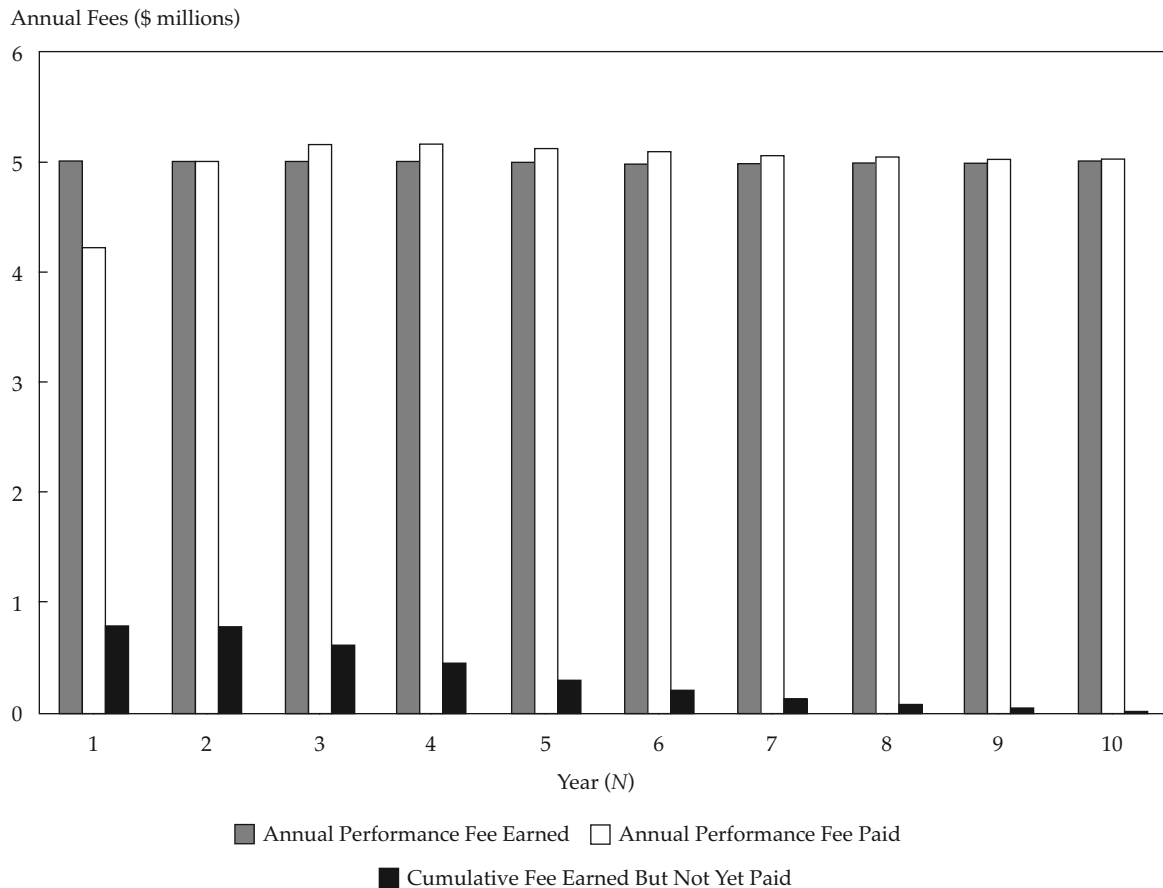


Figure 4. Annual Fees Earned, Paid, and Held Back over Time

(f_N), and \$1 million has been held back because of the quality factor (Q). Over time, the amount held back decreases, and in this example, in Year 3 the fees actually paid (the white bars) go above \$5 million. In the formula, Q serves as a built-in escrow mechanism.

The performance fee is the first and most important component of the fee structure. The intention is that over time the manager will be paid a fixed proportion of the value added.

Base Fee. The next element in the fee structure is that of a minimum annual base fee, or the keeping-the-lights-on component. We started with a short-term base fee, or annual minimum fee (a). The negotiated base-fee rate, B , is the second negotiated parameter in the fee structure and is a percentage of the active risk target (σ). So,

$$a = B \times \sigma.$$

The active risk target is expressed in absolute dollar terms. As an example, consider a \$500 million mandate with 4 percent active risk. It would have an absolute risk target of \$20 million. The reason is that higher tracking error mandates use up more of

a manager's alpha-generating capacity. Structuring the base fee in this manner merely reflects that fact.

Putting It Together. Putting the two components together, the performance fee and the base fee, the actual fee, f , in year N is given by

$$f_N = \max \{ B \times \sigma; (P \times Q_N \times X_N) - F_{N-1} \}.$$

So, the annual fee is the maximum of the two components—the base fee and the performance fee. The base fee is simply the negotiated base rate multiplied by the active risk target, and if that amount exceeds the annual performance fee, the manager earns the base fee. If the performance-fee component is larger, the manager earns the performance fee. The part of the expression to the right of the semicolon expresses the performance fee as the difference between the cumulative total performance fee minus all the fees paid up until the prior year.

Implicit in this fee structure is the idea that the base fee is an advance on future performance fees. The term on the far right, F_{N-1} , is the total cumulative fees paid up until the previous year. That represents all fees paid in prior years, including if the manager received just the base fee. In essence, all

the previous fees paid, including previous base fees, are subtracted when the current year's fee is calculated. Inherent as well is a high-water-mark feature because fees are limited to a proportion of total cumulative value added.

Negotiating Parameters. As part of putting the fee structure together, we gave some thought to the two negotiable parameters—the base rate, B , and the participation rate, P —on the performance side.

Over the long run, the participation rate is the most important parameter because it is the rate that determines what proportion of the total added value goes to the manager relative to the client. The base rate, however, is more interesting because we have to consider when it is equal to the expected performance fee. Put another way, we need to know the breakeven point at which the two sides of the max operator are equal to one another.

In the long term, as the Q factor goes to 1, we can set the two sides of the max operator equal to each other with the base fee on the left side and the performance fee on the right side. Solving for B/P , we see that the ratio of the two negotiable parameters in the fee structure equals the annualized added value divided by the active risk target, which represents a breakeven IR. Following is the full formula:

$$\begin{aligned} B \times \sigma &= (P \times X_N) - F_{N-1} \\ &= (P \times X_N) - P \times X_{N-1} \\ &= P \times (X_N - X_{N-1}) \\ &= P \times x_N \end{aligned}$$

$$B/P = x_N / \sigma = \text{Implied breakeven IR,}$$

where x_N is the value added in year N . B/P can be thought of as an implied breakeven IR. For example, if B is 2.5 percent and P is 25 percent, the implied breakeven IR is 0.1. We would, however, expect the manager to have an IR significantly greater than 0.1; otherwise, the base fee will dominate the performance fee, which can lead to asset-gathering incentives as I mentioned previously.

One additional part of the fee that we thought about was what is fair and appropriate for termination events so as not to create moral hazard issues from either a manager's perspective or from the board's perspective. If the treatment of performance fees in the event of termination is unfair to the manager (e.g., if Q goes to zero on termination), then the manager will want to negotiate the B and the P much higher. So, we decided to define the termination events in advance and then negotiate the parameters. For example, if the manager is fired for breach of contract, Q goes immediately to zero and any accrued performance fees are forfeited. If the manager is let go without cause (e.g., as the result

of a change in strategy), Q would go to 1 and any performance fees being held back would be released to the manager. Finally, if the manager resigns, Q stays at whatever Q was determined to be by virtue of the manager's skill, and the performance fee is paid at that point in time.

Alignment with Design Objectives

As I previously mentioned, we defined some key design objectives in creating the fee structure, and I want to review how the formula meets those objectives.

The first objective was to align the interests of the manager with the client. This is accomplished by tying fees to performance to create alignment. As I just showed, B/P is negotiated in such a way that the base fee does not dominate the performance fee. Otherwise, it would create an asset-gathering incentive for the manager.

Second, the fee structure pays for skill. Of course, clients cannot pay only for skill unless they know exactly what a manager's IR is. We believe the formula does a reasonably good job, at least qualitatively, of defining skill. The quality factor holds back part of the performance fee until a manager demonstrates that the performance is the result of skill and not luck.

Third, the fee structure minimizes moral hazard. We believe it creates a disincentive to exceed capacity limits, which would only be true if every one of a manager's clients were to use this fee structure. In fact, part of the reason for disclosing the fee structure publicly is that it might become more widely adopted. In addition, the fee structure helps to minimize moral hazard because its incentives are continuous.

Finally, the fee structure is simple. It may not seem like it at first, but once you think about it and work with it for a while, it is actually quite intuitive. It has a built-in escrow mechanism, built-in high water marks, and the base fee acts as a prepayment of future performance fees.

Conclusion

When we present this fee structure to managers, most of them are initially puzzled because it looks complicated. They may be a bit skeptical or wonder if we are trying to take advantage of them. But generally, the managers go away, think about it, run simulations, try to map it onto their existing fee structures, and figure out how it really works. The hope is that they become comfortable with the approach and its rationale because the intention is, in fact, noble; it seeks only an alignment of interests.

Overall, the skilled managers like the fee structure. They see the alignment of interests and the simplicity, and they also realize that widespread adoption would be a comparative advantage for them because they believe they have skill. Unfortunately, the forlorn managers are not so thrilled, and

the doomed are distraught because they do not have skill. In our experience, those managers naturally self-select themselves out of searches involving this kind of fee structure.

This article qualifies for 0.5 CE credits.

REFERENCES

Grinold, Richard C., and Ronald N. Kahn. 1999. *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk*. 2nd ed. New York: McGraw-Hill.

Landsburg, Steven E. 1993. *Armchair Economist: Economics and Everyday Life*. New York: The Free Press.

Question and Answer Session

Donald M. Raymond, CFA

Question: For the purpose of the fee structure, why do you exclude historical records when assessing manager skill?

Raymond: There are several reasons. We try to give the managers as unconstrained a mandate as possible, which is invariably going to be different from any historical performance records they may have. There is also the question of composites, comparability, and other issues. But ultimately, the only true sample that matters to us is the one that the manager generates in our account.

Question: What happens if a manager's quality factor (Q) drops or if a manager actually has skill but no longer devotes time to your portfolio?

Raymond: That is definitely a risk. If the manager's Q goes below 0.5, which means he or she has lost money for us, the person overseeing that mandate must decide whether to keep the manager on or not. If a manager stops dedicating the resources to the account that are appropriate, then the obvious solution is to let the manager go.

Question: Have you tried to negotiate clawbacks on incentive fees, and if so, have you been successful?

Raymond: We have not. The reason is that the escrow mechanism built into the fee structure has the effect of holding back the performance fee. So, the fee is not paid out all at once, which means that we don't have to negotiate clawbacks.

Question: Do you apply a similar fee structure to internal teams?

Raymond: We measure the internal teams using the same parameters, and we know what their Q scores are. Our internal compensation system mirrors the fee structure we use with the external managers.

Question: Does this fee structure work for infrastructure, private equity, and real estate?

Raymond: No. One reason is that it is hard to measure the IR. You could possibly do it on private equity, where there is some chance of benchmarking. But in real estate, even the benchmarking questions are challenging.

Question: What flexibility do you allow in negotiating the participation rate (P) and the base fee (B)?

Raymond: There's complete flexibility. Often, the managers will come back with their own proposal; in some cases, they are close to being in an acceptable range, and in some cases, they are way outside of what we consider an acceptable range. We have our own internal simulation tools to map various fee structures into our overall fee structure, and of course, managers do the same. But as I have said, the principle is to define the fee structure up front, include all the termination events and clauses in the agreement, and then negotiate the parameters.

Question: How do you prevent a manager from simply gaming the benchmark that the tracking error is being evaluated against?

Raymond: We have attempted to mitigate that problem by implementing most of our mandates as pure overlays. In fact, the first four managers we hired were long-short global equity managers who had the ability to borrow stocks directly from our internal passive pool, sell them short, and then go long other securities. Their benchmark was zero. Overall, we try to minimize the potential gaming of benchmarks.

Question: What is the median IR for equity managers, fixed-income managers, and managers of other asset classes?

Raymond: I don't want to give specific numbers, although I have observed that there's an arbitrage between the manager's belief in his or her IR and what we believe is a sustainable IR in an asset class. We do believe in the fundamental law of active management, meaning the fewer the constraints and the broader the opportunity set, the higher the expected IR should be.

In the first mandates, as a general principle, we tried to construct as broad an opportunity set for the mandates as possible. So, the initial mandates were global equity along with any other skill we believed the manager had. Some did currency overlays, country overlays, country-sector baskets, or whatever made sense in the context of the portfolio. All of the mandates were under one account that we applied the fee structure to because applying it once is preferable to applying a performance fee structure to the separate strategies.

Question: Please expand on the concept of the base fee as a percentage of the active risk target and how that ties to capacity and capacity limits.

Raymond: The base fee is expressed as a proportion of the total active risk target measured in absolute dollar terms. The example I gave was a \$500 million mandate with 4 percent active risk. Multiplying the two together equals \$20 million of active risk. The reason the base fee should be proportional to that total dollar active risk is that if the active risk is doubled, by either doubling the mandate size with the same active risk target or keeping the mandate size constant and doubling the risk, it will lead to a doubling of the active positions or deviations from the benchmark. So, the active risk target is proportional to the capacity usage, which is why the base fee is tied to the active risk target.

Question: What is the best way to sell this fee structure to a board when the performance fee could be quite substantial, particularly as Q goes up?

Raymond: The fact is that paying large performance fees means that more value is being added. It is a logical argument, but seeing its logic does require a change in perspective on the part of board members.

Question: How do you handle situations that require changes to the investment mandates midstream?

Raymond: That is built into the agreements. For example, if the active risk target of the manager is changed from \$20 million of active risk to \$40 million, calculating Q and the IRs means simply adjusting the return series so that they have equal active risk targets. So, double the returns of the first part

of the series or multiply by half the second part to calculate Q and then apply that to the cumulative added value. Similarly, adjust the base fee so that it is proportional to the active risk target. The B part itself is not adjusted, but the base fee goes up and down with the active risk target.

Question: Are your formulas conditional on a positive track record, or do they assume that returns are normally distributed?

Raymond: We just look at the actual historical record and assume it is normally distributed.

Question: What impact does a correlation between luck and skill have?

Raymond: In the way I defined it, there is no correlation between luck and skill.